



HAL
open science

Logics for Social Networks: Asynchronous Announcements in Orthogonal Structures

Saúl Fernández González

► **To cite this version:**

Saúl Fernández González. Logics for Social Networks: Asynchronous Announcements in Orthogonal Structures. Library and information sciences. Université Paul Sabatier - Toulouse III, 2021. English. NNT: 2021TOU30269 . tel-03683355

HAL Id: tel-03683355

<https://theses.hal.science/tel-03683355v1>

Submitted on 31 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE

En vue de l'obtention du
DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE
Délivré par l'Université Toulouse 3 - Paul Sabatier

Présentée et soutenue par
SAUL FERNANDEZ GONZALEZ

Le 15 décembre 2021

**Logiques pour les réseaux sociaux: annonces asynchrones dans
des structures orthogonales**

Ecole doctorale : **EDMITT - Ecole Doctorale Mathématiques, Informatique et
Télécommunications de Toulouse**

Spécialité : **Informatique et Télécommunications**

Unité de recherche :
IRIT : Institut de Recherche en Informatique de Toulouse

Thèse dirigée par
Philippe BALBIANI

Jury

Mme Alessandra PALMIGIANO, Rapporteur
M. Thomas AGOTNES, Rapporteur
M. Hans VAN DITMARSCH, Examineur
M. Philippe BALBIANI, Directeur de thèse
M. Emiliano LORINI, Président

LOGICS FOR SOCIAL NETWORKS

ASYNCHRONOUS ANNOUNCEMENTS IN ORTHOGONAL STRUCTURES



Saúl Fernández González

PhD Thesis

I.R.I.T
Université de Toulouse
October 2021

Contents

1	Introduction	7
2	Preliminaries	11
2.1	Frames	11
2.2	Modal logics	14
2.3	Dynamic Epistemic Logic	17
I	Indexed Frames	24
3	Orthogonal frames	25
3.1	Examples of indexed frames	27
3.2	Orthogonal frames	29
3.3	Orthogonal structures	34
4	Some case studies	42
4.1	Products.	42
4.2	Subset spaces	43
4.3	STIT logic	45
5	Social Epistemic Logic	59
5.1	Axiomatising $\mathcal{L}(@)$ via canonical models	62
5.2	Finite models	69
5.3	Extensions of SEL	71
5.4	Axiomatisation of $\mathcal{L}(@\downarrow)$	73
5.5	Social Epistemic Logic in orthogonal structures	75
II	Asynchronous Announcements	78
6	Asynchronous Announcement Logic	79
6.1	Syntax	83
6.2	Semantics	88
6.3	The logic AA	99
6.4	Equivalence relations	105
6.5	Comparison with Asynchronous Broadcast Logic	111
7	Partially Synchronised Announcements	119

7.1	Syntax and histories	122
7.2	Semantics	127
7.3	The logic PSAL	128
7.4	PSAL generalises AA	131
7.5	PSAL generalises PAL	132
7.6	Common belief	135
8	Quantification over AA	139
8.1	The logic AAA	140
8.2	Expressivity of AAA	145
8.3	Positive formulas	147
8.4	Axiomatisation of AAA	150
8.5	Asynchronous Action Models	154
9	Dynamic Extensions of SEL	159
9.1	Social Epistemic Logic	160
9.2	Updates	161
9.3	Asynchronous reception of messages	170
10	Conclusions	183
10.1	What we did and did not do	183
10.2	Lost causes and future perspectives	185
	Bibliography	193
	Index	202

Acknowledgements

FIRST AND FOREMOST, I wish to thank my thesis supervisor Philippe Balbiani. Thanks to his support and, more importantly, his tolerance towards my often esoteric way to do mathematics (in spite of his staunch Bourbakism) I am a doctor today. Working with Philippe consisted of him enabling, guiding, improving and giving wings to my ideas for a good 95% of the time and aggressive disagreements about the placement of parentheses for the remaining 5%. It has been an absolute pleasure and, despite the bifurcation of our paths, I look forward to rich collaborations in the future.

Another person to whom I owe a lot from these three years is Hans van Ditmarsch, co-author in several of these investigations. I wish to thank Hans for inviting me to visit him in Nancy. I am also grateful for all the discussions, book recommendations and some of my favourite academic anecdotes¹. Above all, I am grateful for his passion for epistemic logic and his collaborative nature, wonderful counterweights to my often detached and isolationist ways.

Other people who discussed with me the topics of this dissertation include Valentin Shehtman (who gave us a preliminary idea of what much later would become the main isomorphism of Chapter 3), Emiliano Lorini (to whom I owe enriching discussions about STIT logics some of whose fruits made it into Chapter 4), Dr. Antonio Yuste Ginel (whose persistent knowledge of modal logic greatly surpasses my I'll-just-look-it-up-for-the-fifteenth-time attitude), and Mina Pedersen (to whom I owe my interest in social network logics that would end up shaping Part I of this thesis²).

I wish to thank the people who constituted the jury of my thesis: Balbiani, Lorini and van Ditmarsch, already mentioned; Thomas Ågotnes and Alessandra Palmigiano, whose comments and careful reading of this text helped greatly improve it and whose questions during the defence watered the seeds of so many ideas.

I would also like to thank all the people who shared an office with me at one point in the last three years: Dr. Maryam Rostamigiv, Dr. Christos Rantsoudis, Cédric Tarbouriech, Quentin Gougeon, Xinghan Liu, and my good friend Dr. Élise Perrotin, who through her stubborn crocheting and tea-making managed to make Office 303 feel like a home.

¹Hans always travels with his violoncello, which usually forces him to buy an extra plane ticket just for the instrument; since airlines insist of collecting the names of all passengers, even wooden inanimate ones, Hans often writes 'Dr. Violoncello' on its ticket.

²Our intention, thwarted by the pandemic but soon to be realised, was to collaborate on a paper on these topics. This frustrated endeavour ended up becoming half of my dissertation.

I extend my gratefulness to the Amsterdam alumni who made my stay at IRIT more pleasant than it would already have been: Grzesiek Lisowski, who hosted me in his house when I interviewed for this job and gave me my first tour of the city, and Dr. Arianna Novaro, who helped me navigate the Dantesque depths of bureaucracy and was always ready to lend a hand.

Of all people who accompanied me in this endeavour, I am most grateful to Rachael Colley – great friend, flatmate and accomplice.

Due to a rather decaffeinated notion of ‘professionalism’ I have thus far decided to only thank academics. Finishing these lines I feel like I must correct myself: I wish to thank all the important people in my life, in particular those who I met in Toulouse and have made these strange and apocalyptic years some of the best of my life: Scherr, Colley, Fuchiwaki, Vené, Yuste Ginel³ – witnesses, enablers, anchors of my sanity.

I give love to people extremely deserving of it. And I think that’s rather nice.



³This is an extremely non-exhaustive survey.

Abstract

This thesis has two main objects of study, closely related to each other.

On the one hand, we provide and study models for asynchronous transmission and reception of messages. To do this, we utilize the framework of Dynamic Epistemic Logic, a branch of Modal Logic which studies the epistemic state of an agent (i.e. what they know) and how this state changes under several circumstances. One of the better known dynamic epistemic logics is Public Announcement Logic [86], a logic which allows for a notion of receiving a message. In a multi-agent system, this message is received by all agents at the same time, and they all know that the others have received it. In the main chapter of this thesis, we provide a framework for asynchronous announcements, in which the agents might receive the message at different times and be uncertain whether others know the information contained within it.

On the other hand, we study a class of relational structures for modal logics which show up quite often in different areas of the literature: this is the class of *orthogonal frames*. Orthogonal frames are bi-relational structures wherein two distinct points cannot be connected by both relations at the same time. We give a sound and complete logic of orthogonal frames under different restrictions, and we provide decidability results. To illustrate the ubiquity of these structures, we provide multiple examples of frameworks for modal logics which are based on orthogonal frames, and we use some of the results obtained earlier to show how one can further the study of these structures by focusing on their orthogonality.

To finish up, we combine the two areas of study, by taking as a case study the orthogonal framework of Social Epistemic Logic [92]. This is a framework for studying the epistemic state of agents in a social network. We provide different dynamic extensions, and in particular we give a way to model the transmission of announcements asynchronously in a social network.



Résumé

Cette thèse a deux objets d'étude principaux.

D'une part, nous proposons et étudions des modèles de transmission et de réception asynchrones de messages. Pour cela, nous nous plaçons dans le cadre des logiques épistémiques dynamiques - un sous-domaine de la logique modale qui formalise les états épistémiques d'un agent (i.e. ce que l'agent sait) et qui caractérise la façon dont ces états évoluent en différentes circonstances. La plus connue des logiques épistémiques dynamiques est la logique des annonces publiques [86] - une logique dynamique qui considère comme action de base l'action d'effectuer une annonce publique. Dans un système multi-agent, il est dans la connaissance commune des agents que les messages sont reçus par tous les agents au même instant. Dans le chapitre principal de la thèse, nous proposons un modèle d'annonces asynchrones dans lequel les agents peuvent recevoir les annonces à différents instants tout en ignorant si les autres agents ont également reçu ces annonces.

D'autre part, nous étudions une classe de structures relationnelles qui apparaissent assez souvent en logique modale : la classe des cadres orthogonaux. Les cadres orthogonaux sont des structures birelationnelles dans lesquelles deux composantes connexes arbitraires déterminées par les deux relations ont au plus un élément en commun. Pour différentes restrictions de la classe des cadres orthogonaux, nous proposons des axiomatisations correctes et complètes des ensembles de formules valides que ces restrictions déterminent et nous proposons quelques résultats de décidabilité de ces ensembles. Pour illustrer l'ubiquité des cadres orthogonaux, nous proposons des exemples de classes de modèles pour les logiques modales qui sont basées sur eux et nous montrons comment les résultats de la thèse peuvent être utilisés pour étudier ces classes du point de vue de leur orthogonalité.

Enfin, nous combinons les deux parties précédentes dans le contexte de la logique épistémique sociale [92]. Il s'agit d'une logique développée pour l'étude des états épistémiques des agents dans un réseau social. Nous proposons différentes extensions dynamiques de cette logique et, en particulier, nous modélisons la transmission d'annonces asynchrones dans un réseau social.



Introduction

THIS THESIS YOU ARE HOLDING is the result of three years of work devoted to two overlapping sets of questions; the results obtained in this interval are presented here in a narratively aesthetic yet unnervingly unchronological order.

Firstly (yet time-wise lastly), this thesis is about the relation between *indexed frames* and *orthogonal structures*. The former consist of bidimensional sets (which is to say Cartesian products) with two relations defined on them which ‘respect’ one of the coordinates, in the following sense: if a pair of points (x, y) is related to a pair of points (a, b) in an indexed frame, it must be the case that either $x = a$ or $y = b$, depending on which of the relations links them. The latter consists of sets equipped with two relations, which have the property of being orthogonal to each other: this means one can never reach a point y from a different point x by moving simultaneously along both relations.

Two key insights about these structures are explored in these pages. The first one being the fact that indexed frames seem to pop up in unexpected places in the Modal Logic literature: a product of two frames [48] is an indexed frame; a subset space [77] can be seen as an indexed frame; most frameworks for the logic of agency called STIT [20] utilise indexed frames; some recent models for knowledge and the transmission thereof within a social network [92] are built upon these frames. Regarding the last example, this framework for epistemic logics in social networks proves to be an excellent case study and is the subject of two chapters of this thesis. Due to their ubiquity, I believe a good argument can be made towards the importance of an independent study of these bidimensional structures.

The second (and perhaps more relevant) insight is that indexed frames and orthogonal structures are one and the same thing (which is a rather unmathematical way to say they are isomorphic to each other). This observation helps push the above argument: on the one hand, this means that indexed frames are even more commonplace than they seem at first glance (every tree with two relations is an indexed frame!); on the other hand, the categorical equivalence between these families of mathematical structures, along with the apparent unremarkability of orthogonal structures as opposed to indexed frames, gives hope that standard, run-of-the-mill techniques in Modal Logic may be used to deal with more unconventional bidimensional frameworks.




Figure 1.1: An indexed frame (left) and an orthogonal structure (right) which are isomorphic to each other; the full and dotted arrows represent different relations.

Lastly (yet since the start of my doctorate), this thesis is about *asynchronous information exchange*. These investigations are contained within the area of *epistemic logic*. Epistemic logic is a broad term which refers to the formal study of notions such as ‘belief’ and ‘knowledge’. Philosophical questions such as ‘what is the difference between correctly believing and knowing?’ [53] or ‘to which extent can a statement about knowledge of a future event be said to be true?’ [64] fall under the shadow of this parasol; so do more formal questions of the sort of ‘how does one mathematically model uncertainty?’ or ‘how could a robot be made aware that a human believes something false?’ [35].

The ‘modal’ flavour of epistemic logic came by the hand of Hintikka [63] in the form of relational structures which represent knowledge situations. Modal epistemic logic has ever since been a fruitful object of study. Starting in the 80’s, an interest to study notions of epistemic change in these mathematical models emerged: how does an agent adjust her beliefs when she receives novel information which contradicts her previous certainties? [2]; how does one incorporate new factual information to one’s knowledge? [86]

The second part of this thesis is thus situated in the realms of *dynamic epistemic logic*. In those pages we point our lens towards multi-agent situations wherein a message is sent to multiple correspondents who will receive the message in their own time (be it sooner, later or not at all) and independently of each other. Investigations on asynchronicity are not a novelty: multiple forms of asynchronous communication have been thoroughly studied in the distributed computing and temporal logic literature [58, 71, among many others]; however, the study of this particular kind of asynchronicity (the kind which results from separating the sending and receiving of a message) from the perspective of dynamic epistemic logic is rather uncharted (albeit

not completely – see [67]).

The framework for asynchronous announcements presented here could be apt to model communicative situations such as a mass sending of emails; the correspondents expect every other correspondent to eventually read the email, but are uncertain about whether they have read it or when they will read it. With a more involved toolset it could also help model the way posts travel in social networks: when one sends a tweet, or posts an update, one only expects a subset of all users to read it (namely, one’s ‘friends’ on the social network), and certainly (barring some obsessive types) not in the moment it was sent.

One of the case studies used in the first part of the thesis makes a comeback towards the end: combining the results on indexed frames obtained in the first half, and more specifically their application to a particular logic of knowledge in social networks [92], the last chapter of this dissertation presents a logic for social networks with asynchronous messages.

The decision to entitle this thesis after the last chapter responds, once again, to narrative criteria: while the most relevant results in this text are not contained in it (with Chapters 3 and 6 being arguably the juiciest), it does interlace the two main lines of investigation into a rather nice metaphorical ribbon.

All ribboned up, let us get started.



This thesis contains the following:

Chapter 1: Introduction (p. 7), in which we briefly introduce the main subjects of the thesis and give an overview of the chapters, namely this very overview, and which contains the only instance of a chapter cross-referencing itself, just three lines of text ago.

Chapter 2: Preliminaries (p. 11), in which a presumed inexperienced reader is showered with some of the mathematical notions necessary to make sense of the rest of this thesis.

PART I: INDEXED FRAMES, consisting of:

Chapter 3: Orthogonal frames (p. 25), in which the notions of indexed frames and orthogonal structures are explored, the categorical equivalence between these structures is highlighted, and the logic of these structures is given.

Chapter 4: Some case studies (p. 42), in which we point out examples of well-known models in the Modal Logic literature which contain underlying indexed frames and study the shape of their isomorphic orthogonal structures.

Chapter 5: Another case study: Social Epistemic Logic (p. 59), in which the aforementioned framework for knowledge in social networks is discussed and some of the techniques developed in prior chapters are applied to it.

PART II: ASYNCHRONOUS ANNOUNCEMENTS, consisting of:

Chapter 6: Asynchronous Announcement Logic (p. 79), the longest chapter, in which the main framework for asynchronous communication is presented, its sound and complete logic is provided, and multiple observations concerning its syntax and semantics are discussed.

Chapter 7: Partially Synchronised Announcement Logic (p. 119), in which a variant of the above framework is presented in detail, allowing for groups of epistemic agents to receive messages together while being aware that others in the group have received the same message, and in which multiple interpretations for ‘group reception’ are considered and axiomatised.

Chapter 8: Quantifying over asynchronous information change (p. 139), in which yet another variant of the Asynchronous Announcements framework is provided, this time allowing to reason about arbitrary sequences of potential future messages, and the knowledge and beliefs which could be gained when such reasoning is allowed.

And to finish it off:

Chapter 9: Some dynamic extensions of Social Epistemic Logic (p. 159), in which several frameworks for information change in social networks are presented in detail, most notably one allowing for an asynchronous spreading of messages, combining some of the results obtained in preceding chapters.

All errors are mine.¹

¹*At the very least*, those contained in this thesis.

Preliminaries

Truth is a matter of the imagination.

Ursula K. LeGuin, *The Left Hand of Darkness* (1969).

THIS CHAPTER PRESENTS TO THE READER some of the basic mathematical notions which will be discussed in this thesis. The reader familiar with modal logics (and epistemic logics in particular) should feel free to skip this chapter.

We point the unfamiliar readers to the excellent reference texts [28, 42] for precise details.

2.1 Frames

Perhaps the most used concepts in this thesis are the following:

Definition 2.1 (Frame). A *relational structure* (or a *Kripke frame* or simply a *frame*) is a tuple

$$(W, R)$$

where W is a nonempty set and $R = \{R_i : i \in I\}$ is a family of binary relations defined on the set W (that is, each R_i is a subset of W^2).

In the case where I is finite, $I = \{1, \dots, n\}$. we may represent the frame as (W, R_1, \dots, R_n) .

An element $w \in W$ will be called a *world* or a *state*.

For the remainder of this thesis, we will let \mathbf{Prop} be a countable set, whose elements are called *propositional variables* and will be generally represented by the letters p, q, \dots

Definition 2.2 (Model). A *model* is a tuple

$$(W, R, V)$$

where (W, R) is a frame and $V : \mathbf{Prop} \rightarrow 2^W$ is a map which assigns some subset of W to each propositional variable; this map is called a *valuation*.

2.1.1 Morphisms

Let $F = (W, \{R_i\}_{i \in I})$ and $F' = (W', \{R'_i\}_{i \in I})$ be two frames with the same number of relations; let V and V' be two valuations on F and F' respectively.

Definition 2.3 (Bounded morphisms). A *bounded morphism* (or simply a *morphism*) between the frames F and F' is a map

$$f : W \rightarrow W'$$

satisfying the following properties for all $i \in I$, $w, v \in W$ and $v' \in W'$:

- (forth) if $R_i w v$, then $R'_i f(w) f(v)$;
- (back) if $R'_i f(w) v'$, then there exists some $v \in W$ such that $f(v) = v'$ and $R_i w v$.

A *bounded morphism* between the models (F, V) and (F', V') is a map $f : W \rightarrow W'$ satisfying the properties above plus

- (atoms) $w \in V(p)$ if and only if $f(w) \in V'(p)$.

An *isomorphism* between F and F' is a bijection $f : W \rightarrow W'$ such that, for all $i \in I$ and $w, v \in W$,

$$R_i w v \text{ if and only if } R'_i f(w) f(v),$$

whereas an isomorphism between models respect as well the (atoms) property.

Note that an isomorphism is a bounded morphism.

Definition 2.4. [Generated subframe] $(W', \{R'_i\}_{i \in I}, V')$ is said to be a *generated subframe* of $(W, \{R_i\}_{i \in I}, V)$ if $W' \subseteq W$, R'_i and V' are the restrictions of R_i and V to W' , and, for all $w' \in W'$ and $w \in W$, if $R_i w' w$, then $w \in W'$.

2.1.2 Categorical equivalences between frames

The concept of *categorical equivalences* is mentioned a number of times in this thesis. We shall not venture into the dark waters of Category Theory within these preliminaries. We will give, however, some necessary notions.

Definition 2.5. For the purposes of this thesis, a *frame category* $\mathbb{C} = (\text{Obj } \mathbb{C}, \text{Ar } \mathbb{C})$ consists of:

- A class $\text{Obj } \mathbb{C}$ of frames (the *objects* of \mathbb{C} , denoted X, Y, \dots), and
- the family $\text{Ar } \mathbb{C}$ of morphisms between frames in $\text{Obj } \mathbb{C}$ (the *arrows* of \mathbb{C} , denoted $X \xrightarrow{f} Y$).

A *functor* between two categories $F : \mathbb{C} \rightarrow \mathbb{D}$ consists of

- A map which assigns to every object $X \in \text{Obj } \mathbb{C}$ an object $FX \in \text{Obj } \mathbb{D}$, and
- a map which assigns to every arrow $X \xrightarrow{f} X'$ in $\text{Ar } \mathbb{C}$ an arrow

$$FX \xrightarrow{Ff} FX'$$

in $\text{Ar } \mathbb{D}$, such that $F(\text{Id}_X) = \text{Id}_{FX}$ and, for all arrows $X \xrightarrow{f} Y$ and $Y \xrightarrow{g} Z$, $F(f \circ g) = Ff \circ Fg$, where Id_X denotes the identity map, and $f \circ g$ denotes the composition of f and g .

A *categorical equivalence* is a pair of functors $F : \mathbb{C} \rightarrow \mathbb{D}$ and $G : \mathbb{D} \rightarrow \mathbb{C}$ such that GFX is isomorphic to X for all $X \in \text{Obj } \mathbb{C}$ and FGY is isomorphic to Y for every object $Y \in \text{Obj } \mathbb{D}$.

A *categorical isomorphism* is a categorical equivalence (F, G) with the extra property that $FGY = Y$ and $GFX = X$, for all $X \in \text{Obj } \mathbb{C}$ and $Y \in \text{Obj } \mathbb{D}$.

The following Lemma will be useful:

Lemma 2.6. *Given two frame categories \mathbb{C} and \mathbb{D} , if every \mathbb{C} -object is isomorphic as a frame to some \mathbb{D} -object, and every \mathbb{D} -object is isomorphic as a frame to some \mathbb{C} -object, then \mathbb{C} and \mathbb{D} are equivalent.*

Proof. Suppose that for every \mathbb{C} -object X , there exists some isomorphic \mathbb{D} -object Y_X , with $\sigma_X : X \rightarrow Y_X$ being the corresponding isomorphism; conversely, suppose as well that for every $Y \in \text{Obj } \mathbb{D}$ there exists some $X_Y \in \text{Obj } \mathbb{C}$ and an isomorphism $\mu_Y : Y \rightarrow Y_X$.

We define a functor $F : \mathbb{C} \rightarrow \mathbb{D}$ by setting:

- $FX = Y$ for $X \in \text{Obj } \mathbb{C}$;
- given a morphism $X \xrightarrow{f} X'$, Ff is the map $Y_X \rightarrow Y_{X'}$ defined by $Ff(y) = \sigma_{X'} f \sigma_X^{-1}(y)$.

Similarly, for \mathbb{D} -objects Y , we set $GY = X_Y$, and

$$G(Y \xrightarrow{g} Y') = (X_Y \xrightarrow{Gg} X_{Y'}),$$

given by $Gg(x) = \mu_{Y'} g \mu_Y^{-1}(x)$.

We leave it to the reader to check that F and G are functors, and we simply point out that by definition FGY must always be isomorphic to Y and GFX must always be isomorphic to X . ■

2.2 Modal logics

2.2.1 Semantics

Definition 2.7 (Basic modal language). The *basic modal language* consists of all *formulas* that can be constructed by using the truth symbol \top , and the propositional letters in \mathbf{Prop} by combining them with the binary operator \wedge or preceding them by the unary operators \neg and \Box , the latter being called a *modal operator* or *modality*. More formally, the construction of a formula ϕ is given by the rule

$$\phi := p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \Box\phi,$$

where $p \in \mathbf{Prop}$.

Other connectives and modal operators can be defined from these:

$$\perp = \neg\top, \phi \vee \psi = \neg(\neg\phi \wedge \neg\psi), \phi \rightarrow \psi = \neg\phi \vee \psi, \Diamond\phi = \neg\Box\neg\phi.$$

The modal operator \Diamond is called the *dual* of \Box .

One could likewise have a *multi-modal* language, with multiple modal operators $\{\Box_i : i \in I\}$. The notation of modal operators in a modal language need not be the boxes above: in this thesis we will deal with modal operators denoted $K, F, B_a, [Agt], \dots$

Definition 2.8 (Relational semantics: satisfaction, validity). Given a model $M = (W, \{R_i\}_{i \in I}, V)$ and a modal language \mathcal{L} given by

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \Box_i\phi,$$

with $p \in \mathbf{Prop}$ and $i \in I$, we define a *satisfaction relation* “ \models ” between worlds $w \in W$ and formulas $\phi \in \mathcal{L}$ recursively as follows:

$$\begin{aligned} M, w \models p & \quad \text{iff } w \in V(p); \\ M, w \models \top & \quad \text{always}; \\ M, w \models \neg\phi & \quad \text{iff } M, w \not\models \phi; \\ M, w \models \phi_1 \wedge \phi_2 & \quad \text{iff } M, w \models \phi_1 \text{ and } M, w \models \phi_2; \\ M, w \models \Box_i\phi & \quad \text{iff } R_i w v \text{ implies } M, v \models \phi. \end{aligned}$$

We read $M, w \models \phi$ as “the formula ϕ is true (or ‘holds’) at world w ”. If there is no risk of ambiguity we omit the model and simply write $w \models \phi$.

A formula ϕ is *valid* in M , denoted $M \models \phi$, if $M, w \models \phi$ for all $w \in W$; we say that ϕ is *valid* in a class \mathcal{C} of models if $M \models \phi$ for all $M \in \mathcal{C}$.

A formula is *valid in a frame* if it is valid in every model derived from appending a valuation to said frame; we define validity in a class of frames accordingly.

The set of all formulas valid in a class of frames \mathcal{C} , denoted $\mathbf{Log} \mathcal{C}$, is called *the logic of \mathcal{C}* .

The following result will be used throughout this thesis:

Proposition 2.9. *Let $M' = (W', R', V')$ and $M = (W, R, V)$ be two models.*

If M' is a generated submodel of M and $w \in W'$, then for every formula ϕ we have: $M, w \models \phi$ iff $M', w \models \phi$.

If $f : M \rightarrow M'$ is a bounded morphism and $f(w) = w'$, then $M, w \models \phi$ iff $M', w' \models \phi$; in particular, if f is surjective, every formula which can be satisfied in M' can also be satisfied in M .

2.2.2 Some modal logics

Definition 2.10 (Normal modal logics). The *minimal normal modal logic K* is the set of formulas that can be obtained by a *proof* (i.e. a finite sequence of formulas which are either axioms or the result of applying a rule to previous elements in the sequence) utilising the following axioms:

- all tautologies in classical propositional logic;
- (K): $\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$;

and the following rules:

- *modus ponens*: from $\phi \rightarrow \psi$ and ϕ , infer ψ ;
- *necessitation*: from ϕ , infer $\Box\phi$.

A *normal modal logic* is a set of formulas L which contains K and is closed under:

- *modus ponens*: if $\phi \rightarrow \psi \in L$ and $\phi \in L$, then $\psi \in L$;
- *necessitation*: if $\phi \in L$, then $\Box\phi \in L$, and
- *uniform substitution*: if $\phi \in L$, then $\phi[p/\theta] \in L$, where $\phi[p/\theta]$ is obtained by uniformly replacing the occurrences of the proposition letter p in ϕ by the formula θ .

Definition 2.11 (Soundness and completeness). Given a logic L and a class of frames \mathcal{C} , we say that:

- L is *sound* with respect to \mathcal{C} if $L \subseteq \text{Log } \mathcal{C}$, and
- L is *complete* with respect to \mathcal{C} if $L \supseteq \text{Log } \mathcal{C}$.

Equivalently, L is sound and complete with respect to \mathcal{C} if every formula that can be proven in L is valid in \mathcal{C} and, conversely, every formula valid in \mathcal{C} can be proven in L .

We now list some of the modal logics that will be mentioned and used throughout the thesis:

Definition 2.12 (Some normal modal logics).

- $S5$ is the least normal modal logic containing the axioms:
 - (T) $\Box\phi \rightarrow \phi$;
 - (4) $\Box\phi \rightarrow \Box\Box\phi$;
 - (5) $\neg\Box\phi \rightarrow \Box\neg\Box\phi$.
- $S4$ is the least normal modal logic containing the axioms (T) and (4).
- $KD45$ is the least normal modal logic containing the axioms (4), (5), and
 - (D) $\Box\phi \rightarrow \neg\Box\neg\phi$.

We refer the reader to [28, Chapter 4] for proofs of the following results:

- Theorem 2.13.**
1. $S5$ is sound and complete with respect to the class of frames (W, R) such that R is an equivalence relation;
 2. $S4$ is sound and complete with respect to the class of reflexive and transitive frames;
 3. $KD45$ is sound and complete with respect to the class of frames (W, R) which are serial (i.e. for all w there exists v such that Rwv), transitive and Euclidean ($Rwv \ \& \ Rwu$ imply Rvu).

We call the corresponding classes of frames $\mathcal{S}5$, $\mathcal{S}4$ and $\mathcal{KD}45$.

A logic has the *finite model property* (FMP) if it is sound and complete with respect to a class of finite models. All the above logics have this property, with these classes being the subclasses of finite models of $\mathcal{S}5$, $\mathcal{S}4$ and $\mathcal{KD}45$ respectively.

If a logic has the FMP and admits a finitary axiomatisation, then the problem of membership of a formula to this logic is decidable; this is the case for the above logics.

2.3 Dynamic Epistemic Logic

Most of the logics depicted in this thesis are *epistemic logics*, and many of them are *dynamic*. Let us see what these words mean.

2.3.1 ‘Epistemic’

As the name suggests, these are logics which deal with what an agent *believes* or *knows*.

Using the relational semantics outlined above to describe knowledge is an idea pioneered by Hintikka [63]. The intuition is as follows: the elements $w \in W$ represent *possible worlds*, i.e., different potential states of affairs. A link between two of these worlds represents indistinguishability, i.e., an agent’s inability to tell apart two worlds given the information she has at her disposal.

If an epistemic agent does not know whether, say, polar bears eat penguins (let us represent this by a propositional letter p), this will be modelled by having two worlds in our model, w and v , with $w \models p$ and $v \models \neg p$ and Rwv : the agent cannot distinguish between a world in which polar bears do eat penguins and a world in which they do not. If bears eat penguins in all of an agent’s epistemically accessible worlds (i.e., if $v \models p$ for all v such that Rwv), then we say the agent *knows* this proposition.

If we are talking about a logic of knowledge, we use the basic modal language but we use K as our modal box; in a multi-agent situation, where A is a set of agents, we use modalities $\{K_a\}_{a \in A}$; for belief, we use the notations B and B_a for the modal boxes. In the example above, the formula Kp stands for ‘the epistemic agent knows proposition p ’.

Models for knowledge are usually taken to be in the class $\mathcal{S5}$, for the three axioms of $\mathcal{S5}$ correspond to somewhat desirable properties of knowledge: knowledge is factual (an agent can only know true things, represented by axiom (T) $K\phi \rightarrow \phi$), positively introspective (an agent knows the things she knows, represented by (4) $K\phi \rightarrow KK\phi$), and negatively introspective (an agent knows what are the things she does not know, represented by (5) $\neg K\phi \rightarrow K\neg K\phi$).

Similarly, belief can desirably follow the axioms of $\mathcal{KD45}$ and is usually modelled on serial, transitive and Euclidean models: axiom (D) $B\phi \rightarrow \neg B\neg\phi$ ensures an agent does not believe a contradiction; axioms (4) and (5) ensure that if agent (does not) believe something, then she believes that she (does not) believe it.

Some frameworks in the literature drop the negative introspection requirement for knowledge (for instance Hintikka himself rejects this principle [63]) and consider $\mathcal{S4}$ models; some model knowledge using logics in between $\mathcal{S4}$ and

S5 [80, 97]; some present both knowledge and belief modalities and include axioms that describe their interaction [97, 98].

A few of the above frameworks model knowledge and belief on topological spaces instead of relational models; in this text, we stick to the relational semantics.

2.3.2 ‘Dynamic’.

The word ‘dynamic’ in *dynamic epistemic logic* refers to the notion of *change*. Pioneered by Plaza [86] and Gerbrandy and Groeneveld [52] (and inspired by works on information change in linguistics [55] and belief revision [2]), dynamic epistemic logics model how an agent’s epistemic state varies based on things such as the receiving of new information [86], the refinement of her evidence [34], or the formation of new beliefs which contradict her prior beliefs [2].

The following are some of the logics that will be mentioned in this thesis:

Public Announcement Logic (PAL). Given some finite set of epistemic agents A , the language of PAL [86] is built as follows:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid K_a\phi \mid [\phi]\phi,$$

with $p \in \text{Prop}$ and $a \in A$. The dual operators of K and $[\phi]$ are defined as $\hat{K}\phi = \neg K\neg\phi$ and $\langle\phi\rangle\psi = \neg[\phi]\neg\psi$.

The *dynamic* modality $[\phi]$ represents a *public announcement*: this is some information given by the environment and received by all agents simultaneously.

Within this language, for instance, we can represent the following proposition: ‘if Amanda knows that either p or q are true, and she receives the information that p is false, then afterwards she knows q ’, via the formula

$$K_a(p \vee q) \rightarrow [\neg p]K_aq.$$

These formulas are interpreted on $\mathcal{S5}$ models $(W, \{R_a\}_{a \in A}, V)$ as in Def. 2.8, that is:

$$\begin{aligned} M, w, \models p & \quad \text{iff } w \in V(p), \\ & \quad \vdots \\ M, w \models K_a\phi & \quad \text{iff } R_avv \text{ implies } M, v \models \phi, \end{aligned}$$

with the following extra clause for the ‘announcement’ operator:

$$M, w \models [\phi]\psi \quad \text{iff } M, w \models \phi \text{ implies } M^\phi, w \models \psi.$$




Figure 2.1: *Left*: epistemic model representing a two-agent situation where agent a knows that $\neg p \vee q$ is the case but does not know whether p or whether q . The actual world, w , is represented by the grey node. Transitive epistemic links are omitted.

Right: after the formula p is announced, the worlds v and s , wherein $\neg p$ is the case, disappear from the updated model. The agent is left knowing both p and q .

In the last line, the model M^ϕ is the restriction of M to only those worlds in which ϕ is true, i.e.,

$$M^\phi = (W^\phi, \{R_a^\phi\}_{a \in A}, V^\phi), \text{ where } W^\phi = \{v \in W : M, v \models \phi\},$$

and R_a^ϕ, V^ϕ are the corresponding restrictions. This represents the fact that all agents in the system, after receiving the information that ϕ , have ‘eliminated’ all the $\neg\phi$ -worlds from their pool of conceivable possible worlds, and thus have incorporated ϕ to their knowledge.

We represent the above example in Figure 2.1.

The sound and complete logic of PAL consists of:

- the S5 axioms and rules for each of the K_a modalities;
- the following *reduction axioms*:

$$\begin{array}{ll}
 [\phi]\top \leftrightarrow \top; & [\phi]p \leftrightarrow (\phi \rightarrow p); \\
 [\phi]\neg\psi \leftrightarrow (\phi \rightarrow \neg[\phi]\psi); & [\phi](\psi \wedge \chi) \leftrightarrow ([\phi]\psi \wedge [\phi]\chi); \\
 [\phi]K_a\psi \leftrightarrow (\phi \rightarrow K_a[\phi]\psi); & [\phi][\psi]\chi \leftrightarrow [\phi \wedge [\phi]\psi]\chi.
 \end{array}$$

By applying the above reductions recursively, one can depart from a formula in the language of PAL and obtain an equivalent formula in the basic

(announcement-free) modal language. The completeness of PAL is thus a corollary of the completeness of (multi-agent) S5 [59]; likewise for the Finite Model Property. PAL is decidable [42].

Arbitrary Public Announcement Logic (APAL). The language of APAL [7] adds an extra modal operator to the language of PAL:

$$[!] \phi,$$

which is read ‘after any arbitrary announcement, ϕ holds’. This can be used to represent knowability, in situations such as: ‘there is some information that Amanda could receive after which she would know ϕ ’, $\langle ! \rangle K_a \phi$.

Sentences of this language are interpreted on S5 models; the new modality is interpreted as follows:

$$M, w \models [!] \psi \quad \text{iff} \quad M, w \models [\phi] \psi$$

for every formula ϕ in the language of PAL.

The logic APAL has an infinitary axiomatisation [7]¹ and is undecidable: there is no reduction axiom for the ‘arbitrary announcement’ modality [47].

APAL has spawned several variants, such as Group Announcement Logic [1], ‘APAL with memory’ [19], or Arbitrary Arrow Update Logic [43].


Action Model Logic. An *action model* [17] is a tuple

$$\mathcal{E} = (E, S, \text{pre})$$

where E is a finite domain whose elements are called *actions* and $S = \{S_a : a \in A\}$ is a family of binary relations on E indexed by the finite set of agents A , whereas pre assigns a formula ϕ in the language to each $e \in E$: $\text{pre}(e)$ is called the *precondition* of e .


These structures can be used to describe epistemic actions which are more complex than a public announcement. A simple example of such an action from [42] is the following: agents Amanda and Boris are unsure of the truth of a statement p ; in front of Boris, Amanda receives an envelope with the information as to whether p or $\neg p$ is the case, which she opens and reads without letting Boris see. The initial situation of uncertainty can be represented by a model of the shape:

¹The proof of completeness in that text contains an incorrect demonstration of the *Truth Lemma*; a correction was published in [6].



(where the grey node represents the ‘actual world’ and the letters inside the nodes represent the valuation of p at that world).

After the above epistemic action takes place, we would want a situation in which (i) Amanda knows the truth about whether p or $\neg p$ (let us say that p is true), (ii) Boris is still uncertain about p , but (iii) Boris is aware that Amanda knows whether p , i.e.:



Action models have a dual existence as both semantic and syntactic objects which allows for these complex interactions.

Semantically, we can use an action model to update a model as follows:


Definition 2.14. Given a model $M = (W, R, V)$ and an action model $\mathcal{E} = (E, S, \text{pre})$, we define the update $M \otimes \mathcal{E}$ as the model $M \otimes E = (W', R', V')$, where:

- i. W' consists of the (world,action) pairs such that the world satisfies the precondition of the action, i.e.

$$W' = \{(w, e) \in W \times E : M, w \models \text{pre } e\};$$

- ii. $(w, e)R'_a(v, f)$ iff $R_a wv$ and $S_a e f$;
- iii. $(w, e) \in V'(p)$ iff $w \in V(p)$.

In our previous example, we may update the initial model with the action model



to obtain the desired final model. In the drawing above, the squares represent the different actions, the grey square is the ‘actual action’ taking place (corresponding to Amanda learning the fact that p) and the formulas inside the squares represent preconditions.

One can also represent a public announcement using action models. We leave it to the reader to check that, if we update a model M with the singleton action model $\mathcal{E}_\phi = (E, S, \text{pre})$, where $E = \{*\}$, $R_a = \{(*, *)\}$ for all a , and $\text{pre}(*) = \phi$, the resulting model $M \otimes \mathcal{E}_\phi$ is precisely (or rather: isomorphic to) the restriction-after-announcement M^ϕ defined above for PAL.

Syntactically, noting that action models are finite, and thus for a finite set of agents the set of possible action models is countable up to isomorphism, we can consider countably many modal operators $[\mathcal{E}, e]\phi$ for each action model \mathcal{E} and for each $e \in E$ and add them to the language of epistemic logic to obtain the language of Action Model Logic.

This new modal operator is interpreted as follows:

$$M, w \models [\mathcal{E}, e]\phi \text{ iff } M, w \models \text{pre}(e) \text{ implies } M \otimes \mathcal{E}, (w, e) \models \phi.$$

Similarly to PAL, there is a sound and complete axiomatisation of Action Model Logic which employs reduction axioms; see [18] for details.

Arrow Update Logic. *Arrow Update Logic* [68] offers a different way to update epistemic models, less powerful than action models yet more so than public announcements. These ‘arrow updates’ do not alter the domain of the model, but rather constrain which of the links between the different worlds will remain after the update has taken place.

This can be quite useful for logics of knowledge and belief, for instance by making some worlds inaccessible to an agent, but also for modal logics outside of the epistemic umbrella; for instance, in Chapter 9 we will use updates of this type to model the erasure of a ‘social link’ (e.g. a friendship) between two people.

We do not provide many details here; instead, we refer the reader to [68, 37]. We do provide the following definitions:

Given a finite set of agents A , an *arrow update* is a finite set

$$U = \{(\phi_1, a_1, \psi_1), \dots, (\phi_n, a_n, \psi_n)\},$$

where the ϕ_i and ψ_i ’s are formulas and the a_i ’s are agents. This represents the fact that, after the update, agent a_i will only maintain epistemic links that go from worlds satisfying ϕ_i to worlds satisfying ψ_i .

For each such set, an update operation on models is defined:

$$(W, \{R_a\}_{a \in A}, V) \otimes U = (W, \{R_a^U\}_{a \in A}, V),$$

with

$$R_a^U = \{(u, v) \in R_a : \exists(\phi, a, \psi) \in U \text{ s.t. } M, u \models \phi \ \& \ M, v \models \psi\},$$

and a modal operator $[U]\phi$ is added to the multi-agent epistemic language, which is interpreted on epistemic models as follows:

$$M, w \models [U]\phi \text{ iff } M \otimes U, w \models \phi.$$

Again, this logic has a complete axiomatisation (via reduction axioms) and is decidable; see [68].

With these technical preliminaries out of the way, let us move on to the results of this dissertation.

Part I

INDEXED FRAMES

Orthogonal frames

We define and study the notion of an *indexed frame*. This is a bi-dimensional structure consisting of a Cartesian product equipped with relations which only relate pairs if they coincide in one of their components. We show that these structures are quite ubiquitous in modal logic, showing up in the literature as products of frames, subset spaces, or temporal frames for STIT logics. We show that indexed frames are completely characterised by their ‘orthogonal’ relations, and we provide their sound and complete logic.¹



THIS FIRST PART of the thesis is concerned with a certain type of bi-dimensional relational structure which shows up in multiple areas of modal logic. The ubiquity of these structures, I wish to argue, should motivate an independent study of their properties and their logic, the first steps towards which are taken in the present chapter.

In the text we shall call these structures *frames!**indexed frames*. Let us start off by providing two distinct (but ultimately equivalent) definitions of what we mean by that.

Definition 3.1. By *indexed frame* we refer indistinctly to any of the following structures:

- (IF1) Frames $(W_1 \times W_2, R_1, R_2)$ where R_1 and R_2 are binary relations on $W_1 \times W_2$ such that, for all $w_1, w'_1 \in W_1$ and $w_2, w'_2 \in W_2$,
 - $(w_1, w_2)R_1(w'_1, w'_2)$ implies $w_2 = w'_2$, and
 - $(w_1, w_2)R_2(w'_1, w'_2)$ implies $w_1 = w'_1$.
- (IF2) Tuples (W_1, W_2, R^1, R^2) where $R^1 = \{R_w^1 : w \in W_2\}$ is a family of binary relations on W_1 indexed by the elements of W_2 , and $R^2 = \{R_w^2 : w \in W_1\}$ is a family of binary relations on W_2 indexed by the elements of W_1 .

¹This chapter is based on the first four sections of the paper

Philippe Balbiani and Saúl Fernández González. Orthogonal frames and indexed relations. In *Logic, Language, Information, and Computation, 27th International Workshop WoLLIC 2021*, pages 219–234. Springer, 2021.

It is quite straightforward to see how these two definitions refer to the same type of structure²:

Given a frame of the form (IF1), we define families of relations R^1 and R^2 as follows: for every $w \in W_2$:

$$w_1 R_w^1 w'_1 \text{ iff } (w_1, w) R_1(w'_1, w)$$

and for every $w \in W_1$:

$$w_2 R_w^2 w'_2 \text{ iff } (w, w_2) R_2(w, w'_2).$$

With this, we obtain a frame of the form (IF2). Conversely, given a frame in the form (IF2) we obtain a (IF1) frame by defining relations R_1 and R_2 on $W_1 \times W_2$ as follows:

$$(w_1, w_2) R_1(w'_1, w'_2) \text{ iff } w_2 = w'_2 \text{ and } w_1 R_{w_2}^1 w'_1$$

and

$$(w_1, w_2) R_2(w'_1, w'_2) \text{ iff } w_1 = w'_1 \text{ and } w_2 R_{w_1}^2 w'_2.$$

Having these bi-dimensional structures at hand, one can interpret formulas over a bi-modal language

$$\phi ::= p \mid \perp \mid (\phi \wedge \phi) \mid \neg \phi \mid \Box_1 \phi \mid \Box_2 \phi$$

with respect to pairs $(w_1, w_2) \in W_1 \times W_2$ as follows:

$$\begin{aligned} \text{(IF1)} \quad (w_1, w_2) \models \Box_1 \phi & \text{ iff } (w_1, w_2) R_1(w'_1, w_2) \text{ implies } (w'_1, w_2) \models \phi; \\ (w_1, w_2) \models \Box_2 \phi & \text{ iff } (w_1, w_2) R_2(w_1, w'_2) \text{ implies } (w_1, w'_2) \models \phi; \end{aligned}$$

$$\begin{aligned} \text{(IF2)} \quad (w_1, w_2) \models \Box_1 \phi & \text{ iff } w_1 R_{w_2}^1 w'_1 \text{ implies } (w'_1, w_2) \models \phi; \\ (w_1, w_2) \models \Box_2 \phi & \text{ iff } w_2 R_{w_1}^2 w'_2 \text{ implies } (w_1, w'_2) \models \phi. \end{aligned}$$

It is easy to see how these semantics are equivalent via the above transformations.

We start this chapter by illustrating that indexed frames show up quite often in the literature. In order to put forward this argument, we provide in Section 3.1 examples of well-known models in different areas of modal logic which are indexed frames. The examples in the first section will be recovered as the subject of the next two chapters. In Section 3.2 we show that the property of ‘orthogonality’ (i.e., the fact that each point in the model is uniquely determined by the pair of connected components to which it belongs) is necessary and sufficient to characterize indexed frames, and we use this property

²Or, if one prefers, to two categorically isomorphic classes of structures.

to provide their sound and complete logic. In Section 3.3 we enrich our language with modalities \blacksquare_1 and \blacksquare_2 which fix w_2 (resp. w_1) and quantify over all points in W_1 (resp. W_2). We provide a sound and complete logic for this extended language.

3.1 Examples of indexed frames

Let us see some well-known structures that are either indexed frames or generated subframes thereof. We will use the term ‘indexed relation’ to informally refer to a relation defined on a Cartesian product which ‘respects’ one of the coordinates.³ We will slightly abuse our definition to include frames with more than one ‘indexed relation’ for each component.

Example 3.2 (Products). [48, Chapter 3] The *product* of two Kripke frames (W_1, R_1) and (W_2, R_2) is the frame

$$(W_1, R_1) \times (W_2, R_2) = (W_1 \times W_2, R_H, R_V),$$

where the *horizontal* and *vertical* relations R_H and R_V are defined as follows:

$$(w_1, w_2)R_H(w'_1, w'_2) \quad \text{iff } w_2 = w'_2 \text{ and } w_1 R_1 w'_1, \text{ and}$$

$$(w_1, w_2)R_V(w'_1, w'_2) \quad \text{iff } w_1 = w'_1 \text{ and } w_2 R_2 w'_2$$

One can see that products very closely adjust to the (IF1) definition; in fact, indexed frames can be seen as a generalization of products.

Indeed, a product could be defined as an (IF2) indexed frame (W_1, W_2, R^1, R^2) with the extra property that, for all $w_1, w'_1 \in W_1$ and $w_2, w'_2 \in W_2$, $R^1_{w_2} = R^1_{w'_2}$ and $R^2_{w_1} = R^2_{w'_1}$.




Figure 3.1: Example of a product of two frames; the dotted and filled arrows represent different relations.

³In the sense that it only connects points which share the i -th coordinate, for fixed $i \in \{1, 2\}$.

Example 3.3 (Subset spaces). In its most basic form [77], a *subset space* is a tuple consisting of a nonempty set X and some collection \mathcal{O} of nonempty subsets of X .

One can define a valuation $V : \text{Prop} \rightarrow 2^X$ and interpret formulas of a bimodal language including \Box and K modalities on a subset space with respect to a pair (x, U) such that $x \in U$ and $U \in \mathcal{O}$ as follows:

$$\begin{aligned} x, U \models p & \quad \text{iff } x \in V(p); \\ & \quad \vdots \\ x, U \models K\phi & \quad \text{iff } y, U \models \phi \text{ for all } y \in U \\ x, U \models \Box\phi & \quad \text{iff } x, V \models \phi \text{ for all } V \subseteq U \text{ such that } x \in V \ \& \ V \in \mathcal{O}. \end{aligned}$$

The semantics above naturally defines two indexed relations on the graph $\mathcal{O}_X := \{(x, U) : x \in U \ \& \ U \in \mathcal{O}\}$, namely:

$$\begin{aligned} (x, U) \equiv_K (y, V) & \quad \text{iff } U = V; \\ (x, U) \geq_{\Box} (y, V) & \quad \text{iff } x = y \text{ and } U \supseteq V \end{aligned}$$

Clearly, the standard Kripke semantics on the frame $(\mathcal{O}_X, \equiv_K, \geq_{\Box})$ (let us call this a *subset space frame*) are the exact semantics above, and moreover this subset space frame is (a generated subframe of) an indexed frame.

Example 3.4 (Social Epistemic Logic). *Social Epistemic Logic* (SEL) is a multi-modal framework to model knowledge within social networks, introduced in [92]. Its language contains, in addition to propositional variables p, q, \dots , *nominal variables* n, m, \dots , an artefact borrowed from Hybrid Logic [4]. It has operators $K\phi$ and $F\phi$ to express things such as ‘‘I know ϕ ’’ and ‘‘all my friends ϕ ’’, and, in addition, it presents an operator $@_n\phi$ for each nominal n to express ‘‘ ϕ is true of the agent named by n ’’.

The models for SEL are of the form $(W, A, \{\sim_a\}_{a \in A}, \{\succ_w\}_{w \in W}, V)$, where each \sim_a is an ‘epistemic indistinguishability’ equivalence relation for agent a on the set of possible worlds W , and each \succ_w is a ‘social’ symmetric and irreflexive relation, representing which pairs of agents in the set A are ‘friends’ at world w . The valuation V assigns subsets of $W \times A$ to propositional variables p and, for a nominal n , $V(n)$ is of the form $W \times \{a\}$ for some a ; it is then said that ‘‘ n is the name of a ’’, denoted $a = \underline{n}_V$.

For the semantics, we read formulas with respect to a pair of a world and an agent as follows:

$$\begin{aligned} (w, a) \models K\phi & \quad \text{iff } (v, a) \models \phi \text{ for all } v \text{ such that } w \sim_a v; \\ (w, a) \models F\phi & \quad \text{iff } (w, b) \models \phi \text{ for all } b \text{ such that } a \succ_w b; \\ (w, a) \models @_n\phi & \quad \text{iff } (w, \underline{n}_V) \models \phi. \end{aligned}$$

$(W, A, \{\sim_a\}_{a \in A}, \{\succ_w\}_{w \in W})$ is clearly an (IF2) indexed frame and even the $@_n$ modality can be interpreted via the ‘‘indexed’’ relation: $(w, a)R_n(v, b)$ iff $w = v$ and $b = \underline{n}_V$.

Its equivalent (IF1) form is $(W \times A, \sim, \succ)$, where $(w, a) \sim (v, b)$ iff $a = b$ and $w \sim_a v$, and $(w, a) \succ (v, b)$ iff $w = v$ and $a \succ_w b$.

Example 3.5 (STIT logic). The logic of *seeing-to-it-that* or *STIT* was first studied in a series of papers culminating in [20] and has shown up in the literature with many variations; in most cases, the different models for STIT are quite explicitly indexed frames or present indexed relations. The one we showcase here is (a slightly simplified version of) a *Kamp frame*, discussed in [33, 101].

A *Kamp frame* is a tuple $(W, \mathcal{O}, \{\sim_t\}_{t \in T}, \{\sim_{t,i}\}_{t \in T, i \in \text{Agt}})$, where each ‘world’ $w \in W$ has a ‘timeline’ associated to it, this being a linear order $\mathcal{O}(w) = (T_w, <_w)$. T is the union of all the T_w ’s. For each $t \in T$, the relations \sim_t and $\sim_{t,i}$ are equivalence relations defined on the set $\{w : t \in T_w\}$. A *Kamp model* is a Kamp frame along with a valuation $V : \text{Prop} \rightarrow 2^T$.

Sentences in a language including a necessity operator \Box , agency operators $[i]$ for $i \in \text{Agt}$ and a temporal operator G are read with respect to pairs (t, w) such that $t \in T_w$ as follows:

$$\begin{aligned} (t, w) \models \Box\phi & \text{ iff } (t, w') \models \phi \text{ for all } w' \sim_t w; \\ (t, w) \models [i]\phi & \text{ iff } (t, w') \models \phi \text{ for all } w' \sim_{t,i} w; \\ (t, w) \models G\phi & \text{ iff } (t', w) \models \phi \text{ for all } t' >_w t. \end{aligned}$$

While this does not exactly adjust to the definitions of indexed frame above, one sees how this structure can be defined as (a generated subframe of) the (IF2) indexed frame

$$(W, T, \{\sim_t\}_{t \in T}, \{\sim_{t,i}\}_{t \in T, i \in \text{Agt}}, \{<_w\}_{w \in W}).$$

(We are slightly bending our definition of ‘indexed frame’ here and allowing for multiple families of relations indexed by the elements of T .)

We can easily ‘rewrite’ these relations to be defined on (a subset of) $W \times T$ in the (IF1) way:

$$\begin{aligned} (t, w) \equiv_{\Box} (t', w') & \text{ iff } t = t' \ \& \ w \sim_t w' \\ (t, w) \equiv_i (t', w') & \text{ iff } t = t' \ \& \ w \sim_{t,i} w' \\ (t, w) <_G (t', w') & \text{ iff } w = w' \ \& \ t <_w t'. \end{aligned}$$

All the frames showcased in this section share one property: namely that of *orthogonality*. This property is explained and studied in the next section.

3.2 Orthogonal frames

The relations R_1 and R_2 in an indexed frame $(W_1 \times W_2, R_1, R_2)$ are ‘orthogonal’ to each other, in the sense that there cannot be two distinct points

connected by both R_1 and R_2 . Indeed, if there is an R_i path from (w_1, w_2) to (w'_1, w'_2) (i.e. if they belong to the same R_i -connected component⁴), then $w_j = w'_j$ for $j \neq i$ and, in consequence, if there are both R_1 paths and R_2 paths between these pairs, then $(w_1, w_2) = (w'_1, w'_2)$. In the present section we shall see that this property fully characterises indexed frames.

For the remainder of this text, given a relation R , we let R^* denote the least equivalence relation containing R , i.e., the equivalence relation induced by the connected components of R .

Definition 3.6 (Orthogonal Frames). A birelational frame (W, R_1, R_2) is said to be *orthogonal* if there exist equivalence relations \equiv_1 and \equiv_2 on W satisfying

$$\text{FO1. } R_i \subseteq \equiv_i \text{ for } i = 1, 2;$$

$$\text{FO2. } \equiv_1 \cap \equiv_2 = Id_W,$$

where Id_W is the identity relation on W .

An orthogonal frame (W, R_1, R_2) is said to be *fully orthogonal* if these relations also satisfy

$$\text{FO3. } \equiv_1 \circ \equiv_2 = W^2.$$

We note the following:

Lemma 3.7. *A frame (W, R_1, R_2) is orthogonal if and only if it satisfies*

$$\text{O1. } R_1^* \cap R_2^* = Id_W.$$

Proof. From right to left, it suffices to take $\equiv_i = R_i^*$; from left to right, if \equiv_1 and \equiv_2 satisfy FO1 and FO2, then $Id_W \subseteq R_1^* \cap R_2^* \subseteq \equiv_1 \cap \equiv_2 = Id_W$. ■

Note as well that, if such a pair of equivalence relations exists, it is not necessarily unique: consider the frame (W, R_1, R_2) where $W = \mathbb{N}^2$ and $R_1 = R_2 = Id_W$; the pair of equivalence relations (\equiv_1, \equiv_2) , where $(n_1, n_2) \equiv_i (m_1, m_2)$ iff $n_i = m_i$ satisfies properties FO1 – FO3; however, the pair (W^2, Id_W) does as well.

Proposition 3.8. *(W, R_1, R_2) is isomorphic to an indexed frame if and only if it is fully orthogonal.*

Proof. Let $(W_1 \times W_2, R_1, R_2)$ be an indexed frame. Then the relations \equiv_1 and \equiv_2 , defined as

$$(w_1, w_2) \equiv_i (w'_1, w'_2) \text{ iff } w_j = w'_j \text{ (where } \{i, j\} = \{1, 2\}),$$

⁴Given a frame (W, R) , a set $C \subseteq W$ is *connected* if, for any two points $w_0, w_n \in C$, there is a sequence $(w_0, w_1, \dots, w_{n-1}, w_n)$ such that Rw_iw_{i+1} or $Rw_{i+1}w_i$ for all i ; a *connected component* is a maximal connected set. The connected components are exactly the equivalence classes of the least equivalence relation containing R .

satisfy FO1, FO2, and FO3.

Conversely, suppose such relations exist and let $[w]_i$ denote the equivalence class of w under \equiv_i . By FO2 and FO3, given any two elements $w, v \in W$, there is exactly one element in the intersection $[w]_2 \cap [v]_1$: let x_{wv} denote this unique element. Consider the frame $(W/\equiv_2 \times W/\equiv_1, R_1, R_2)$, where, for $i = 1, 2$,

$$([w]_2, [v]_1)R_i([w']_2, [v']_1) \text{ if and only if } x_{wv}R_ix_{w'v'}.$$

This is an indexed frame, for if $x_{wv}R_1x_{w'v'}$, then $x_{wv} \equiv_1 x_{w'v'}$, and since $v \equiv_1 x_{wv}$ and $x_{w'v'} \equiv_1 v'$, this gives $[v]_1 = [v']_1$. We reason analogously for R_2 . It is isomorphic to (W, R_1, R_2) via the map $f([w]_2, [v]_1) = x_{wv}$. For injectivity, note that if $x_{wv} \neq x_{w'v'}$, then either $w \not\equiv_2 w'$ or $v \not\equiv_1 v'$. For surjectivity, note that $w = x_{ww}$ for all $w \in W$. Finally, note that we have defined the map in such a way that $([w]_2, [v]_1)R_i([w']_2, [v']_1)$ iff $f([w]_2, [v]_1)R_if([w']_2, [v']_1)$. ■

For the next two definitions, let L_1 and L_2 be two normal modal logics, each with one modal operator, \Box_1 and \Box_2 respectively.

Definition 3.9. A frame (W, R_1, R_2) is a $[L_1, L_2]$ -frame if $(W, R_i) \models L_i$, for $i = 1, 2$ (i.e., if $(W, R_i) \models \phi$ for all $\phi \in L_i$).

Definition 3.10. The *fusion logic* $L_1 \oplus L_2$ is the least normal modal logic containing L_1 and L_2 .

A well-known result is the following:

Theorem [48, Thm. 4.1]. $L_1 \oplus L_2$ is the logic of $[L_1, L_2]$ -frames.

We have:

Proposition 3.11. An orthogonal $[L_1, L_2]$ -frame (W, R_1, R_2) is a generated subframe of a fully orthogonal $[L_1, L_2]$ -frame.

Proof. Given an orthogonal $[L_1, L_2]$ -frame (W, R_1, R_2) , let $[w, v]$ be an abbreviation for the pair $\langle R_2^*(w), R_1^*(v) \rangle$ (where $R_i^*(x)$ represents the equivalence class of x under R_i^* , i.e., the R_i -connected component to which x belongs). We extend W to the set

$$W' = W \cup \{x_{[w,v]} : w, v \in W, R_2^*(w) \cap R_1^*(v) = \emptyset\},$$

i.e., we add one element for each pair of connected components which have an empty intersection, and we extend the relations R_i as follows:

- if $F_\bullet \models L_i$, then $R'_i = R_i$;




Figure 3.2: Example of a fully orthogonal frame (top left) and construction of its isomorphic indexed frame (bottom right). The horizontal dotted lines represent R_1 , whose corresponding equivalence classes, h_1, h_2 and h_3 are represented on the top right. Note that the dotted arrows (R_1) are contained in the dotted boxes (\equiv_1). Similarly for the relation R_2 (vertical, dashed lines) and the corresponding equivalence classes of \equiv_2 (v_1, v_2, v_3 , represented by the dashed boxes on the bottom left). Note that the intersection $h_i \cap v_j$ is always a singleton. We thus make every pair (h_i, v_j) correspond to one point of the original frame, and we draw the relations accordingly (bottom right).

- if $F_\circ \models L_i$, then $R'_i = R_i \cup \{(x, x) : x \in W' \setminus W\}$;

where F_\bullet is the irreflexive singleton frame $(\{*\}, \emptyset)$, and F_\circ is the reflexive singleton frame $(\{*\}, \{(*, *)\})$. (The reader might recall that every normal modal logic is valid in either F_\bullet or F_\circ ; this is a consequence of a classical result by Makinson [75].)

Note that, in either case, no elements of W are related to any elements of $W' \setminus W$ and thus (W, R_1, R_2) is a generated subframe of (W', R'_1, R'_2) .

We define

$$\begin{aligned} \equiv'_1 &= (R_1 \cup \{(v, x_{[w,v]}) : v \in W\})^*, \text{ and} \\ \equiv'_2 &= (R_2 \cup \{(w, x_{[w,v]}) : w \in W\})^*. \end{aligned}$$

Note that \equiv'_1 and \equiv'_2 satisfy conditions FO1 – FO3 of Def. 3.6, and therefore (W', R'_1, R'_2) is a fully orthogonal frame. Finally, for $i = 1, 2$, (W', R'_i) is the disjoint union of the L_i -frame (W, R_i) with a family of singleton L_i -frames, and thus it is an L_i -frame. \blacksquare

Proposition 3.12. *The fusion logic $L_1 \oplus L_2$ is the logic of orthogonal $[L_1, L_2]$ -frames.*

Proof. The proof in [48, Thm. 4.1] of the fact that

the logic of frames (W, R_1, R_2) such that $(W, R_i) \models L_i$ for $i = 1, 2$
is the fusion $L_1 \oplus L_2$

utilises the construction of a *dovetailed frame* in order to prove that any formula ϕ which is *consistent* in $L_1 \oplus L_2$ (in the sense that $\neg\phi \notin L_1 \oplus L_2$) is satisfiable in an $[L_1, L_2]$ -frame. It is a recursive process done as follows: first, one obtains a consistent formula in the language of L_1 by rewriting $\diamond_2\phi$ with ‘surrogate’ propositional variables $p^1_{\diamond_2\psi_1}, \dots, p^1_{\diamond_2\psi_n}$ in place of its maximal subformulas preceded by \diamond_2 . Then one constructs a rooted L_1 -frame satisfying ϕ . Whenever a point in this frame satisfies a surrogate variable $p^1_{\diamond_2\psi}$, one rewrites ψ in the language of L_2 by using surrogates $q^2_{\diamond_1\theta_1}, \dots, q^2_{\diamond_1\theta_m}$ and makes this point the root of an L_2 -frame satisfying this formula. One repeats this process, alternating between L_1 -formulas and L_2 -formulas until one obtains a rooted frame satisfying ϕ at the root.

We point the interested reader to [48] for more precise details about this construction; for clarity, we provide a simple example from [48], using the formula

$$\phi = p \wedge \diamond_1(\neg p \wedge \diamond_2 p) \wedge \diamond_2(\neg p \wedge \diamond_1(p \wedge \diamond_2 p)).$$

We rewrite ϕ as $p \wedge \diamond_1(\neg p \wedge \mathbf{q}^2) \wedge \mathbf{r}^2$, where \mathbf{q}^2 is a ‘surrogate’ for $\diamond_2 p$, \mathbf{r}^2 for $\diamond_2(\neg p \wedge \mathbf{q}^1)$, \mathbf{q}^1 for $\diamond_1(p \wedge \mathbf{s}^2)$, and \mathbf{s}^2 for $\diamond_2 p$.

We construct a rooted L_1 -frame satisfying the rewritten formula (top left of Fig. 3.3); we make the node satisfying \mathbf{r}^2 into the root of an L_2 -frame satisfying its surrogate formula $\diamond_2(\neg p \wedge \mathbf{q}^1)$ and the \mathbf{q}^2 node into the root of a frame satisfying $\diamond_2 p$ (top right); we then proceed similarly with \mathbf{q}^1 (bottom left) and finally with \mathbf{s}^2 (bottom right) to find a $[L_1, L_2]$ -frame satisfying ϕ at its bottom point.

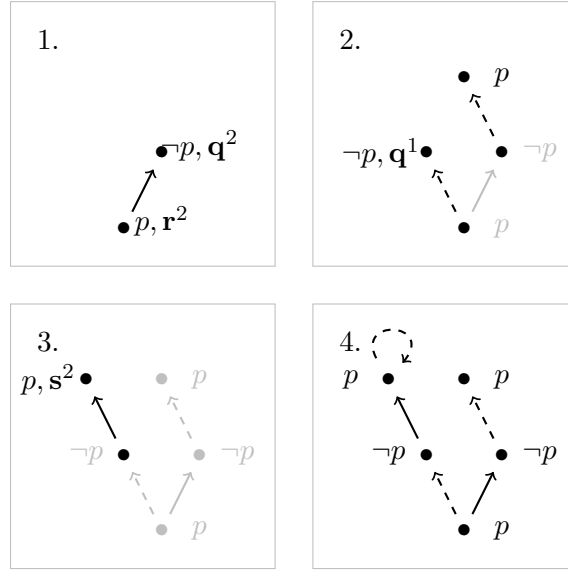


Figure 3.3: ‘Dovetailed’ construction of a frame for $\phi = p \wedge \diamond_1(\neg p \wedge \diamond_2 p) \wedge \diamond_2(\neg p \wedge \diamond_1(p \wedge \diamond_2 p))$.

For our current purposes it suffices to point out that the ‘dovetailed’ frames obtained by this method are always orthogonal, for this construction does not allow for two distinct points x and y to be reachable from each other by both R_1 and R_2 . ■

As an immediate consequence of Propositions 3.11 and 3.12:

Theorem 3.13. *The logic of $[L_1, L_2]$ -indexed frames is the fusion $L_1 \oplus L_2$.*

3.3 Orthogonal structures

In the definition for fully orthogonal frames (Def. 3.6) we demand the existence of equivalence relations which are supersets of the two given relations

and satisfy the properties of full orthogonality. These relations are not made explicitly part of the structure and are not taken into account when defining the logic.

In this section we consider structures $(X, R_1, R_2, \equiv_1, \equiv_2)$ satisfying FO1, FO2 and FO3, and we study the logic of these frames when we add modal operators to our language to explicitly account for the orthogonal equivalence relations.

Let us first note the following fact:

Lemma 3.14 (Generalized orthogonal frames). *If (W, R_1, R_2) is a frame such that there exist equivalence relations \equiv_1 and \equiv_2 on W satisfying*

- FO1. $R_i \subseteq \equiv_i$,
- FO2. $\equiv_1 \cap \equiv_2 = Id_W$, and
- FO3'. $(\equiv_1 \circ \equiv_2) = (\equiv_2 \circ \equiv_1)$,

then (W, R_1, R_2) is a disjoint union of fully orthogonal frames.

Proof. We leave it to the reader to check that FO3' implies that $\equiv_1 \circ \equiv_2$ is an equivalence relation. Let W' be an equivalence class of $\equiv_1 \circ \equiv_2$. For $i = 1, 2$, let R'_i and \equiv'_i be the restrictions of R_i and \equiv_i to W' . It is routine to check that (FO1) $R'_i \subseteq \equiv'_i$, (FO2) $\equiv'_1 \cap \equiv'_2 = Id_{W'}$, and (FO3) $\equiv'_1 \circ \equiv'_2 = (W')^2$.

Therefore, for each equivalence class W' modulo $\equiv_1 \circ \equiv_2$, (W', R'_1, R'_2) is a fully orthogonal frame and (W, R_1, R_2) is equal to the disjoint union $\bigcup_{W' \in W/\equiv_1 \circ \equiv_2} (W', R'_1, R'_2)$. ■

Definition 3.15. An *orthogonal structure* is a tuple $(W, R_1, R_2, \equiv_1, \equiv_2)$, where (W, R_1, R_2) is a birelational frame and \equiv_1 and \equiv_2 are equivalence relations on W satisfying FO1, FO2, and FO3' above. A *standard orthogonal structure* satisfies moreover

- FO3. $\equiv_1 \circ \equiv_2 = W^2$.

A tuple satisfying FO1 and FO3' is called a *semistructure*.

We define a semantics for (semi)structures $(W, R_1, R_2, \equiv_1, \equiv_2)$ with respect to a language containing operators \square_i and \blacksquare_i for $i = 1, 2$ as follows:

- $w \models \square_i \phi$ iff, for all v , $wR_i v$ implies $v \models \phi$;
- $w \models \blacksquare_i \phi$ iff, for all v , $w \equiv_i v$ implies $v \models \phi$.

Below we use the following notation: $L_1 \oplus L_2$ is the fusion of the logics L_1 and L_2 , defined in the previous section; given a logic L and a formula ϕ in the language of L , we let $L \oplus \phi$ denote the least normal modal logic containing $L \cup \{\phi\}$. Given a logic L in a unimodal language and given a modal operator \square , we let L_\square be the logic L expressed in the language using \square as its modal operator (e.g. $S5_{\blacksquare_1}$ represents the logic S5 for the \blacksquare_1 operator, and so on).

A very standard canonical model argument shows that:

Proposition 3.16. *The sound and complete logic of semistructures is*

$$K_{\square_1} \oplus K_{\square_2} \oplus S5_{\blacksquare_1} \oplus S5_{\blacksquare_2} \oplus \blacksquare_1 \blacksquare_2 \phi \leftrightarrow \blacksquare_2 \blacksquare_1 \phi \oplus \blacksquare_i \phi \rightarrow \square_i \phi \quad (i = 1, 2). \quad (Log_{\neg})$$

Moreover, if L_1 and L_2 are canonical unimodal logics, the logic of semistructures $(W, R_{1,2}, \equiv_{1,2})$ such that $(W, R_i) \models L_i$ for $i = 1, 2$ is

$$L_1 \oplus L_2 \oplus S5_{\blacksquare_1} \oplus S5_{\blacksquare_2} \oplus \blacksquare_1 \blacksquare_2 \phi \leftrightarrow \blacksquare_2 \blacksquare_1 \phi \oplus \blacksquare_i \phi \rightarrow \square_i \phi \quad (i = 1, 2). \quad (Log_{\neg}^{L_1 L_2})$$

Proof sketch. This uses the rather commonplace technique of *canonical models*; we point the reader to [29, Chapter 4] for full details and we simply offer a sketch here:

Let X be the set of maximal consistent sets⁵ of formulas in the language. We define the relations $xR_i y$ iff, for all ϕ , $\square_i \phi \in x$ implies $\phi \in y$ and $x \equiv_i y$ iff, for all ϕ , $\blacksquare_i \phi \in x$ implies $\phi \in y$.

Then one proves the *Truth Lemma*, showing that, given the valuation $V(p) = \{x \in X : p \in x\}$, it is the case that $x \models \phi$ iff $\phi \in x$.

We note that the logic $Log_{\neg}^{L_1 L_2}$ is canonical, for canonicity is preserved by fusions [69, Cor. 6] and the addition of ‘Sahlqvist axioms’ [29, Chapter 4], which include the axioms above. This canonicity ensures that $(X, R_i) \models L_i$; the S5 axioms for the \blacksquare_i ’s ensure that \equiv_i is an equivalence relation; the axiom $\blacksquare_1 \blacksquare_2 \phi \leftrightarrow \blacksquare_2 \blacksquare_1 \phi$ ensures that FO3’ is satisfied; finally, the axioms $\blacksquare_i \phi \rightarrow \square_i \phi$ ensure FO1.

Therefore $(X, R_{1,2}, \equiv_{1,2})$ is a semistructure and any consistent formula ϕ can be satisfied in it. \blacksquare

Let us call these logics Log_{\neg} and $Log_{\neg}^{L_1 L_2}$ respectively. Let us now see that Log_{\neg} is also the logic of orthogonal structures (and, in turn, of ‘standard’ structures).

Recall that if a *bounded morphism* (Def. 2.3) between frames F' and F is surjective, then every formula which is satisfiable in F' can be satisfied in F . (See [29, Thm. 3.14] for details).

We shall show that a semistructure is always the image of a bounded morphism departing from an orthogonal structure, which in turn will let us prove that the logic of orthogonal structures is the above.

The proof below utilises the notion of a *matrix enumeration*.

⁵ *Consistent* here means: there are no $\phi_1, \dots, \phi_n \in X$ such that $\phi_1 \wedge \dots \wedge \phi_n \rightarrow \perp \in L$. If a set is maximally consistent, for every formula ϕ , it is the case that either $\phi \in X$ or $\neg \phi \in X$.

Definition 3.17. Given sets I and X , an I -matrix enumeration of X is a map $f : I \times I \rightarrow X$ such that, for any fixed $i_0 \in I$, both maps

$$j \in I \mapsto f(i_0, j) \in X \text{ and } j \in I \mapsto f(j, i_0) \in X$$

are surjective.

Lemma 3.18. *If $\text{card}(I) \geq \text{card}(X)$, there exists an I -matrix enumeration of X .*

Proof. We distinguish two cases:

(i) X is finite. Suppose $X = \{0, \dots, n-1\}$ and $\text{card}(I) \geq n$. Choose n distinct elements $i_1, \dots, i_n \in I$. Set

$$\begin{aligned} f(i_k, i_j) &\equiv (k+j) \bmod n && \text{for } k, j \in \{1, \dots, n\}; \\ f(i_k, i) &= f(i, i_k) = k-1 && \text{for } k \in \{1, \dots, n\}, i \notin \{i_1, \dots, i_n\}; \\ f(i, j) &\text{arbitrary} && \text{for } i, j \notin \{i_1, \dots, i_n\}. \end{aligned}$$

We denote “ $a \equiv b \bmod c$ ” (and we say a is congruent with b modulo c) iff a is the unique natural number strictly smaller than c such that $b-a$ is a multiple of c . Let us see that f is a matrix enumeration. Set $i \in I$. If $i \notin \{i_1, \dots, i_n\}$, then $f(i, i_1) = 0, \dots, f(i, i_n) = n-1$. If $i = i_k$ for some $k \in \{1, \dots, n\}$, then $f(i_k, i_1), \dots, f(i_k, i_n)$ are congruent modulo n with $k+1, \dots, k+n$ and thus range over the set $\{0, \dots, n-1\}$. In both cases, the map $f(i, \cdot)$ is surjective; this reasoning is analogous for the map $f(\cdot, i)$.

(ii) X is infinite. In this case, I is obviously infinite as well. Let $\{I_1, I_2\}$ be a partition of I into two sets which are equipotent to I itself.⁶ Let $f_1 : I_1 \rightarrow X$ and $f_2 : I_2 \rightarrow X$ be two surjections. Let $\{k, l\} = \{1, 2\}$. The map f , defined as

$$f(i, j) = \begin{cases} f_k(i) & \text{if } i, j \in I_k \\ f_l(j) & \text{if } i \in I_k, j \in I_l, k \neq l \end{cases}$$

is the desired enumeration. Indeed, suppose $i \in I_1$. Note that the image of the map $f(i, \cdot)$ includes the set $\{f(i, j) : j \in I_2\}$, which is the image of f_2 , which is X . We reason analogously for the case $i \in I_2$ and for the map $f(\cdot, j)$. ■

With this:

Proposition 3.19. *A semistructure is a bounded morphic image of an orthogonal structure.*

⁶Note that the existence of such partition is equivalent to the fact that I is equipotent to $I \times \{0, 1\}$; this requires the Axiom of Choice for uncountable cardinalities [100].

Proof. Let $(W, R_{1,2}, \equiv_{1,2})$ be a semistructure. Let I be a set of indices with the same cardinality as W (this could be, e.g., W itself).

Let us consider the quotient set $W/\equiv_{1 \cap \equiv_2}$. Let us fix a matrix enumeration $f_{[w]} : I \times I \rightarrow [w]$ of each equivalence class $[w] \in W/\equiv_{1 \cap \equiv_2}$. We use w_{ij} as a shorthand for $f_{[w]}(i, j)$. Note that it is always the case that $w \equiv_k w_{ij}$ for $k = 1, 2$.

We define the following relations on the set $W' = W/\equiv_{1 \cap \equiv_2} \times I^2$:

$$\begin{aligned} ([w], i_1, i_2) \equiv'_1 ([v], j_1, j_2) & \text{ iff } w \equiv_1 v \text{ and } i_2 = j_2; \\ ([w], i_1, i_2) \equiv'_2 ([v], j_1, j_2) & \text{ iff } w \equiv_2 v \text{ and } i_1 = j_1; \\ ([w], i_1, i_2) R'_1 ([v], j_1, j_2) & \text{ iff } w_{i_1 i_2} R_1 v_{j_1 j_2} \text{ and } i_2 = j_2; \\ ([w], i_1, i_2) R'_2 ([v], j_1, j_2) & \text{ iff } w_{i_1 i_2} R_2 v_{j_1 j_2} \text{ and } i_1 = j_1. \end{aligned}$$

Let us see that this is an orthogonal structure. Indeed,

- FO1. $([w], i_1, i_2) R'_1 ([v], j_1, j_2)$ implies $w_{i_1 i_2} R_1 v_{j_1 j_2}$ and $i_2 = j_2$; the former implies $w_{i_1 i_2} \equiv_1 v_{j_1 j_2}$. This means that we have $w \equiv_1 v$ and, since $i_2 = j_2$, this gives $([w], i_1, i_2) \equiv'_1 ([v], j_1, j_2)$. The argument that shows that $R'_2 \subseteq \equiv'_2$ is identical.
- FO2. If $([w], i_1, i_2) \equiv'_k ([v], j_1, j_2)$ for $k = 1, 2$, then $i_1 = j_1$, and $i_2 = j_2$, and $(w, v) \in \equiv_1 \cap \equiv_2$, which implies $[w] = [v]$. Therefore, $\equiv'_1 \cap \equiv'_2 = Id_{W'}$.
- FO3'. If $([w], i_1, i_2) (\equiv'_1 \circ \equiv'_2) ([u], j_1, j_2)$, then $w (\equiv_1 \circ \equiv_2) u$. This, plus property (FO3') of the semistructure, implies that there exists some v' such that $w \equiv_2 v' \equiv_1 u$. But then

$$([w], i_1, i_2) \equiv'_2 ([v'], i_1, j_2) \equiv'_1 ([u], j_1, j_2).$$

This shows that $(\equiv'_1 \circ \equiv'_2) \subseteq (\equiv'_2 \circ \equiv'_1)$; the converse inclusion is analogous.

Finally, the map

$$([w], i_1, i_2) \in W/\equiv_{1 \cap \equiv_2} \times I^2 \mapsto w_{i_1 i_2} \in W$$

is a surjective bounded morphism. The fact that it is surjective comes from the matrix enumerations: given $w \in W$, fix any $i_0 \in I$ and there exists some $j \in I$ such that $w_{i_0 j} = f_{[w]}(i_0, j) = w$. For the forth condition, $([w], i_1, i_2) \equiv'_k ([v], j_1, j_2)$ implies $w_{i_1 i_2} \equiv_k v_{j_1 j_2}$ and $([w], i_1, i_2) R'_k ([v], j_1, j_2)$ implies $w_{i_1 i_2} R_k v_{j_1 j_2}$, by definition. For the back condition, if $w_{i_1 i_2} R_1 v$, then there exists an index $j \in I$ such that $f_{[v]}(j, i_2) = v$ (because $f_{[v]}$ is a matrix enumeration and thus the map $f_{[v]}(\cdot, i_2)$ is surjective), and, by definition,

$$([w], i_1, i_2) R'_1 ([v], j, i_2).$$

An analogous argument can be made for R_2, \equiv_1 and \equiv_2 . ■

As a consequence:

Theorem 3.20. *The sound and complete logic of standard orthogonal structures is Log_{\perp} ,*

$$K_{\square_1} \oplus K_{\square_2} \oplus S5_{\blacksquare_1} \oplus S5_{\blacksquare_2} \oplus \blacksquare_1 \blacksquare_2 \phi \leftrightarrow \blacksquare_2 \blacksquare_1 \phi \oplus \blacksquare_i \phi \rightarrow \square_i \phi.$$

Proof. Consequence of Propositions 3.16 and 3.19. ■

Remark 3.21. The construction in the proof above respects many properties of the R_i relations: for instance, if R_i is reflexive, transitive, symmetric, Euclidean (among other properties) then so is R'_i . This means that this technique can be used to prove that $Log_{\perp}^{L_1 L_2}$ is the logic of indexed structures (W, R_i, \equiv_i) where $(W, R_i) \models L_i$ for a large family of logics that includes S4, S5, KD45. I conjecture that the result is true for any pair L_1, L_2 of Kripke-complete unimodal logics.

Let us now define a semantics for this extended language directly on indexed frames $(W_1 \times W_2, R_1, R_2)$, taking advantage of the isomorphism between indexed frames and fully orthogonal frames given in the proof of Proposition 3.8. The fact that the isomorphic image of the equivalence classes of the ‘orthogonal’ equivalence relations are sets of the form $W_1 \times \{w_2\}$ and $\{w_1\} \times W_2$ allows us to consider the \blacksquare modalities as coordinate-wise ‘universal modalities’; that is to say, if we interpret formulas of the extended language on indexed frames as follows:

$$\begin{aligned} (w_1, w_2) \models \blacksquare_1 \phi & \text{ iff } (v, w_2) \models \phi \text{ for all } v \in W_1, \text{ and} \\ (w_1, w_2) \models \blacksquare_2 \phi & \text{ iff } (w_1, v) \models \phi \text{ for all } v \in W_2, \end{aligned}$$

then we have that:

Proposition 3.22. *Log_{\perp} is the sound and complete logic of indexed frames for the language including \square_i and \blacksquare_i operators.*

Proof. Soundness is routine. For completeness, given a formula $\phi \notin Log_{\perp}$, it suffices to use Thm. 3.20 to find a standard orthogonal structure $(W, R_{1,2}, \equiv_{1,2})$ that refutes ϕ , construct the indexed frame $(W/\equiv_2 \times W/\equiv_1, R_1, R_2)$ isomorphic to (W, R_1, R_2) via Prop. 3.8 and note that the equivalence relations \cong_1 and \cong_2 , defined for $i = 1, 2$ as

$$([w]_2, [v]_1) \cong_i ([w']_2, [v']_1) \text{ iff } x_{wv} \equiv_i x_{w'v'},$$

relate two pairs if and only if their j -th coordinate coincides, for $j \neq i$. ■

We finish this section by pointing out the fact that Log_{\perp} enjoys the Finite Model Property with respect to semistructures, orthogonal structures and indexed frames, in the following sense:

Proposition 3.23. *If $\phi \notin \text{Log}_{\perp}$, then there exists a finite indexed frame refuting ϕ .*

Proof. This involves a rather standard filtration argument. (See [29, Chapter 2] for details on this technique).

Given a consistent formula ϕ , we let $(W, R_{1,2}, \equiv_{1,2})$ be a semistructure satisfying ϕ at a point w_0 , we let Γ be a finite set of formulas closed under subformulas such that $\phi \in \Gamma$, and we define an equivalence relation $w \sim_{\Gamma} v$ iff for all $\psi \in \Gamma$, ($w \models \psi$ iff $v \models \psi$). We define relations in the quotient set W/\sim_{Γ} as follows: for $i = 1, 2$,

$[w]_{\Gamma} \equiv'_i [v]_{\Gamma}$ iff, for all $\blacksquare_i \psi \in \Gamma$, ($w \models \blacksquare_i \psi$ iff $v \models \blacksquare_i \psi$), and

$[w]_{\Gamma} R'_i [v]_{\Gamma}$ iff $[w]_{\Gamma} \equiv'_i [v]_{\Gamma}$ and for all $\square_i \psi \in \Gamma$ ($w \models \square_i \psi$ implies $v \models \psi$).

We leave it to the reader to check that the resulting tuple is a semistructure and a filtration and therefore that $[w_0]_{\Gamma} \models \phi$. We can then use Prop. 3.19 and Lemma 3.14 to obtain an indexed frame satisfying ϕ . ■

I conjecture that, if L_1 and L_2 have the Finite Model Property, then for all $\phi \notin \text{Log}_{\perp}^{L_1 L_2}$ there exists a finite $[L_1, L_2]$ -semistructure (perhaps a finite $[L_1, L_2]$ -indexed frame) refuting ϕ ; this problem, however, remains open.

Discussion. We have identified a structure that shows up with relative frequency in different areas of modal logic; we have argued that an independent study of this structure is warranted and have taken the first steps towards it.

These structures have been shown to be completely characterised by the ‘orthogonality’ of their relations. Proofs of completeness of frameworks based on indexed frames are not particularly easy to tackle; as an example, we point the reader to the completeness proof of SEL in [90]. I hope that the above observations about orthogonality will help facilitate some of these proofs.

Some work remains to be done and many questions are open. Among these are the following:

- Is $\text{Log}_{\perp}^{L_1 L_2}$ the logic of orthogonal structures $(W, R_1, R_2, \equiv_1, \equiv_2)$ such that $(W, R_i) \models L_i$, for any pair of Kripke-complete logics L_1 and L_2 ? Can a formula $\phi \notin \text{Log}_{\perp}^{L_1 L_2}$ be refuted in a finite indexed frame whenever L_1 and L_2 have the FMP? I conjecture an affirmative answer to these questions, albeit further research will be necessary to resolve them.
- Perhaps the most salient question: how does one generalise the definitions and results in this chapter to the n -dimensional case? The reader may find that there are two reasonable generalisations of the notion of indexed frames to the n -th dimension:

- (A) $(W_1 \times \dots \times W_n, R_1, \dots, R_n)$ such that
 $(w_j)_{j=1}^n R_i (v_j)_{j=1}^n$ implies $w_j = v_j$ for all $j \neq i$;
- (B) $(W_1 \times \dots \times W_n, R_1, \dots, R_n)$ such that
 $(w_j)_{j=1}^n R_i (v_j)_{j=1}^n$ implies $w_i = v_i$.

Out of these two, I suggest (A) is more appropriate, for it does not make much sense to apply (B) to $n = 1$, and (A) is the only one which still generalises n -dimensional products. Many of the results of this chapter may translate relatively easily to the n -dimensional case, whereas some may not. Future work shall be devoted to this question.

Some case studies

4

Using the ‘orthogonality’ results from the last chapter, we provide necessary and sufficient conditions for an arbitrary Kripke frame to be isomorphic to a product or a ‘subset space frame’. We also provide some categorical equivalences among different classes of models for STIT logic.¹

Unrealised futures are just branches of the past: dry branches.

Italo Calvino, *Invisible Cities*.

IN THIS SHORT CHAPTER we return to some of the examples in Section 3.1 of the previous Chapter and, with the help of our orthogonality results above, we abstract from the “indexed frame” definition and give necessary and sufficient conditions on orthogonal frames to be isomorphic to these structures.

4.1 Products.

We recall the definition of product from Example 3.2:

Definition 4.1. The *product* of two Kripke frames (W_1, R_1) and (W_2, R_2) is the frame

$$(W_1 \times W_2, R_1^H, R_2^V),$$

where the *horizontal* and *vertical* relations R_1^H and R_2^V are defined as follows:

$$\begin{aligned} (w_1, w_2)R_1^H(w'_1, w'_2) & \text{ iff } w_2 = w'_2 \text{ and } w_1 R_1 w'_1, \text{ and} \\ (w_1, w_2)R_2^V(w'_1, w'_2) & \text{ iff } w_1 = w'_1 \text{ and } w_2 R_2 w'_2 \end{aligned}$$

We have:

¹This chapter is a substantial extension of the last section of the paper

Philippe Balbiani and Saúl Fernández González. Orthogonal frames and indexed relations. In *Logic, Language, Information, and Computation, 27th International Workshop WoLLIC 2021*, pages 219–234. Springer, 2021.

Proposition 4.2. *A frame (X, R_1, R_2) is isomorphic to a product of Kripke frames if and only if there exist two equivalence relations \equiv_1 and \equiv_2 such that:*

- FO1. $R_i \subseteq \equiv_i$, for $i = 1, 2$;
- FO2. $\equiv_1 \cap \equiv_2 = Id_X$;
- FO3. $\equiv_1 \circ \equiv_2 = X^2$, and
- P1. $(R_i \circ \equiv_j) = (\equiv_j \circ R_i)$, for $i \neq j$.

Proof. That a product $(W_1, R_2) \times (W_2, R_2)$ satisfies these properties is trivial: it suffices to use the equivalence relations $(w_1, w_2) \equiv_i (v_1, v_2)$ iff $w_j = v_j$ for $\{i, j\} = \{1, 2\}$.

Now let us consider a frame (W, R_1, R_2) satisfying the properties above and let x_{wv} denote the unique element in $[w]_2 \cap [v]_1$ (as in the proof of Prop. 3.8). This frame satisfies, for all $w, w', v, v' \in W$: $x_{wv} R_1 x_{w'v}$ iff $x_{wv'} R_1 x_{w'v}$. Indeed, if $x_{wv} R_1 x_{w'v}$, since $x_{w'v} \equiv_2 x_{w'v'}$, then by (P1) there must exist some y such that $x_{wv'} \equiv_2 y R_1 x_{w'v'}$; this y must be $x_{wv'}$, for $y \in [x_{wv'}]_2 \cap [x_{w'v'}]_1 = [w]_2 \cap [v']_1$. We can thus define a relation on W/\equiv_2 as $[w]_2 R'_1 [w']_2$ iff $x_{wv} R_1 x_{w'v}$ for some v (equivalently: for all v). We proceed similarly to define a relation R'_2 on W/\equiv_1 : $[v]_1 R'_2 [v']_1$ iff $x_{wv} R_1 x_{wv'}$ for some (for all) w .

The product $(W/\equiv_2, R'_1) \times (W/\equiv_1, R'_2)$ is equal to $(W/\equiv_2 \times W/\equiv_1, R_1, R_2)$, isomorphic to (W, R_1, R_2) by Prop. 3.8. \blacksquare

4.2 Subset spaces

Recall the notion of a *subset space frame* from Example 3.3:

Definition 4.3. A *subset space* is a tuple (X, \mathcal{O}) such that $\mathcal{O} \subseteq \mathcal{P}(X)$ and both X and \mathcal{O} are nonempty.

A subset space induces a *subset space frame* $(\mathcal{O}_X, \equiv_K, \geq_\square)$ where

- $\mathcal{O}_X := \{(x, U) : x \in U \in \mathcal{O}\}$,
- $(x, U) \equiv_K (y, V)$ iff $U = V$, and
- $(x, U) \geq_\square (y, V)$ iff $x = y$ and $U \supseteq V$.




Figure 4.1: Example of a subset space (left) and its associated subset space frame (right).

Given a bimodal language containing K and \square operators, the semantics for subset spaces is as follows:

$$(x, U) \models K\phi \quad \text{iff } (y, U) \models \phi \text{ for all } y \in U;$$

$$(x, U) \models \square\phi \quad \text{iff } (x, V) \models \phi \text{ for all } V \in \mathcal{O} \text{ such that } x \in V \subseteq U.$$

Note that this coincides with the Kripke semantics on subset space frames using the above relations.

We have:

Proposition 4.4. *A frame (W, R_K, R_\square) is isomorphic to a subset space frame if and only if it satisfies*

SS1. R_K is an equivalence relation;

SS2. R_\square is a partial order (i.e. reflexive, transitive and antisymmetric);

SS3. $R_\square \circ R_K \subseteq R_K \circ R_\square$,

and there exists an equivalence relation \equiv_\square such that

FO1. $R_\square \subseteq \equiv_\square$;

FO2. $R_K \cap \equiv_\square = Id_W$,

and, moreover,

SS4. $([R_K \circ \equiv_\square](u) \supseteq [R_K \circ \equiv_\square](v) \text{ and } u \equiv_\square v) \text{ imply } uR_\square v$;

SS5. $[R_K \circ \equiv_\square](u) = [\equiv_\square \circ R_K](v) \text{ implies } uR_K v$.

Proof. For the left-to-right direction, given a subset space frame we consider the relations $(x, U)R_K(y, V)$ iff $U = V$, $(x, U)R_\square(y, V)$ iff $x = y$ and $U \supseteq V$, and $(x, U) \equiv_\square(y, V)$ iff $x = y$. We note that

$$(R_K \circ \equiv_\square)(x, U) = \{(x', U') \in \mathcal{O}_X : x' \in U\},$$

and we leave it to the reader to check that this satisfies all the properties in Prop. 4.4.

Let us now consider a frame with these properties. We let $[\cdot]_{\square}$ and $[\cdot]_K$ denote the equivalence classes of \equiv_{\square} and R_K . Let us define the subset space

$$\begin{aligned} X &= W / \equiv_{\square} = \{[w]_{\square} : w \in W\} \\ \mathcal{O} &= \{U_v : v \in W\}, \text{ where } U_v = \{[w]_{\square} \in X : v[R_K \circ \equiv_{\square}]w\}. \end{aligned}$$

Note that $[w]_{\square} \in U_v$ if and only if $[w]_{\square} \cap [v]_K \neq \emptyset$.

By FO2, an intersection $[w]_{\square} \cap [v]_K$ of two equivalence classes is at most a singleton. Let us map an element $([w]_{\square}, U_v)$ in the graph of (X, \mathcal{O}) to the unique element in $[w]_{\square} \cap [v]_K$. This is a bijection whose inverse maps each $w \in W$ to $([w]_{\square}, U_w)$. We define relations \equiv_K and \geq_{\square} on this graph as in Example 3.3 and, to show that this map is an isomorphism, it suffices to show that

$$\begin{aligned} wR_Kv &\text{ iff } ([w]_{\square}, U_w) \equiv_K ([v]_{\square}, U_v), \text{ and} \\ wR_{\square}v &\text{ iff } ([w]_{\square}, U_w) \geq_{\square} ([v]_{\square}, U_v). \end{aligned}$$

We start with the second item. From left to right, if $wR_{\square}v$, then $[w]_{\square} = [v]_{\square}$ by FO1, and let us see that $U_w \supseteq U_v$. If $[y]_{\square} \in U_v$, then there is a unique element $x \in [y]_{\square} \cap [v]_K$. But since $wR_{\square}vR_Kx$, it follows by SS3 that there must exist some x' such that $wR_Kx'R_{\square}x$. Since $x' \equiv_{\square} x$, by FO1, and $x \equiv_{\square} y$, it follows that $x' \in [w]_K \cap [y]_{\square}$, and thus $[y]_{\square} \in U_w$. From right to left, it suffices to see that $U_w \supseteq U_v$ and $w \equiv_{\square} v$ implies $wR_{\square}v$. But this follows directly from (SS4), noting that $U_w \supseteq U_v$ implies $[R_K \circ \equiv_{\square}](w) \supseteq [R_K \circ \equiv_{\square}](v)$.

For the first item it suffices to show that wR_Kv iff $U_w = U_v$. The left-to-right direction is immediate from the definition of U_w , whereas the right-to-left direction follows from SS5. \blacksquare

4.3 STIT logic

In [33], the authors compare three different (albeit ultimately equivalent) semantics for STIT logic. One of these types of models was briefly discussed in Example 3.5.

In this Section we shall properly introduce STIT logic (in the version depicted in [33]), we shall show that the three types of models depicted in that paper are based on orthogonal frames (some with minimal ‘rewritings’), and finally we shall see that two of these families of models are categorically equivalent (in a truth-preserving manner), after which we shall add a simple condition to make them both equivalent to the third.

In this section we deal with orthogonal frames that do not have two relations, but rather two different families of relations which are ‘orthogonal’ to each other. We thus need to extend our definition of orthogonal frame:

Definition 4.5. In the present section, we will say that a multi-relational frame $(W, R_1, \dots, R_n, S_1, \dots, S_m)$ is *orthogonal* if there exist two equivalence relations \equiv_R, \equiv_S on W such that

- FO1. $\equiv_R \cap \equiv_S = Id_W$;
 FO2. $R_i \subseteq \equiv_R$ and $S_j \subseteq \equiv_S$ for $1 \leq i \leq n$ and $1 \leq j \leq m$.

We adjust our definition of ‘indexed frame’ accordingly:

Definition 4.6. An *indexed frame* (for the remainder of this section) is a frame of the form $(W_1 \times W_2, R_1, \dots, R_n, S_1, \dots, S_m)$ such that the R_i relations only connect pairs of points which share their second component, and the S_j relations only connect pairs which share their first.

4.3.1 The language of *seeing-to-it-that*

Although there are papers by the same authors on this topic that precede it, [20] can be considered the foundational text of STIT logic.

The logic of *seeing-to-it-that* is a multi-agent logic of action and agency, wherein modal operators are introduced in order to encode in one way or another the notion that an agent ‘sees to it that’ a certain outcome is achieved. Models for STIT generally present some notion of branching time, where at each point in time many possible futures are accessible and agents have the power to ensure some of these possible futures will not ‘become real’.

The variant of STIT presented in [33] is built upon the language \mathcal{L}_{STIT} , defined as follows:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid [i]\phi \mid [Agt]\phi \mid \Box\phi \mid G\phi \mid H\phi,$$

where $p \in \mathbf{Prop}$, a countable set of propositional variables, and $Agt = \{1, \dots, n\}$ is a finite set of agents, and $i \in Agt$.

The $[i]\phi$ operator is the aforementioned ‘STIT operator’, and is read ‘agent i makes sure that ϕ ’ (or ‘... *sees to it that* ϕ ’). $[Agt]\phi$ is read ‘all agents, acting together, see to it that ϕ ’.

$G\phi$ and $H\phi$ are temporal operators, meant to be read ‘it will always be the case (resp. ‘it has always been the case’) that ϕ ’; their duals are $F\phi$ and $P\phi$, respectively.

$\Box\phi$ is a necessity operator, which reads ‘ ϕ is necessarily true’.²

²We point the philosophically inclined reader to [31] for some interesting remarks on the combination of this operator with the future operator F , which allows for a distinction between an ‘Ockhamist’ view of the future tense (wherein ‘ p will be the case’= Fp) and an ‘Antactualist’ one (wherein ‘ p will be the case’= $\Box Fp$).

4.3.2 Three types of STIT models

Three families of models for this language are presented in [33], which we recall below:

T-STIT frames.

Definition 4.7. A *temporal Kripke STIT frame*, or *T-STIT frame*, is a tuple

$$(X, \equiv_{\square}, \{\equiv_i\}_{i \in \text{Agt}}, \equiv_{\text{Agt}}, \prec_G)$$

where X is a nonempty set of points, \equiv_{\square} , \equiv_i and \equiv_{Agt} are equivalence relations and \prec_G is a serial and transitive relation satisfying:

- C1. For all i , $\equiv_i \subseteq \equiv_{\square}$;
- C2. if $x_1 \equiv_{\square} \dots \equiv_{\square} x_n$, then there exists some $y \in X$ such that $x_i \equiv_i y$ for all i ;
- C3. $\equiv_{\text{Agt}} = \bigcap_{i \in \text{Agt}} \equiv_i$.
- C4. if $x \prec_G y$ and $x \prec_G z$, then either $y \prec_G z$ or $z \prec_G y$ or $y = z$;
- C5. if $y \prec_G x$ and $z \prec_G x$, then either $y \prec_G z$ or $z \prec_G y$ or $y = z$;
- C6. $(\prec_G \circ \equiv_{\square}) \subseteq (\equiv_{\text{Agt}} \circ \prec_G)$;
- C7. if $x \equiv_{\square} y$, then $x \not\prec_G y$.

A *T-STIT model* is a T-STIT frame along with a valuation $V : \text{Prop} \rightarrow 2^X$ satisfying:

- C8. If $x \equiv_{\square} y$, then $x \in V(p)$ iff $y \in V(p)$.

Note that we can substitute C4 and C5 for the following condition:

- C45. (X, \prec_G) is a disjoint union of linearly ordered sets.

Frames of this type are obviously orthogonal: all the equivalence relations are contained in \equiv_{\square} whereas the order \prec_G is contained in the equivalence relation \prec_G^* determined by its connected components. Each of these connected components is a maximal chain (by C4 and C5), and condition C7 ensures that the intersection of one of these with an equivalence class of \equiv_{\square} is at most a singleton.

Formulas of \mathcal{L}_{STIT} are interpreted on these models with respect to a point $x \in X$ as one might expect:

$$\begin{aligned} x \models \Box\phi & \quad \text{iff } x \equiv_{\Box} y \text{ implies } y \models \phi; \\ x \models [i]\phi & \quad \text{iff } x \equiv_i y \text{ implies } y \models \phi; \\ x \models [Agt]\phi & \quad \text{iff } x \equiv_{Agt} y \text{ implies } y \models \phi; \\ x \models G\phi & \quad \text{iff } x \prec_G y \text{ implies } y \models \phi; \\ x \models H\phi & \quad \text{iff } x \succ_G y \text{ implies } y \models \phi. \end{aligned}$$

Kamp Agent Frames.

Definition 4.8 (Kamp frame). A *Kamp agent frame* is a tuple

$$(W, T, \mathcal{O}, \{\sim_t\}_{t \in T}, \{\sim_{t,i}\}_{t \in T, i \in Agt}, \{\sim_{t,Agt}\}_{t \in T}),$$

where W is a nonempty set, T is an infinite set, and $\mathcal{O} : W \rightarrow 2^T \times T^2$ is a function which assigns to each $w \in W$ an infinite linearly ordered set $(T_w, <_w)$, with $\bigcup_{w \in W} T_w = T$.

Relations \sim_t , $\sim_{t,i}$ and $\sim_{t,Agt}$ are equivalence relations defined on the set $\{w : t \in T_w\}$ satisfying:

- K1. For all $t \in T, i \in Agt$, $\sim_{t,i} \subseteq \sim_t$;
- K2. if $w_1 \sim_t \dots \sim_t w_n$, then there exists some $v \in W$ such that $w_i \sim_{t,i} v$ for all i ;
- K3. $\sim_{t,Agt} = \bigcap_{i \in I} \sim_{t,i}$;
- K4. $t <_w t'$ and $w \sim_{t'} w'$ imply $w \sim_{t,Agt} w'$;
- K5. if $w \sim_t w'$, then $\{t' \in T_w : t' <_w t\} = \{t' \in T_{w'} : t' <_{w'} t\}$.


A *Kamp agent model* is a Kamp frame along with a valuation $V : Prop \rightarrow 2^T$.

Formulas of \mathcal{L}_{STIT} are read with respect to pairs (w, t) such that $t \in T_w$ as follows:

$$\begin{aligned} (w, t) \models \Box\phi & \quad \text{iff } w \sim_t v \text{ implies } (v, t) \models \phi; \\ (w, t) \models [i]\phi & \quad \text{iff } w \sim_{t,i} v \text{ implies } (v, t) \models \phi; \\ (w, t) \models [Agt]\phi & \quad \text{iff } w \sim_{t,Agt} v \text{ implies } (v, t) \models \phi; \\ (w, t) \models G\phi & \quad \text{iff } t <_w t' \text{ implies } (w, t') \models \phi; \\ (w, t) \models H\phi & \quad \text{iff } t >_w t' \text{ implies } (w, t') \models \phi. \end{aligned}$$

Remark 4.9. It is of note that condition K5 is not included in the definition of Kamp frames of [33]. Instead, it is incorrectly claimed that K5 follows

from K1, K3 and K4. To see why this is not the case, let $W = \{w, v\}$, $T = T_w = T_v = \mathbb{Z}$, let $<_w$ be the natural ordering of the integers, let $<_v$ be the inverse of $<_w$, and let $\sim_t = \sim_{t,i} = \sim_{t,Ag t} = W^2$.



We leave it to the reader to check that this satisfies all properties of the definition above except (K5). After a discussion with one of the authors of [33], I have added (K5) to the above definition.

The above definition, although it presents families of indexed relations, does not strictly define an indexed frame on account of the \mathcal{O} function. However, one only needs to tweak it a bit to obtain a definition of what amounts to an indexed frame:

Definition 4.10 (Kamp (indexed) frame). A *Kamp indexed frame* is a tuple

$$(X, <, \sim, \{\sim_i\}_{i \in Agt}, \sim_{Agt}),$$

where X is a nonempty subset of some Cartesian product $W \times T$ such that each set $X_w := \{t : (w, t) \in X\}$ is infinite, $<$ is a strict order relation such that each set $\{w\} \times X_w$ is linearly ordered, and \sim, \sim_i, \sim_{Agt} are equivalence relations satisfying:

KI01. $(w, t) \sim (w', t')$ implies $t = t'$;

KI02. $(w, t) < (w', t')$ implies $w = w'$;

and, moreover,

KI1. $\sim_i \subset \sim$ for all $i \in Agt$;

KI2. if $(w_1, t) \sim \dots \sim (w_n, t)$, then there is some v such that $(w_i, t) \sim_i (v, t)$ for all i ;

KI3. $\sim_{Agt} = \bigcap_{i \in Agt} \sim_i$;

KI4. $(w, t) < (w', t')$ and $(w, t') \sim (w', t')$ imply $(w, t) \sim_{Agt} (w', t)$;

KI5. if $(w, t) \sim (w', t)$, then $\{t' \in X_w : (w, t') < (w, t)\} = \{t' \in X_{w'} : (w', t') < (w', t)\}$.

A *Kamp indexed model* is a Kamp indexed frame along with a valuation $V : \text{Prop} \rightarrow 2^T$. The semantics of \mathcal{L}_{STIT} on these models is the usual relational semantics, as one might expect:

$$\begin{array}{ll} (w, t) \models \Box\phi & \text{iff } (w, t) \sim (w', t) \text{ implies } (w', t) \models \Box\phi; \\ (\text{resp. } [i]\phi, [Agt]\phi) & (\text{resp. } \sim_i, \sim_{Agt}); \\ (w, t) \models G\phi & \text{iff } (w, t) < (w', t) \text{ implies } (w', t) \models \phi; \\ (\text{resp. } H\phi) & (\text{resp. } >); \end{array}$$

The following should also be unsurprising:

Lemma 4.11. *Definitions 4.8 and 4.10 define two categorically isomorphic classes of frames.*

Indeed, given a Kamp indexed frame (Def. 4.10), define an ordinary Kamp frame (Def. 4.8) by setting $T_w = X_w$, $t <_w t'$ iff $(w, t) < (w, t')$, $w \sim_t w'$ iff $(w, t) \sim (w', t)$ and so on; conversely, given a Kamp frame according to the first definition, define a Kamp indexed frame by setting $X = \{(w, t) : t \in T_w\}$, $(w, t) < (w, t')$ iff $t <_w t'$, and $w \sim_t w'$ iff $(w, t) \sim (w', t)$. One sees that the semantics above are also equivalent.

One of the main aims of this section is to prove the following result:

Theorem 4.12. *The class of T-STIT frames and the class of Kamp frames are categorically equivalent.*

It is of note that this result is incidentally proven in [33], albeit never explicitly stated. Indeed, the authors prove the fact that

T-STIT frames and Kamp frames satisfy the same formulas

by defining a truth-preserving model transformation from a T-STIT frame into a Kamp frame and vice versa; these transformations define the functors that give the categorical equivalence. However, thanks to Def. 4.10 and the results of the previous Chapter, we can give here an easier proof of this result:

Proof of Thm. 4.12. We prove the following two items:

(i) *Every Kamp frame is a T-STIT frame.* We use Def. 4.10 and we see that it satisfies all the postulates of Def. 4.7. (C1), (C2) and (C3) correspond directly to (KI1), (KI2) and (KI3); (C4) and (C5), in their combined form (C45), correspond to the fact that each X_w is linearly ordered; (C6) follows from (KI4) and (KI5), for if $(w, t) < (w, t') \sim (w', t')$, then we have that $(w', t) \in X$ and $(w', t) < (w', t)$ by (KI5), and that $(w, t) \sim (w', t)$ by (KI4); finally, the orthogonality condition (C7) corresponds to the fact that, if

$$(w, t)(\sim \cap <)(w', t'),$$

then $(w, t) = (w', t')$ against the irreflexivity of $<$.


(ii) *Every T-STIT frame is isomorphic to a Kamp frame.* We know from Prop. 3.8 that a T-STIT frame is isomorphic to the generated subframe of

$$(X / \prec_G^* \times X / \equiv_{\square}, \equiv'_{\square}, \{\equiv'_i\}_{i \in \text{Agt}}, \equiv'_{\text{Agt}}, \prec'_G)$$

given by the subset $X' = \{([x]_{\prec_G^*}, [y]_{\equiv_{\square}}) : [x]_{\prec_G^*} \cap [y]_{\equiv_{\square}} \neq \emptyset\}$. By an argument analogous to the one above, we see that this subframe satisfies the postulates of Def. 4.10. We leave most details of this argument to the reader and simply show that this isomorphic structure satisfies (KI4) and (KI5):

For cleanliness we denote $[x, y] := ([x]_{\prec_G^*}, [y]_{\equiv_{\square}})$ and we denote by a_{xy} the unique element of $[x]_{\prec_G^*} \cap [y]_{\equiv_{\square}}$. Suppose $[x, y] \prec'_G [x', y'] \equiv'_{\square} [x', y']$. Then we have $a_{xy} \prec_G a_{x'y'} \equiv_{\square} a_{x'y'}$. By (C6) there is some $b \in X$ such that $a_{xy} \equiv_{\text{Agt}} b \prec_G a_{x'y'}$. But since $b \equiv_{\square} a_{xy} \equiv_{\square} y$ and $b \prec_G^* a_{x'y'} \prec_G^* x'$, we have that $b \in [x']_{\prec_G^*} \cap [y]_{\equiv_{\square}}$ thus $b = a_{x'y'}$. We thus have that $[x, y] \equiv'_{\text{Agt}} [x', y]$ (KI4), and that $[x', y] \prec'_G [x', y']$ (KI5).

Items (i) and (ii) give the functors which define the categorical equivalence between the categories of Kamp and T-STIT frames, as showcased in Table 4.1. ■



$$\begin{aligned}
 F(X, \sim, \{\sim_i\}_{i \in \text{Agt}}, \sim_{\text{Agt}}, <) &= (X, \sim, \{\sim_i\}_{i \in \text{Agt}}, \sim_{\text{Agt}}, <); \\
 Ff &= f.
 \end{aligned}$$

$$\begin{aligned}
 G(X, \equiv_{\square}, \{\equiv_i\}_{i \in \text{Agt}}, \equiv_{\text{Agt}}, \prec_G) &= (X', \equiv'_{\square}, \{\equiv'_i\}_{i \in \text{Agt}}, \equiv'_{\text{Agt}}, \prec'_G); \\
 Gg([x]_{\prec_G^*}, [y]_{\equiv_{\square}}) &= g(a_{xy}).
 \end{aligned}$$

Table 4.1: *Categorical equivalence between the categories of Kamp and T-STIT frames. F is the identity functor, whereas G gives the isomorphic Kamp frame described in (ii) above; f and g are morphisms between Kamp frames and T-STIT frames respectively.*

Choice B-trees. Let us discuss the last of the three structures considered in [33], and give a class of indexed frames categorically isomorphic to it.

Definition 4.13. A *choice B-tree* is a tuple

$$(M, \prec, B, \{\sim_{m,i}\}_{m \in M, i \in \text{Agt}}, \{\sim_{m,\text{Agt}}\}_{m \in M}),$$

where

- (M, \prec) is an infinite tree (i.e., \prec is a binary relation on M that is serial, transitive, irreflexive, and past-linear³);
- B is a *bundle*, i.e., a collection of maximal chains of (M, \prec) (these shall be called *histories* from now on) such that for all $m \in M$ there exists some $h \in B$ such that $m \in h$, and
- the relations $\sim_{m,i}$ and $\sim_{m,\text{Agt}}$ are equivalence relations on the set $B_m := \{h \in B : m \in h\}$ such that
 - B1. For all $h_1, \dots, h_n \in B_m$, there is some $h \in B_m$ such that $h_i \sim_{m,i} h$ for all i ;
 - B2. $\sim_{m,\text{Agt}} = \bigcap_{i \in \text{Agt}} \sim_{m,i}$;
 - B3. $m \prec m'$ and $h, h' \in B_{m'}$ imply $h, h' \in B_m$ and $h \sim_{m,\text{Agt}} h'$.

Valuations assign subsets of M to propositional variables. Formulas are read with respect to pairs (m, h) with $m \in h \in B$ as follows:

$$\begin{aligned} (m, h) \models \Box\phi & \quad \text{iff } (m, h') \models \phi \text{ for all } h' \in B_m; \\ (m, h) \models [i]\phi & \quad \text{iff } h \sim_{m,i} h' \text{ implies } (m, h') \models \phi; \\ (m, h) \models [\text{Agt}]\phi & \quad \text{iff } h \sim_{m,\text{Agt}} h' \text{ implies } (m, h') \models \phi; \\ (m, h) \models G\phi & \quad \text{iff } m \prec m' \text{ and } m' \in h \text{ imply } (m', h) \models \phi; \\ (m, h) \models H\phi & \quad \text{iff } m \succ m' \text{ and } m' \in h \text{ imply } (m', h) \models \phi. \end{aligned}$$

Again, these frames are very similar to indexed frames but not quite. Let us now establish a class of indexed frames which is categorically isomorphic to them.

Definition 4.14. An (*indexed*) *choice B-tree* will be a tuple

$$(X, \equiv, \equiv_i, \equiv_{\text{Agt}}, \prec)$$

such that:

³In the sense that the set $\{m' : m' \prec m\}$ is linearly ordered for all m .

- i. X is a nonempty subset of some cartesian product $M \times B$ with full projection to the first component (i.e. for all $m \in M$, there is some $h \in B$ such that $(m, h) \in X$).

(For $m \in M$ and $h \in B$, we define the projections

$$B_m := \{h \in B : (m, h) \in X\} \text{ and } M_h := \{m \in M : (m, h) \in X\}.)$$

- ii. $<$ is a serial irreflexive partial order on X such that:

- $<_1$. $(m, h) < (m', h')$ implies $h = h'$;
- $<_2$. for each $h \in B$, the set $\{(m', h') \in X : h' = h\}$ is a chain;
- $<_3$. for all $h, h' \in B_m$,
 $\{m' \in M_h : (m', h) < (m, h)\} = \{m' \in M_{h'} : (m', h') < (m, h')\}$;
- $<_4$. $M_h \subseteq M_{h'}$ implies $h = h'$.

- iii. \equiv is a universal relation on the second component, i.e.,

$$\text{EQ. } (m, h) \equiv (m', h') \text{ iff } m = m';$$

- iv. $\equiv_i \subseteq \equiv$ for all $i \in \text{Agt}$;

- v. Moreover, it satisfies the following:

- BI1. $x_1 \equiv \dots \equiv x_n$ implies $\bigcap_i \equiv_i(x_i) \neq \emptyset$;
- BI2. $\equiv_{\text{Agt}} = \bigcap_{i \in \text{Agt}} \equiv_i$;
- BI3. $(m, h) < (m', h) \equiv (m', h')$ implies $(m, h) \equiv_{\text{Agt}} (m, h')$.

Given an indexed choice B-tree defined in this way, along with a valuation $V : \text{Prop} \rightarrow 2^M$, we interpret formulas on pairs $(m, h) \in X$ by using the \equiv relation to read the \square modality, \equiv_i for $[i]$, \equiv_{Agt} for $[\text{Agt}]$, $<$ for G and \succ for H .

And we have:

Lemma 4.15. *The two above definitions define categorically equivalent classes of frames; moreover, this categorical equivalence is truth-preserving.*

Proof. Given an indexed choice B-tree \mathcal{B} according to the second definition, we define \prec on M as follows:

$$m \prec m' \text{ iff } \exists h : (m, h) \in X \ \& \ (m', h) \in X \ \& \ (m, h) < (m', h).$$

We likewise define equivalence relations on B_m as follows:

$$h \sim_{m,x} h' \text{ iff } (m, h) \equiv_x (m, h'), \text{ for } x \in \text{Agt} \cup \{\text{Agt}\}.$$

Finally, we let the set B' be given by $\{M_h : h \in B\}$.

It holds that $\mathcal{B}' = (M, \prec, B', \{\sim_{m,i}\}, \{\sim_{m,\text{Agt}}\})$ is a choice B-tree according to the first definition.

Indeed, let us see that the \prec relation is indeed irreflexive, transitive and past-linear:

Irreflexivity is obvious, whereas transitivity is given by (\prec_3): indeed, if $m \prec m' \prec m''$, then for all $h'' \in B_{m''}$ we have that $(m'', h'') > (m', h'') > (m, h'')$, and thus $m \prec m''$.

Past-linearity follows from (\prec_2) and (\prec_3): if $m', m'' \prec m$ and $m' \neq m''$, then for any $h \in B_m$, we have that $m', m'' \in M_h$ and thus either $(m', h) < (m'', h)$ or $(m', h) > (m'', h)$, whence $m' \prec m''$ or $m'' \prec m'$.

Each M_h is a maximal chain in \prec : it is a chain by (\prec_2), and it is maximal by (\prec_3). Indeed, if $M_h \cup \{m^*\}$ is a chain for some $m^* \notin M_h$, then either $m^* \prec m$ for some $m \in M_h$, which implies by (\prec_3) that $m^* \in M_h$ (a contradiction), or $m \prec m^*$ for all $m \in M_h$. The latter implies, again using (\prec_3) that $M_h \subseteq M_{h^*}$ for any $h^* \in B_{m^*}$ and thus, by (\prec_4), $h = h^*$: noting that $h^* \notin B_m$, another contradiction.

The set $B' = \{M_h : h \in B\}$ is thus a bundle and it is routine to check that this bundled tree satisfies properties B1 – B3 as an immediate consequence of properties BI1 – BI3 above, as well as the fact that $(m, h) \models \phi$ iff $(m, M_h) \models \phi$.

The converse is straightforward: given a choice B-tree (first definition)

$$\mathcal{B}' = (M, \prec, B, \sim_{m,i}, \sim_{m,\text{Agt}}),$$

we define the following relations on the set $\{(m, h) \in M \times B : m \in h\}$:

$$\begin{aligned} (m, h) &\equiv (m', h') && \text{iff } m = m'; \\ (m, h) &\equiv_x (m', h') && \text{iff } m = m' \text{ and } h \sim_{x,m} h', \text{ for } x \in \text{Agt} \cup \{\text{Agt}\}; \\ (m, h) &< (m', h') && \text{iff } h = h' \text{ and } m \prec m'. \end{aligned}$$

It is routine to check that this structure \mathcal{B} satisfies all the properties of the second definition.

Let F be the functor which transforms an indexed choice B-tree \mathcal{B} into a regular choice B-tree \mathcal{B}' , and G the functor which does the opposite, as described above. The reader may check that $FG(\mathcal{B}')$ is isomorphic to \mathcal{B}' , and $GF(\mathcal{B})$ is isomorphic to \mathcal{B} . ■

Until the end of the section, we refer exclusively to the second definition of both choice B-trees and Kamp frames.

We have:

Proposition 4.16. *Every choice B-tree is a Kamp frame (and thus a T-STIT frame).*

Proof. It is routine to check that a tuple satisfying the properties of Def. 4.14 will satisfy the properties of Def. 4.10; left to the reader. ■

The converse need not hold; we can have a Kamp frame which is not isomorphic to a B-tree. However,

Proposition 4.17. *A Kamp frame is isomorphic to a choice B-tree iff it satisfies:*

\langle'_4 . *If $X_w \subseteq X_{w'}$ and $(w, t) \sim (w', t)$ for all $t \in X_w$, then $w = w'$.*

Proof. Let $\mathcal{K} = (X, \sim, \sim_i, \sim_{Agt}, \prec)$ (where $X \subseteq W \times T$) be an (indexed) Kamp frame satisfying the above property. We construct an isomorphic (indexed) choice B-tree $\mathcal{B} = (Y, \equiv, \equiv_i, \equiv_{Agt}, <)$ (where $Y \subseteq B \times M$) as follows:

- Let $B = W$ and let M be the set

$$\{(t, [w, t]) : (w, t) \in X\},$$

where $[w, t]$ represents the equivalence class of (w, t) under the \sim relation.

For shortness, given $(w, t) \in X$, we use t_w to denote the element $(t, [w, t]) \in M$. Note that we have defined this in a way such that

$$t_w = t_{w'} \text{ iff } (w, t) \sim (w', t'). \quad (4.1)$$

- Let $Y = \{(w, t_w) : (w, t) \in X\}$.
- Let $(w, t_w) \equiv (w', t_{w'})$ (resp. \equiv_i, \equiv_{Agt}) iff $(w, t) \sim (w', t')$ (resp. \sim_i, \sim_{Agt}).

Note that by the above observation (4.1), we have that $(w, t_w) \equiv (w', t_{w'})$ iff $t_w = t_{w'}$, satisfying the (EQ) condition of Def. 4.10.

- Finally, let $(w, t_w) < (w', t_{w'})$ iff $(w, t) \prec (w', t')$.

We leave it to the reader to check that \mathcal{B} satisfies the properties of Def. 4.10. The map $(w, t) \mapsto (w, t_w)$ is a bijection; we have defined the relations on \mathcal{B} so that it is an isomorphism.

Conversely, suppose \mathcal{K} is a Kamp frame which does not satisfy condition \langle'_4 , i.e, there are two distinct w, w' such that $X_w \subseteq X_{w'}$ and $(w, t) \sim (w', t)$

for all $t \in X_w$. Suppose towards contradiction that \mathcal{K} is isomorphic to some B-tree \mathcal{B} , via some isomorphism f .

For some t , let (m, h) be the image under this isomorphism of (w, t) , and let (m', h') be $f(w', t)$. Since $(w, t) \sim (w', t)$, we thus have that $(m, h) \equiv (m', h')$, which entails $m = m'$ by (EQ). Since $f(w, t) \neq f(w', t)$, it must then be the case that $h \neq h'$; let us see that $M_h \subseteq M_{h'}$ to reach a contradiction.

Let $n \in M_h$. By property $(<_2)$, we must then have $(n, h) <^* (m, h)$ and thus, by isomorphism, it must be that $f^{-1}(n, h) \prec^* (w, t)$. This means that $f^{-1}(n, h) = (w, t')$ for some $t' \in X_w$. Now, since $t' \in X_w$ and $(w, t') \sim (w', t')$, it must be the case that $f(w', t') = (n, h^*)$ for some h^* . But since the set $\{w'\} \times X_{w'}$ is linearly ordered, we have that $(w', t') \prec^* (w', t)$, whence it must be that $h^* = h'$. Therefore $n \in M_{h'}$: we reach a contradiction. \blacksquare

Now it is almost systematic to use the above result plus the functor defined in Thm. 4.12 to determine necessary and sufficient conditions for a T-STIT frame to be isomorphic to a choice B-tree.

Given a T-STIT frame $\mathcal{X} = (X, \equiv, \equiv_i, \equiv_{Agt}, \prec_G)$ and two sets $A, B \subseteq X$, we say that A is *pseudo-contained* in B , notation $A \sqsubseteq B$, if for all $a \in A$ there exists some $b \in B$ such that $a \equiv b$. We have the following:

Corollary 4.18. *A T-STIT frame is isomorphic to a choice B-tree iff it satisfies, for all $x, y \in X$:*

$$\prec''_4. \text{ If } [x]_{\prec_G^*} \sqsubseteq [y]_{\prec_G^*}, \text{ then } [x]_{\prec_G^*} = [y]_{\prec_G^*}.$$

Proof. Let $\mathcal{K} = (X, <, \sim, \{\sim_i\}_{i \in Agt}, \sim_{Agt})$ be a Kamp indexed frame satisfying $(<''_4)$. Let us see that, considered as a T-STIT frame, it satisfies $(<''_4)$. Suppose $[(w, t)]_{\prec^*} \sqsubseteq [(w', t')]_{\prec^*}$. This means that for all $(w, t_0) \in [(w, t)]_{\prec^*}$, there is some $(w', t'_0) \in [(w', t')]_{\prec^*}$ such that $(w, t_0) \sim (w', t'_0)$. But if this is the case, it must be that $t_0 = t'_0$. Since, for all $t_0 \in X_w$, we have by linearity that $(w, t_0) \in [(w, t)]_{\prec^*}$, the above observation implies that $X_w \subseteq X_{w'}$ and that $(w, t_0) \sim (w', t_0)$ for all $t_0 \in X_w$; by $(<''_4)$, this gives $w = w'$ and thus $[(w, t)]_{\prec^*} = [(w', t')]_{\prec^*}$.

Conversely, let $\mathcal{X} = (X, \equiv_{\square}, \equiv_i, \equiv_{Agt}, \prec_G)$ be a T-STIT frame satisfying $(<''_4)$, let us consider its isomorphic Kamp frame $\mathcal{K}' = (X', \equiv'_{\square}, \equiv'_i, \equiv'_{Agt}, \prec'_G)$, as described in the proof of Thm. 4.12, and let us see that it satisfies the $(<''_4)$ property. Indeed, suppose that $X'_{[x]_{\prec_G^*}} \subseteq X'_{[x']_{\prec_G^*}}$ and that $[x, y] \equiv'_{\square} [x', y]$ for all y such that $[x, y] \in X'$. In particular, take any $y \in [x]_{\prec_G^*}$: since $y \in [x]_{\prec_G^*} \cap [y]_{\equiv_{\square}}$, we have that $[y]_{\equiv_{\square}} \in X'_{[x]_{\prec_G^*}}$, and thus, by assumption, $[y]_{\equiv_{\square}} \in X'_{[x']_{\prec_G^*}}$ and $[x, y] \equiv'_{\square} [x', y]$. By the definition of the isomorphism, this

gives $a_{xy} \equiv_{\square} a_{x'y}$. Recall that a_{xy} is the unique element in the intersection $[x]_{\prec_G^*} \cap [y]_{\equiv_{\square}}$, whence $a_{xy} = y$; on the other hand, $a_{x'y}$ belongs to $[x']_{\prec_G^*}$. We have thus shown that $[x]_{\prec_G^*} \subseteq [x']_{\prec_G^*}$ which, by the (\prec_4'') property, gives that $[x]_{\prec_G^*} = [x']_{\prec_G^*}$, and thus (\prec_4') holds. ■

Discussion. We have given necessary and sufficient conditions for certain classes of general orthogonal frames to be categorically equivalent to classes of models for some of the different logics presented as examples in Section 3.1. The remaining example provided in Section 3.1, namely the framework of Social Epistemic Logic, is the subject of the next chapter.

The authors of [33] make the following remark in their conclusion:

The alternative semantics from Section 2 [referring to T-STIT frames] are appealing because they are convenient mathematical objects. Indeed, the semantics based on [T-STIT frames] quantifies over time-points only, and that based on Kamp agent frames quantifies over times and worlds. These entities are introduced as primitives, not as defined set-theoretical objects. As a consequence, the two semantics keep quantification at first-order level and this in turn allows for the application of convenient techniques, such as Sahlqvist techniques for proving completeness, that do not apply to more common structures for indeterminist time and agency.

We share the hope that these equivalences will assist in applying conventional techniques to some unconventional two-dimensional frameworks.

Some open questions remain:

- There have been a lot of investigations into the logic of *semi-products*, i.e. arbitrary subframes of products (see e.g. [48, Chapter 9]), which have interesting properties somewhat different from products themselves. In the counterpart framework presented in this chapter, semi-products would be orthogonal frames (as opposed to fully orthogonal), with some extra properties which are worth investigating.
- Can the conditions for isomorphism to a subset space frame be simplified?
- Some variations on subset space logics consider families of subsets which are closed under intersection [77] or which are topologies [51, 38, for instance]. What further restrictions does one have to add to obtain a result analogous to Prop. 4.4 for these structures? In the latter case, is there a relation between these properties and the point-free topologies of (e.g.) [85]?

- Can more STIT frameworks be generalised to be seen as orthogonal frames?

Another case study: Social Epistemic Logic

We return to the framework of SEL, which was introduced in [93] within a proposal for a multi-modal framework called ‘Epistemic Logic of Friendship’ (and later ‘Social Epistemic Logic’), allowing for both an epistemic accessibility relation and a ‘friendship’ relation.

The models for SEL have the shape of an indexed frame, although their use of logical machinery is broader. The set of agents is encoded in the semantics, and these agents are named using nominal variables (a notion borrowed from hybrid logic) with the novelty that these nominals only refer to the elements of one of the sets. Sano [90] provided an axiomatisation for a fragment of the language. We give a simplified proof of this result and we axiomatise an extension of this fragment. Results on decidability and the Finite Model Property are provided.

We conclude with a small study on orthogonal structures which are isomorphic to these models.¹



Good old Ulises was a ticking bomb, and what was worse, socially speaking, was that everyone knew or could sense that he was a ticking bomb and no one wanted him to get too close, for obvious and forgivable reasons.

Roberto Bolaño, *The Savage Detectives* (1998).

I N THE PRESENT CHAPTER we come back to the framework of Social Epistemic Logic, briefly alluded to in Example 3.4.

This framework, introduced by Seligman, Liu and Girard in [92],² uses an (IF2) indexed frame as a basis for modelling the epistemic state of agents in

¹Many of the results in this chapter (except for those in the novel Section 5.5) were originally published in:

Philippe Balbiani and Saúl Fernández González. Indexed frames and hybrid logics. In *Advances in Modal Logic*, 2020,

although this chapter extends and improves on those results significantly.

²In this introductory paper the framework is referred to as ‘Epistemic Friendship Logic’. It was later re-baptised by [105] as ‘Social Epistemic Logic’ or SEL; here we stick to the latter nomenclature.

a social network as well as their social connections. We shall reintroduce and expand on this framework below.

Let us start this recapitulation with the most basic bimodal language \mathcal{L} , defined as:

$$\phi ::= p \mid \perp \mid \neg\phi \mid (\phi \wedge \phi) \mid K\phi \mid F\phi,$$

where $p \in \mathbf{Prop}$, a countable set of propositional variables. These propositional variables are meant to represent indexical propositions related to a certain agent (to use the example given in [92], p can mean ‘I am in danger’). K is meant to be read as an epistemic modality (*I know p*), whereas F is a ‘friendship’ modality (*all my friends p*). We use \hat{K} and \hat{F} as the duals of these operators, i.e. $\hat{K}\phi := \neg K\neg\phi$ and $\hat{F}\phi := \neg F\neg\phi$.

Models consist of indexed frames along with a valuation, i.e., tuples of the form (W, A, \sim, \succ, V) , where W and A are nonempty sets (“states” and “agents”, respectively), $\sim = \{\sim_a : a \in A\}$ is a family of binary relations on W indexed by A ($\sim_a \subseteq W^2$ represents agent a ’s epistemic accessibility), and $\succ = \{\succ_w : w \in W\}$ is a family of binary relations on A indexed by W (each representing which agents are ‘friends’ at world w). $V : \mathbf{Prop} \rightarrow 2^{W \times A}$ is a valuation.

We interpret formulas of \mathcal{L} with respect to pairs $(w, a) \in W \times A$, as follows:

$$\begin{aligned} (w, a) \models K\phi & \text{ iff } (v, a) \models \phi \text{ for all } v \text{ such that } w \sim_a v; \\ (w, a) \models F\phi & \text{ iff } (w, b) \models \phi \text{ for all } b \text{ such that } a \succ_w b. \end{aligned}$$

To illustrate this, see the diagram in Figure 5.1. It represents a situation with three agents, Alice, Bob and Charlie, wherein at world w Alice has a friend with the property p (represented by the grey nodes) yet she does not know that:




Figure 5.1: A model for Social Epistemic Logic.

Indeed, it holds that $w, a \models \hat{F}p \wedge \neg K\hat{F}p$. We could also express more complex things such as “Alice does not know Bob and Charlie are friends”. In order to do this, we would need to extend the language.

Naming the agents. Let us go back to the notion “Alice does not know Bob and Charlie are friends”. In order to express this in our language, we need to name the agents. This is done in [92] via the introduction of *nominal variables* and a modality $@_n$, directly imported from hybrid logic: see [4, 30, 49, 83]. The language $\mathcal{L}(@)$ extends \mathcal{L} with the atom n and the operator $@_n\phi$, where n belongs to \mathbf{Nom} , a countable set of nominal variables.

Definition 5.1. Models for $\mathcal{L}(@)$ are of the shape $\mathfrak{M} = (W, A, \sim, \succ, V)$, where (W, A, \sim, \succ) is an indexed frame and $V : \mathbf{Prop} \cup \mathbf{Nom} \rightarrow 2^{W \times A}$ is a valuation function with the property that, for each $n \in \mathbf{Nom}$, $V(n) = W \times \{a\}$ for some $a \in A$. We refer to this unique a as $a = \underline{n}_V$ (or $a = \underline{n}$ if there is no risk of ambiguity).

A model is *named* whenever, for each $a \in A$, there exists $n \in \mathbf{Nom}$ such that $\underline{n} = a$. (Note that, in a named model, A is at most countable.)

We interpret formulas of $\mathcal{L}(@)$ in named models with respect to pairs consisting of worlds $w \in W$ and agents $a \in A$ as follows:

$$\begin{aligned} (w, a) \models n & \quad \text{iff } (w, a) \in V(n) \\ & \quad (\text{iff } \underline{n} = a); \\ (w, a) \models @_n\phi & \quad \text{iff } M, w, \underline{n} \models \phi \\ & \quad (\text{iff } \phi \text{ is true of the agent named by } n). \end{aligned}$$

A complete axiomatisation of $\mathcal{L}(@)$ was provided for the first time by Sano in [90]. The proof of completeness works (roughly) as follows: first, a cut-free tree sequent calculus is introduced, which is then shown to be sound and complete. Then Sano shows that a formula which is provable in the Hilbert-style system can be converted into a provable tree sequent and, conversely, that from a provable tree sequent one can obtain a formula which is derivable in the Hilbert-style system.

In the conclusion of [90] it is suggested that finding a proof of this result using canonical models is an interesting area of future research. We present such a proof in Section 5.1, along with a proof that the logic possesses the finite model property (Section 5.2).

Back to friendship logic. For most of this chapter we ignore many of the constraints imposed in [92] upon the models in order to make them a realistic framework for a logic of knowledge and friendship, namely: the set of agents A should be finite, the epistemic relations \sim_a should be equivalence relations, the friendship relations \succ_w should be symmetric and irreflexive, and, optionally, it should be the case that an agent always knows who her friends are (if $w \sim_a v$ and $a \succ_w b$, then $a \succ_v b$). We address these properties

in Section 5.3 and use all the previous results to provide a logic for the exact class of models proposed in [92]. (It is worth noting that, although we are currently sticking to the \sim and \asymp symbols to maintain the notation of [92, 90], the reader should not assume until Section 5.3 that they denote equivalence or symmetric relations.)

Another extension. Another operator from hybrid logic is considered in [92]. The operator $\downarrow x.\phi$ allows to name the current agent x , making it possible to refer to it indexically. The resulting extension of $\mathcal{L}(@)$, let us call it $\mathcal{L}(@\downarrow)$, allows to express things like “I have a friend who knows n is friends with me”, $\downarrow x.\hat{F}K@_n\hat{F}x$. In Section 5.4 we provide a sound and complete axiomatization for $\mathcal{L}(@\downarrow)$.

For the time being, and going into the next section, we stick to the simpler extension $\mathcal{L}(@)$: this amounts to considering a set $\text{Nom} = \{n, m, \dots\}$ of nominal variables to our language and considering the language:

$$\phi ::= p|n|\perp|\neg\phi|(\phi \wedge \phi)|K\phi|F\phi|@_n\phi,$$

where $p \in \text{Prop}$, $n \in \text{Nom}$.

At the end of this chapter (Section 5.5) we provide a semantics for this extension based on orthogonal structures, as opposed to indexed frames.

5.1 Axiomatising $\mathcal{L}(@)$ via canonical models

It is proven in [90], via an argument that employs a tree sequent calculus, that the logic of $\mathcal{L}(@)$ is the system **SEL**, defined in the table below:

(Taut)	all propositional tautologies	(MP)	from ϕ and $\phi \rightarrow \psi$, infer ψ
(K_K)	$K(\phi \rightarrow \psi) \rightarrow (K\phi \rightarrow K\psi)$	(Nec $_K$)	from ϕ , infer $K\phi$
(K_F)	$F(\phi \rightarrow \psi) \rightarrow (F\phi \rightarrow F\psi)$	(Nec $_F$)	from ϕ , infer $F\phi$
($K_@$)	$@_n(\phi \rightarrow \psi) \rightarrow (@_n\phi \rightarrow @_n\psi)$	(Nec $_@$)	from ϕ , infer $@_n\phi$
(Ref)	$@_nn$	(Selfdual)	$\neg @_n\phi \leftrightarrow @_n\neg\phi$
(Elim)	$@_n\phi \rightarrow (n \rightarrow \phi)$	(Agree)	$@_n@_m\phi \rightarrow @_m\phi$
(Back)	$@_n\phi \rightarrow F@_n\phi$	(DCom)	$@_nK@_n\phi \leftrightarrow @_nK\phi$
(Rigid $_=$)	$@_nm \rightarrow K@_nm$	(Rigid $_{\neq}$)	$\neg @_nm \rightarrow K\neg @_nm$
(Name)	From $@_n\phi$ infer ϕ , where n is fresh in ϕ		
(LBG)	From $L(@_n\hat{F}m \rightarrow @_m\phi)$ infer $L(@_nF\phi)$, m fresh in $L(@_nF\phi)$.		

In the last line of the above table, the *necessity forms* $L(\#)$ are defined as:

$$L ::= \#|\phi \rightarrow L|@_nKL,$$

and $L(\psi)$ is the result of substituting the unique occurrence of the symbol $\#$ in $L(\#)$ by the formula ψ .

In this section we present a novel proof of this result using canonical models. To do this, we consider instead the logic SEL^+ , obtained by replacing the rule (LBG) in SEL by the following infinitary rule:

$$\text{(LBG}^+) \quad \text{From } L(@_n \hat{F}m \rightarrow @_m \phi) \text{ for all } m \text{ fresh in } L(@_n F\phi), \\ \text{infer } L(@_n F\phi).$$

The following Lemma can be proven by a straightforward induction on derivations:

Lemma 5.2. *SEL and SEL^+ prove the same formulas.*

We thus prove completeness of SEL^+ . The following validities will be useful:

Proposition 5.3. *The following are derivable in SEL:*

- (T1) $\vdash @_m @_n \phi \leftrightarrow @_n \phi$;
- (T2) $\vdash n \rightarrow (@_n \phi \leftrightarrow \phi)$;
- (T3) $\vdash @_n m \rightarrow (@_n \phi \leftrightarrow @_m \phi)$;
- (T4) $\vdash @_n m \leftrightarrow @_m n$;
- (T5) $\vdash @_n (\phi \rightarrow \psi) \leftrightarrow (@_n \phi \rightarrow @_n \psi)$;
- (T6) $\vdash @_n m \rightarrow (\phi[k/n] \leftrightarrow \phi[k/m])$, where $\phi[k/n]$ is the formula obtained from ϕ by replacing each occurrence of k by n .
- (T7) $\vdash @_n m \rightarrow @_i K @_n m$, and $\vdash @_n \neg m \rightarrow @_i K @_n \neg m$;
- (T8) $\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_n \hat{F}\phi$;
- (T9) $\vdash @_n F\psi \wedge @_n \hat{F}m \rightarrow @_m \psi$;
- (R1) if $\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow \psi$, then $\vdash @_n \hat{F}\phi \rightarrow \psi$, with $m \neq n$ fresh in ϕ and ψ .

Proof. (T1) to (T6) are proven in Prop. 3 of [90] and Lemma 2 of [30].

$$\begin{array}{l} \text{(T7)} \quad \vdash @_n m \rightarrow @_i K @_n m. \\ \hline \vdash @_n m \rightarrow K @_n m \quad (\text{Rigid}_=) \\ \vdash @_i @_n m \rightarrow @_i K @_n m \quad (\text{K}_@ + \text{Nec}_@) \\ \vdash @_n m \rightarrow @_i K @_n m \quad (\text{T1}) \end{array}$$

The derivation of $\vdash @_n \neg m \rightarrow @_i K @_n \neg m$ is identical but using ($\text{Rigid}_\neq + \text{Selfdual}$) in the first step.

(T8) $\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_n \hat{F}\phi.$	
$\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_n \hat{F}m \wedge F@_m \phi$	(Back)
$\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_n \hat{F}m \wedge @_n F@_m \phi$	(Nec _@ +K _@ +T1)
$\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_n \hat{F}(m \wedge @_m \phi)$	(by modal reasoning: $\Box A \wedge \Diamond B \rightarrow \Diamond(A \wedge B)$)
$\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_n \hat{F}\phi$	(by T2: $\vdash m \wedge @_m \phi \rightarrow \phi$)
(T9) $\vdash @_n F\psi \wedge @_n \hat{F}m \rightarrow @_m \psi.$	
$\vdash @_n F\psi \wedge @_n \hat{F}m \rightarrow @_n \hat{F}(m \wedge \psi)$	(modal reasoning: $\Box A \wedge \Diamond B \rightarrow \Diamond(A \wedge B)$)
$\vdash m \wedge \psi \rightarrow @_m \psi$	(T2)
$\vdash @_n F\psi \wedge @_n \hat{F}m \rightarrow @_n \hat{F}@_m \psi$	(two above lines)
$\vdash \hat{F}@_m \psi \rightarrow @_m \psi$	(dual of Back)
$\vdash @_n F\psi \wedge @_n \hat{F}m \rightarrow @_m \psi$	(two above lines plus (T1))

Before showing (R1), let us show this rule:

(Name') If $\vdash \phi \rightarrow @_m \psi$ and m is fresh, then $\vdash \phi \rightarrow \psi.$	
$\vdash \phi \rightarrow @_m \psi$	(Premise)
$\vdash @_m \phi \rightarrow @_m @_m \psi$	(Nec _@ +K _@)
$\vdash @_m \phi \rightarrow @_m \psi$	(Agree)
$\vdash @_m(\phi \rightarrow \psi)$	(T5)
$\vdash \phi \rightarrow \psi$	(Name)

With this:

(R1) If $\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow \psi,$ then $\vdash @_n \hat{F}\phi \rightarrow \psi,$ with $m \neq n$ fresh in ϕ and $\psi.$	
$\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow \psi$	(Premise)
$\vdash @_i @_n \hat{F}m \wedge @_i @_m \phi \rightarrow @_i \psi$	(Nec _@ +K _@ , i fresh)
$\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_i \psi$	(T1)
$\vdash @_n \hat{F}m \wedge @_m \phi \rightarrow @_m @_i \psi$	(Nec _@ +K _@ +T1)
$\vdash @_n \hat{F}m \rightarrow @_m(\phi \rightarrow @_i \psi)$	(T5)
$\vdash @_n F(\phi \rightarrow @_i \psi)$	(BG)
$\vdash @_n \hat{F}\phi \rightarrow @_n \hat{F}@_i \psi$	($\Box(A \rightarrow B) \rightarrow (\Diamond A \rightarrow \Diamond B)$)
$\vdash @_n \hat{F}\phi \rightarrow @_n @_i \psi$	(Back)
$\vdash @_n \hat{F}\phi \rightarrow @_i \psi$	(T1)
$\vdash @_n \hat{F}\phi \rightarrow \psi$	(Name')

■

We will say that a formula in $\mathcal{L}(@)$ is a *named formula* whenever it is of the form $@_n\phi$. A *BCN formula* is a Boolean combination of named formulas, and we use *BCN* to denote the set of such formulas. The following is an immediate consequence of (T1), (T5) and (Selfdual):

Corollary 5.4. *If $\phi \in BCN$, $n \in \text{Nom}$, then $\vdash @_n\phi \leftrightarrow \phi$.*

A formula ϕ is *consistent* if $\neg\phi$ is not derivable. The following lemma will be useful later.

Lemma 5.5. *If n does not occur in ϕ , then ϕ is consistent if and only if $@_n\phi$ is consistent.*

Proof. If ϕ is inconsistent we have $\vdash \neg\phi$ and thus by (Nec_@), $\vdash @_n\neg\phi$, which by (Selfdual) gives that $\vdash \neg @_n\phi$. If $@_n\phi$ is inconsistent then $\vdash \neg @_n\phi$ which by (Selfdual) means $\vdash @_n\neg\phi$ and thus, by (Name), $\vdash \neg\phi$. ■

Now we can start our completeness proof. The two above results allow us to focus only on BCN formulas. A *theory* is a set of BCN formulas T such that:

- i. $\text{SEL}^+ \cap BCN \subseteq T$;
- ii. T is closed under Modus Ponens;
- iii. If $L(@_n\hat{F}m \rightarrow @_m\phi) \in T$ for all $m \neq n$ not occurring in L or in ϕ , then $L(@_nF\phi) \in T$.

A theory is *consistent* whenever $@_n\perp \notin T$ for all n .³ It is easy to see that $\text{SEL}^+ \cap BCN$ is the least consistent theory. A consistent theory is *maximal* if no proper superset of it is a consistent theory.

Lemma 5.6. *Given a theory T and $n \in \text{Nom}$, the set*

$$T_{K_n} := \{\psi \in BCN : \vdash \psi \leftrightarrow @_n\phi \text{ for some } @_nK\phi \in T\}$$

is a theory.

³Note that if $@_n\perp \in T$ for some n , then $@_m\perp \in T$ for all m , on account of the fact that $\vdash @_n\perp \leftrightarrow @_m\perp$.

Proof. Note the following: for any $\phi \in BCN$, we have that $\phi \in T_{K_n}$ iff $@_n K\phi \in T$. Indeed, if $\phi \in T_{K_n}$, then $\vdash \phi \leftrightarrow @_n \psi$ for some $@_n K\psi \in T$. But then, using (Nec_K), (Nec_@) and (DCom) in that order we obtain $\vdash @_n K\phi \leftrightarrow @_n K\psi$, and thus $@_n K\phi \in T$. The other direction is trivial and uses that $\vdash @_n \phi \leftrightarrow \phi$. With this:

Condition i. If $\phi \in \text{SEL}^+ \cap BCN$ and $n \in \text{Nom}$ then $@_n K\phi \in \text{SEL}^+ \cap BCN$ (by applying two Nec rules) and thus $@_n K\phi \in T$, so $\phi \in T_{K_n}$.

Condition ii. If ϕ and $\phi \rightarrow \psi \in T_{K_n}$, then $@_n K\phi, @_n K(\phi \rightarrow \psi) \in T$ and, by applying the K axioms and modus ponens, $@_n K\psi \in T$, and thus $\psi \in T_{K_n}$.

Condition iii. If $L(@_k \hat{F}m \rightarrow @_m \phi) \in T_{K_n}$ for all fresh m , then $@_n KL(@_k \hat{F}m \rightarrow @_m \phi) \in T$ for all fresh m , and thus, since $@_n KL$ is a necessity form, $@_n KL(@_k F\phi) \in T$, whence $L(@_k F\phi) \in T_{K_n}$. ■

Lemma 5.7. *Given a theory T and a formula $\phi \in BCN$, the set*

$$T_\phi := \{\psi \in BCN : \phi \rightarrow \psi \in T\}$$

is a theory containing T and including the formula ϕ , and it is consistent whenever T is consistent and $\neg\phi \notin T$.

Proof. Condition i. If $\psi \in \text{SEL}^+ \cap BCN$, then $\phi \rightarrow \psi \in \text{SEL}^+ \cap BCN$, thus $\psi \in T_\phi$.

Condition ii. Follows from classical propositional logic.

Condition iii. Follows from the fact that, if L is a necessity form, so is $\phi \rightarrow L$.

The fact that $\phi \in T_\phi \supseteq T$ is because $\vdash \phi \rightarrow \phi$ and $\vdash \psi \rightarrow (\phi \rightarrow \psi)$. If $\neg\phi \notin T$, then $@_n \neg\phi \notin T$, thus $@_n(\phi \rightarrow \perp) \notin T$. Using the K axiom and $\vdash \phi \leftrightarrow @_n \phi$, we obtain $\phi \rightarrow @_n \perp \notin T$, and thus $@_n \perp \notin T_\phi$. ■

Now,

Lemma 5.8 (Lindenbaum's lemma). *A consistent theory can be extended to a maximal consistent theory.*

Proof. Let T_0 be a consistent theory and $(\phi_k)_{k \in \omega}$ be an enumeration of BCN where each formula occurs infinitely many times.

Given a consistent theory T_k , we define a consistent theory T_{k+1} (which extends T_k) as follows:

- If $\neg\phi_k \notin T_k$, then $T_{k+1} = (T_k)_{\phi_k}$.
- If $\neg\phi_k \in T_k$, then:

- If $\neg\phi_k$ is of the form $\neg L(@_n F\phi)$, then for some fresh m it must be the case that $L(@_n \hat{F}m \rightarrow @_m\phi) \notin T_k$, for otherwise we would have by Condition iii. that $L(@_n F\phi) \in T_k$, contradicting its consistency. Then we set $T_{k+1} = (T_k)_{\neg L(@_n \hat{F}m \rightarrow @_m\phi)}$.
- Otherwise, $T_{k+1} = T_k$.

Let $T = \bigcup_{k \in \omega} T_k$. Then T is a maximal consistent theory. Consistency is obvious, for each T_k is consistent. Maximality comes from the fact that, for every formula ϕ_k , either $\neg\phi_k$ was already in T_k , or ϕ_k was added to T_{k+1} , therefore it cannot have consistent supersets closed under modus ponens. To see that it is a theory, it suffices to check that Condition iii. is satisfied. And indeed, if $L(@_n F\phi) \notin T$, then $\neg L(@_n F\phi) \in T_k$ for some k . Consider some $k' > k$ such that $\phi_{k'} = \neg L(@_n F\phi)$. Then, by construction, $T_{k'+1}$ must contain $\neg L(@_n \hat{F}m \rightarrow @_m\phi)$ for some fresh m , and therefore it is not the case that $L(@_n \hat{F}m \rightarrow @_m\phi) \in T$ for all fresh m . ■

Let MCT denote the set of maximal consistent theories. Given $T, S \in MCT$, and $n \in \text{Nom}$, we define: $T \sim_n S$ iff $T_{K_n} \subseteq S$.

Lemma 5.9 (Diamond Lemma). *Let $T \in MCT$. We have:*

- i. *If $@_n \hat{K}\phi \in T$, then there exists $S \in MCT$ such that $T \sim_n S$ and $@_n\phi \in S$.*
- ii. *If $@_n \hat{F}\phi \in T$, then there is some $m \neq n$ fresh in ϕ such that $@_n \hat{F}m \wedge @_m\phi \in T$.*

Proof. i. Take the consistent theory $(T_{K_n})_{@_n\phi}$ and extend it to the desired successor using Lindenbaum's lemma. Note that T_{K_n} is consistent, for if not, $@_n\perp \in T_{K_n}$, and thus $@_n K @_n\perp \in T$. But, since $@_n\perp$ is equivalent to \perp , this means that $@_n K \perp \in T$, contradicting $@_n \hat{K}\phi \in T$. Note moreover that $\neg @_n\phi \notin T_{K_n}$, for if that was the case, $@_n K \neg @_n\phi \in T$, which is equivalent to $\neg @_n \hat{K}\phi \in T$: contradiction. Thus $(T_{K_n})_{@_n\phi}$ is consistent.

ii. If $@_n \hat{F}m \wedge @_m\phi \notin T$ for all fresh m , then $\neg(@_n \hat{F}m) \vee \neg(@_m\phi) \in T$ for all fresh m , and thus, by logical equivalence, $@_n \hat{F}m \rightarrow @_m\neg\phi \in T$ for all fresh m , which entails $@_n F\neg\phi \in T$, and therefore $\neg @_n \hat{F}\phi \in T$. ■

Lemma 5.10. *Let $i \in \text{Nom}$. If $T \sim_i S$ then, for any $n, m \in \text{Nom}$, we have: $@_n m \in T$ if and only if $@_n m \in S$.*

Proof. By (T7) of Prop. 5.3: if $@_n m \in T$, then $@_i K @_n m \in T$, which entails $@_i @_n m \in S$, and therefore, by the (Agree) axiom, $@_n m \in S$. If $@_n m \notin T$, by maximal consistency and the (Selfdual) axiom we have that $@_n \neg m \in T$ and we can proceed similarly to obtain that $@_n \neg m \in S$ and thus $@_n m \notin S$. ■

Let ϕ_0 be a consistent formula and let us build a model satisfying it. Take a nominal n_0 not occurring in ϕ_0 and note that $@_{n_0}\phi_0$ is a consistent BCN formula (by Lemma 5.5) and thus the consistent theory $(BCN \cap \text{SEL}^+)_{@_{n_0}\phi_0}$ can be extended (by Lindenbaum's lemma) to $T_0 \in MCT$.

Let W be the set of elements reachable from T_0 by the \sim_n relations, i.e.

$$W = \{S \in MCT : T_0 = S_0 \sim_{n_1} S_1 \sim_{n_2} \dots \sim_{n_k} S_k = S \\ \text{for some } n_1, \dots, n_k \in \text{Nom}, S_0, \dots, S_k \in MCT\}.$$

Note that this construction guarantees (by Lemma 5.10) that for any $T \in W$, $@_n m \in T$ iff $@_n m \in T_0$. Note moreover that the theorems

$$\vdash @_n n \text{ (Ref)}; \vdash @_n m \leftrightarrow @_m n \text{ (T4)}; \vdash @_n m \wedge @_m i \rightarrow @_n i \text{ (conseq. of T3)}$$

guarantee that the binary relation on **Nom** defined as $n \equiv m$ iff $@_n m \in T_0$ is an equivalence relation. Let $[n]$ denote the equivalence class of $n \in \text{Nom}$ and let $A = \{[n] : n \in \text{Nom}\}$.

For $[n] \in A$, we define $\sim_{[n]} = \sim_n$. Let us see that this is well-defined, which amounts to showing that $\sim_n = \sim_m$ whenever $n \equiv m$. But given $T, S \in W$, and $n \equiv m$, the fact that $@_n m \in T \cap S$ paired with (T3) give us that, for all formulas ϕ , $@_n K\phi \in T$ iff $@_m K\phi \in T$, and $@_n \phi \in S$ iff $@_m \phi \in S$, which entails $T \sim_n S$ iff $T \sim_m S$.

For $T \in W$ we define

$$[n] \succ_T [m] \text{ iff } @_n \hat{F}m \in T.$$

Let us see that this definition does not depend on the choice of representative for the equivalence classes: suppose $@_n \hat{F}m \in T$ and take $n' \in [n], m' \in [m]$. We have that $@_{n'} \hat{F}m \in T$, by (T3), and therefore, by (T6), $@_{n'} \hat{F}m' \in T$.

Finally we define a valuation by setting

$$V(p) = \{(T, [n]) \in W \times A : @_n p \in T\}, \quad p \in \text{Prop}; \\ V(n) = \{(T, [n]) : T \in W\}, \quad n \in \text{Nom}.$$

Note that we have defined V so that $\underline{n} = [n]$. We have that

$$\mathfrak{M}^C = (W, A, \sim_{[n] \in A}, \succ_{T \in W}, V)$$

is a named model and, moreover:

Lemma 5.11 (Truth Lemma). *For any formula $\phi \in \mathcal{L}(@)$, it is the case that $\mathfrak{M}^C, T, [n] \models \phi$ if and only if $@_n \phi \in T$.*

Proof. By induction on ϕ . For the case $\phi = m \in \text{Nom}$ we recall that $\underline{m} = [m]$. For the case $\phi = K\psi$, we use the Diamond Lemma. For the case $\phi = F\psi$, we use the Diamond Lemma for one direction and (T9) for the other. ■

With this:

Theorem 5.12. SEL^+ (and therefore SEL) is complete with respect to the class of (not necessarily finite) named indexed models.

Proof. If ϕ_0 is consistent, so is $@_{n_0}\phi_0$ for n_0 not occurring in ϕ_0 , and thus we can construct \mathfrak{M}^C as above and we have that $\mathfrak{M}^C, T_0, [n_0] \models \phi_0$. ■

5.2 Finite models

The following also holds:

Theorem 5.13. SEL is complete with respect to the class of finite named indexed models.

Proof. This amounts to showing that, if a formula ϕ_0 is satisfied in a model $M = (W, A, \sim, \succ, V)$, then there is a finite submodel which satisfies it. Suppose $M, w_0, a_0 \models \phi_0$.

We define $\text{nom } \phi_0$ to be the (finite) set of nominal variables occurring in ϕ_0 . We define relations R , S and A_n (for $n \in \text{Nom}$) on $W \times A$ as follows:

$$\begin{aligned} (w, a)\text{R}(w', a') & \text{ iff } a = a' \text{ and } w \sim_a w', \\ (w, a)\text{S}(w', a') & \text{ iff } w = w' \text{ and } a \succ_w a', \\ (w, a)\text{A}_n(w', a') & \text{ iff } w = w' \text{ and } a' = \underline{n}. \end{aligned}$$

We will consider chains starting at (w_0, a_0) , of the form

$$\alpha = (w_0, a_0) \xrightarrow{\text{T}_1} (w_1, a_1) \dots \xrightarrow{\text{T}_k} (w_k, a_k),$$

with $k \geq 0$, $\text{T}_i \in \{\text{R}, \text{S}, \text{A}_n : n \in \text{Nom}\}$ and $(w_{i-1}, a_{i-1})\text{T}_i(w_i, a_i)$ for $1 \leq i \leq k$. We shall say that such a chain has *length* k (and thus (w_0, a_0) is a chain of length 0). We will call $\text{last } \alpha = (w_k, a_k)$.

Given a formula ϕ , we let $\text{mod } \phi$ be the total number of K , F and $@_n$ modalities occurring in ϕ , and we let $N = \text{mod } \phi_0$.

We shall construct a finite set of chains of length up to N , in N steps. Let $F_0 = \{(w_0, a_0)\}$. For $0 \leq k \leq N - 1$, suppose F_k is a finite set of chains of length k . Let F_{k+1} be a finite set of minimal cardinality satisfying the following property for all $\alpha \in F_k$ and all $\text{T} \in \{\text{R}, \text{S}, \text{A}_n : n \in \text{nom } \phi_0\}$:

for any $(w, a) \in W \times A$, if $(\text{last } \alpha)\top(w, a)$, then there exists an element $(w', a') \sim_{\phi_0} (w, a)$ such that $\alpha \xrightarrow{\top} (w', a') \in F_{k+1}$,

where \sim_{ϕ_0} is the equivalence relation

$(w, a) \sim_{\phi_0} (w', a')$ iff
for all $\psi \in \text{subf } \phi_0$, $M, w, a \models \psi$ iff $M, w', a' \models \psi$,

and $\text{subf } \phi_0$ is, naturally, the set of subformulas of ψ .

It is not hard to see that there is a set of cardinality at most

$$(2 + N) \cdot |F_k| \cdot 2^{|\text{subf } \phi_0|}$$

satisfying this property. Indeed, for any of the $|F_k|$ choices of α and $2 + N$ choices of \top , F_{k+1} will contain an element $\alpha \xrightarrow{\top} (w, a)$ for (at most) one representative of each of the (at most) $2^{|\text{subf } \phi_0|}$ equivalence classes of \sim_{ϕ_0} .

Let F be the closure of $F_0 \cup \dots \cup F_n$ under the following property:

if $\alpha \in F$, $\text{length}(\alpha) < N$, $w \in W$ and $a \in A$ occur in F , and $(\text{last } \alpha)\top(w, a)$, then $\alpha \xrightarrow{\top} (w, a) \in F$.

Obviously, $F_0 \cup \dots \cup F_n$ is finite, and so is F .

We construct our finite model $M^f = (W^f, A^f, R^f, S^f, V^f)$ where W^f and A^f are the restrictions of W and A to those elements occurring in F , i.e.,

$$W^f = \{w \in W : w \text{ occurs in } F\}; \quad A^f = \{a \in A : a \text{ occurs in } F\};$$

and R^f , S^f and V^f are the corresponding restrictions of R , S , and V . The following holds:

Lemma 5.14. *Let $\alpha \in F$ be a chain of length k , i.e.,*

$$\alpha = (w_0, a_0) \xrightarrow{\top_1} (w_1, a_1) \dots \xrightarrow{\top_k} (w_k, a_k),$$

with $\top_i \in \{R, S, A_n : n \in \text{nom } \phi_0\}$. Let ϕ be a subformula of ϕ_0 such that $\text{mod } \phi \leq N - k$. Then, $M, w_k, a_k \models \phi$ if and only if $M^f, w_k, a_k \models \phi$.

This Lemma is proven below and it finishes the proof of our theorem, for it now suffices to apply it to the chain (w_0, a_0) of length 0 to obtain $M^f, w_0, a_0 \models \phi_0$. ■

Proof of Lemma 5.14. By induction on ϕ . The cases for $\phi = p$, $\phi = n$ and $\phi = \perp$ are trivial, and so is the inductive step for $\phi = \neg\psi$.

Case $\phi = \psi_1 \wedge \psi_2$. If $M, w_k, a_k \models \psi_1 \wedge \psi_2$, then $M, w_k, a_k \models \psi_i$ for $i = 1$ and 2 . But then, since $\text{mod } \psi_i \leq \text{mod } \phi \leq N - k$, we have by induction hypothesis that $M^f, w_k, a_k \models \psi_i$ and thus $M^f, w_k, a_k \models \phi$. The converse is analogous.

Case $\phi = K\psi$. Suppose that $M^f, w_k, a_k \models K\psi$ and take w such that $w_k \sim_{a_k} w$. Note that $k < N$ because $N - k \geq \text{mod } K\psi > 0$, and thus F_{k+1} is defined and contains an element $\alpha \xrightarrow{R} (w_{k+1}, a_{k+1})$ such that $a_{k+1} = a_k$, $w_k \sim_{a_k} w_{k+1}$ (and therefore $w_k \sim_{a_k}^f w_{k+1}$) and $(w_{k+1}, a_{k+1}) \sim_{\phi_0} (w, a_k)$. We have that $M^f, w_{k+1}, a_{k+1} \models \psi$ and, since

$$N - (k + 1) = N - k - 1 \geq \text{mod}(K\psi) - 1 = \text{mod } \psi,$$

induction hypothesis gives us that $M, w_{k+1}, a_{k+1} \models \psi$. By the \sim_{ϕ_0} relation, this means that $M, w, a_k \models \psi$, and we have thus proven that $M, w_k, a_k \models K\psi$.

Conversely, suppose $M, w_k, a_k \models K\psi$ and $R_{a_k}^f w_k w$. In this case we have that $w_k \sim_{a_k} w$ and thus $M, w, a_k \models \psi$. Since $\alpha \xrightarrow{R} (w, a_k) \in F$ and its length is $k + 1$, and since $N - (k + 1) \geq \text{mod } \psi$, the induction hypothesis applies and we have that $M^f, w, a_k \models \psi$. This entails $M^f, w_k, a_k \models K\psi$.

The cases for $\phi = F\psi$ and $\phi = @_n\psi$ are completely analogous to the above reasoning, using S and A_n in place of R. \blacksquare

5.3 Extensions of SEL

In [92] some assumptions are made about the epistemic and social relations in the models. The epistemic relations \sim_a are equivalence relations, whereas the friendship relation \succ_w is irreflexive and symmetric.

One would expect, for instance, that if the relations \sim_a that give the semantics of the knowledge modality K are reflexive, transitive and symmetric, then this modality should follow the axioms of S5, namely:

$$\vdash K\phi \rightarrow \phi; \quad \vdash K\phi \rightarrow KK\phi; \quad \vdash \phi \rightarrow K\neg K\neg\phi.$$

Similarly, if \sim_a is a preorder, the extra axioms of S4 (i.e. the first two above), should be included to the logic. Let $\text{SEL} + \text{S5}_K$ denote the logic resulting from adding these three axioms to SEL, and let $\text{SEL} + \text{S4}_K$ be the logic resulting from adding the first two. And indeed:

Theorem 5.15 ([90]). *SEL \oplus S5_K is sound and complete with respect to the class of models where the \sim_a are equivalence relations. Moreover, SEL \oplus S4_K*

is sound and complete with respect to the class of models where each \sim_a is a preorder.

The proof of this result in [90] consists in adding corresponding rules to the tree sequent calculus and showing that a provable formula in the Hilbert-style system can be transformed into a provable sequent and vice versa. With the canonical models presented in this text this proof becomes quite straightforward. First, note that thanks to (T5) the following are easily provable in $\text{SEL} + \text{S5}_K$ (and the first two in $\text{SEL} + \text{S4}_K$):

$$\vdash @_n K\phi \rightarrow @_n \phi; \quad \vdash @_n K\phi \rightarrow @_n KK\phi; \quad \vdash @_n \phi \rightarrow @_n K\neg K\neg\phi.$$

With this, the proof of the following lemma is straightforward:

Lemma 5.16. *If the axioms of S5 for K (resp. S4) are present in the logic, each relation \sim_n in the canonical model is an equivalence relation (resp. a preorder).*

Remark 5.17. Given that $@_n$ distributes over $\rightarrow, \wedge, \vee, \neg$, one can see that there are many examples of formulas ϕ defining a certain frame property from which it is trivial to compute a formula $@_n\psi$ defining the same property in the \sim_n relations of indexed frames. Some obvious questions arise: is this true of any Sahlqvist formula? Can we adapt the notion of Sahlqvist formula to this setting and prove an analogue of the Sahlqvist Completeness Theorem ([28, Thm. 4.42])? We conjecture the answer is affirmative.

Similarly, as pointed out by [90] the following axioms encode irreflexivity and symmetry of the friendship relation \succ_w :

$$(\text{irr}) \quad \neg @_n \hat{F}n \quad (\text{sym}) \quad @_n \hat{F}m \rightarrow @_m \hat{F}n$$

The proof of this lemma is also straightforward:

Lemma 5.18. *If (irr) and (sym) are present in the logic, each relation $\succ_{\mathcal{T}}$ in the canonical model is irreflexive and symmetric.*

The rest of the completeness proof proceeds as above. We note as well that the proof of Thm. 5.13 works without complication, for the property of \sim_a being an equivalence relation (or a preorder), and the property of \succ_w being irreflexive and symmetric are maintained under finite submodels. If we desire a reflexive ‘friendship’ relation instead, it suffices to add instead of (irr) the axiom

$$(\text{refl}) \quad @_n \hat{F}n;$$

one sees trivially that adding (refl) to the logic results in a reflexive \succ_T relation.

We thus have a complete axiomatisation of validity in the class of models proposed by [92]:

Theorem 5.19. $\text{SEL} \oplus \text{S5}_K \oplus (\text{irr}) \oplus (\text{sym})$ is the logic of finite indexed frames (W, A, \sim, \succ) where each \sim_a is an equivalence relation and each \succ_w is irreflexive and symmetric.

Finally, an optional further constraint is that an agent should not doubt who her own friends are. To achieve this, one would consider frames with the property: if $w \sim_a v$, then $a \succ_w b$ implies $a \succ_v b$. We will call these *k-y-f frames* (for “know your friends”). It is again very easy to check that, by adding to the logic the axiom

$$(\text{kyf}) \quad \hat{F}m \rightarrow K\hat{F}m,$$

the resulting canonical model is a k-y-f frame. The fact of being k-y-f is also maintained under finite submodels, so the finite model property is kept for this extended logic.

5.4 Axiomatisation of $\mathcal{L}(@\downarrow)$

In [92] another operator is borrowed from hybrid logic, namely $\downarrow x.\phi$, which names the current agent ‘ x ’, allowing to refer to her indexically. We now have, on top of **Prop** and **Nom**, a countable set $\text{SVar} = \{x, y, \dots\}$ of *state variables*. $\mathcal{L}(@\downarrow)$ is simply $\mathcal{L}(@)$ extended with x and $\downarrow x.\phi$, where $x \in \text{SVar}$. Formulas are read on named indexed models with respect to triples (g, w, a) , where $g : \text{SVar} \rightarrow A$ is an *assignment function*, as follows:

$$M, g, w, a \models x \quad \text{iff } g(x) = a;$$

$$M, g, w, a \models \downarrow x.\phi \quad \text{iff } M, g_a^x, w, a \models \phi,$$

where $g_a^x(y) = g(y)$ for $y \neq x$ and $g_a^x(x) = a$.

Given a formula ϕ and a nominal n , we define $\phi[x/n]$ to be the formula resulting from replacing each *free* occurrence of x in ϕ by n . Formally:

Definition 5.20. Given $x \in \text{SVar}$, $n \in \text{Nom}$ and $\phi \in \mathcal{L}(@\downarrow)$:

$$\phi[x/n] = \phi \text{ if } \phi = p \in \text{Prop}, \perp, m \in \text{Nom} \text{ or } y \in \text{SVar} \setminus \{x\}; \quad x[x/n] = n;$$

$$(\phi \wedge \psi)[x/n] = \phi[x/n] \wedge \psi[x/n]; \quad (\downarrow x.\phi)[x/n] = \downarrow x.\phi;$$

$$(B\phi)[x/n] = B(\phi[x/n]) \text{ if } B = \neg, K, F, @_m, \text{ or } \downarrow y \ (y \neq x);$$

With this, we can define the logic of the fragment $\mathcal{L}(@\downarrow)$:

Definition 5.21. SEL_\downarrow is the logic containing the axioms and rules of SEL plus the following axiom and rule:

- (DA) $\text{@}_n(\downarrow x.\phi \leftrightarrow \phi[x/n])$.
(FV) from $\phi[x/n]$ (with n fresh in ϕ), infer ϕ .

The fact that (DA) is sound can be checked by just unpacking the semantics. The soundness of the (FV) rule is a consequence of the following Lemma, whose proof is an easy induction on ϕ :

Lemma 5.22. *Let $\phi \in \mathcal{L}(\text{@}\downarrow)$ and n be fresh in ϕ . Let $\mathfrak{M} = (W, A, \sim, R, V)$ be a model and g an assignment. We define a new valuation in \mathfrak{M} by: $V'(n) = W \times \{g(x)\}$, $V'(m) = V(m)$ for $n \neq m$, $V'(p) = V(p)$ for $p \in \text{Prop}$. Let $M' = (W, A, \sim, R, V')$. Then $M, w, a, g \models \phi$ iff $M', w, a, g \models \phi[x/n]$.*

For completeness we shall use these two lemmas; respectively an application of the (FV) rule, and a straightforward induction on ϕ :

Lemma 5.23. *If ϕ is consistent and n_1, \dots, n_k are fresh, then $\phi[x_1/n_1] \dots [x_k/n_k]$ is consistent.*

Lemma 5.24. *Let \mathfrak{M} be a model, ϕ be a formula, g an assignment and $x_1, \dots, x_k \in \text{SVar}$. Let $n_1, \dots, n_k \in \text{Nom}$ such that $\underline{n}_i = g(x_i)$. Then*

$$M, w, a, g \models \phi \text{ iff } M, w, a, g \models \phi[x_1/n_1] \dots [x_k/n_k].$$

Now, we construct our canonical model exactly like before with one caveat: our sets MCT will only contain *BCN formulas without free variables* (i.e. *BCN sentences*). We prove the following variant of the Truth Lemma:

Proposition 5.25. *Let g be an assignment and ϕ a formula whose free variables are x_1, \dots, x_k . Let $[n_i] = g(x_i)$. Then*

$$M, T, [n], g \models \phi \text{ iff } \text{@}_n \phi[x_1/n_1] \dots [x_n/n_k] \in T.$$

Proof. First we note that if a formula has no free variables, the assignment g does not play a role in the semantics (and thus $M, T, [n], g \models \psi$ iff $M, T, [n], g' \models \psi$ for any g, g') and, with this in mind, we first prove:

$$\text{If } \psi \text{ is a sentence, then } M, T, [n], g \models \psi \text{ iff } \text{@}_n \psi \in T. \quad (*)$$

This suffices to prove our result: let x_1, \dots, x_k be all the free variables of ϕ . Then $M, T, [n], g \models \phi$ if and only if (by Lemma 5.24, noting that $g(x_i) = [n_i] = \underline{n}_i$)

$$M, T, [n], g \models \phi[x_i/n_i]_{i=1}^k,$$

if and only if (by the result we just proved, noting that $\phi[x_i/n_i]_{i=1}^k$ has no free variables) $@_n\phi[x_i/n_i]_{i=1}^k \in T$.

We prove (*) by induction on the length of ψ . It is exactly like the proof of Lemma 5.11, with one extra induction step:

We have that $@_n\downarrow x.\psi \in T$ if and only if (by the (DA) axiom) $@_n\psi[x/n] \in T$, if and only if (by induction hypothesis, since $\psi[x/n]$ has no free variables) $M, T, [n], g \models \psi[x/n]$, if and only if (because the choice of g does not affect the truth value of a sentence) $M, T, [n], g_n^x \models \psi[x/n]$, if and only if (by Lemma 5.24) $M, T, [n], g_n^x \models \psi$, which is the same as $M, T, [n], g \models \downarrow x.\psi$. ■

With this we can prove completeness:

Theorem 5.26. SEL_{\downarrow} is complete with respect to indexed models.

Proof. Suppose ϕ_0 is a consistent formula. Let x_1, \dots, x_k be the free variables of ϕ_0 and n_0, n_1, \dots, n_k fresh. Then $\phi_0[x_1/n_1] \dots [x_k/n_k]$ is a consistent sentence (by Lemma 5.23) and so is

$$@_{n_0}\phi_0[x_1/n_1] \dots [x_k/n_k]$$

(by Lemma 5.5). We extend this to $T_0 \in MCT$, we construct the corresponding canonical model and we let g be any assignment such that $g(x_i) = [n_i]$. Then we have by Prop. 5.25 that $M, T_0, [n_0], g \models \phi_0$. ■

5.5 Social Epistemic Logic in orthogonal structures

In Chapter 3, the class of ‘fully orthogonal frames’ was introduced (Def. 3.6) and it was shown that the elements of this class are isomorphic to indexed frames (Prop. 3.8).

Recall (Def. 3.15) that an *orthogonal structure* is simply a full orthogonal frame wherein the two equivalence relations are given explicitly.

Let us define a semantics for Social Epistemic Logic on fully orthogonal structures of the form $(X, R_K, R_F, \equiv_A, \equiv_W)$, where $R_K \subseteq \equiv_A$ and $R_F \subseteq \equiv_W$. The equivalence classes of \equiv_A and \equiv_W will represent agents and worlds respectively.

We shall be using SEL models (W, A, \sim, \asymp) with the constraints discussed in Section 5.3, namely: \sim must be an equivalence relation and \asymp must be symmetric and irreflexive. Therefore, via the isomorphism in Prop. 3.8, one easily sees the following:

Lemma 5.27. *Let $(X, R_K, R_F, \equiv_A, \equiv_W)$ be a fully orthogonal structure. The frame (X, R_K, R_F) is isomorphic to a SEL frame if and only if it satisfies*

- (FO1) $R_K \subseteq \equiv_A$ and $R_F \subseteq \equiv_W$,
- (FO2) $\equiv_A \cap \equiv_W = Id_X$,
- (FO3) $\equiv_A \circ \equiv_W = X^2$,

and, besides:

- (SEL1) R_K is an equivalence relation, and
- (SEL2) R_F is symmetric and irreflexive.

Recall (Prop. 3.8) that the corresponding isomorphic SEL model will be $(X/\equiv_W, X/\equiv_A, R_K, R_F)$, where R_K relates two pairs of equivalence classes if and only if the unique elements in the intersection of each pair are related by R_K (and likewise for R_F).

Now let us consider how a valuation must act upon this model. For a SEL model we demand that each $V(n)$ must be of the form $W \times \{a\}$ for some unique agent $a \in A$. Via the isomorphism outlined above, we can see, for the image of a valuation V defined on an orthogonal structure $(X, R_K, R_F, \equiv_A, \equiv_W)$ to be a valid valuation on a SEL model, we want the image of the set $V(n)$ to be $X/\equiv_W \times \{[y]_A\}$ for some $y \in X$. But the pre-image of this set is precisely $[y]_A$.

We thus demand the following property:

- (SEL3) $V(n) \in X/\equiv_A$ for all n .

For each nominal n and $x \in X$, we let n_x denote the unique element in $[x]_W \cap V(n)$.

Recall that SEL models need to be named. A *named model* is a model wherein every agent has a name, i.e., for all $a \in A$, there exists a nominal n such that $a = \underline{n}$. In these isomorphic structures, the corresponding notion of ‘named model’ translates to: for all $y \in X$, there exists n such that $V(n) = [y]_A$, or, equivalently,

- (SEL4) for all $x \in X$, there exists $n \in Nom$ such that $x \in V(n)$.

With all this we can define a semantics for Social Epistemic Logic on full orthogonal models $(X, R_K, R_F, \equiv_A, \equiv_W, V)$ where R_K, R_F and V satisfy the constraints (SEL1) – (SEL4) above as follows:

- $x \models F\phi$ iff $xR_F y$ implies $y \models \phi$;
- $x \models K\phi$ iff $xR_K y$ implies $y \models \phi$;
- $x \models n$ iff $x \in V(n)$ (iff $x = n_x$);
- $x \models @_n\phi$ iff $n_x \models \phi$.

In his recent PhD thesis, Zhen Liang [105] considers a ‘non-rigid’ variant of SEL which can assign different names to agents in each possible world. This is imposed via the following, weaker, constraint on the valuation:

for every nominal n and each world w , there exists a unique agent $a \in A$ such that $(w, a) \in V(n)$.

In the isomorphic structures above, this translates to a constraint weaker than (SEL3), namely:

(SEL3’) for each n and each $x \in X$,
the intersection $[x]_W \cap V(n)$ is a singleton.

The proof of completeness given in this chapter of (standard, rigid) SEL using ‘indexed canonical models’ (Theorems 5.12 and 5.19) does not seem to do the trick when it comes to non-rigid models. Completeness of ‘non-rigid’ SEL was proven in [105] by means of an involved step-by-step construction, but a proof of this result using canonical models remains, at the moment of this writing, an open problem. I conjecture that the semantics above could assist in this endeavour.

Discussion. This chapter has studied several aspects of the framework introduced in [93]. We have as well provided axiomatisations for the fragment $\mathcal{L}(@\downarrow)$, on top of a novel proof of completeness of SEL for the fragment $\mathcal{L}(@)$, for which we have given decidability.

Some interesting directions for future work include studying the decidability of $\mathcal{L}(@\downarrow)$, resolving the conjecture in Remark 5.17, and using the results in the last section to offer a canonical model proof of completeness for non-rigid SEL.

Part II

*ASYNCHRONOUS
ANNOUNCEMENTS*

Asynchronous Announcement Logic

We propose a multi-agent epistemic logic of asynchronous announcements, where truthful announcements are publicly sent but individually received by agents, in the order in which they were sent. On top of epistemic modalities, this logic contains dynamic modalities for both making announcements and for receiving them. What an agent believes is a function of her initial uncertainty and of the announcements she has received. Beliefs need not be truthful, because announcements already made may not yet have been received. As announcements are true when sent, certain message sequences can be ruled out.

We provide a complete axiomatization for this *asynchronous announcement logic* (AA). It is a reduction system that also demonstrates that any formula in AA is equivalent to one without dynamic modalities, just as for public announcement logic.

We provide some results for the class of models $\mathcal{S5}$ and we conclude this investigation by comparing the framework of AA with previous approaches to asynchronous information broadcast.¹



– Echo. Echo! ECHO!!

Heard in Toulouse: girl of about seven years of age, yelling directly at a brick wall, frustratedly waiting for something to happen.

HOW DOES THE KNOWLEDGE of an agent change in a multi-agent system wherein agents act independently and may keep their own time?

The topic of the second part of this thesis is *asynchronicity*.

One type of asynchronicity stems from agents being uncertain about the number of actions which have taken place: this is asynchronicity due to *partial observation*. Many DEL scenarios present this kind of asynchronicity: for instance, in gossip protocols agents communicate by calling each other, so that a may have called b without another agent c noticing that the call took

¹Although Section 6.5 is novel, the rest of this chapter is based on the paper

Philippe Balbiani, Hans van Ditmarsch, and Saúl Fernández González. Asynchronous announcements. *ACM Transactions in Computational Logic*, 23(2), 2022.

place. (See e.g. [3].) Another example: in the ‘One hundred prisoners and a light bulb’ riddle agents communicate asynchronously by individually toggling a light bulb out of sight and hearing of other agents [44]. As a final example, the ‘immediate snapshot algorithm’, wherein agents are unaware of other agents possibly simultaneously accessing a shared memory location, has been modelled in DEL by [54].

Other notions of asynchronous knowledge under conditions of temporal uncertainty have been investigated in depth in distributed computing [21, 50, 70, 71, 78] and in temporal epistemic logics [32, 58, 81, 87].

However, this chapter (and indeed the remainder of this dissertation) is concerned with a different kind of asynchronicity: namely, the one which occurs when the sending and receiving of messages are separate, so that the receiver of a message is uncertain about the moment it was sent.

The asynchronous reception of messages broadcast by the environment has (seemingly) only been modelled in DEL by [67] (see also [66, 91] by the same authors). Here, much like [67], we assume that announcements are still broadcast to all agents, but individually received. But unlike [67] our epistemic notion is interpreted over past messages only, and we provide an axiomatization by way of a reduction to the modal fragment, just as for public announcement logic [86]. Similarities and differences between the present framework and that of [67] are highlighted in Section 6.5.

Let us present an example illustrating our approach before diving into the technicalities.

Example 6.1. Two agents, Anxélica (a) and Bertu (b), know the truth about two propositional variables, p and q . Suppose that Anxélica knows whether p and Bertu knows whether q , and this is common knowledge between them. Let us say p and q are both true.

We can encode this uncertainty in a model, as shown in Fig. 6.1(i).

After the announcement of $p \vee q$, Bertu does not know that p is true but Anxélica considers it possible that he knows. We can formalise this using Public Announcement Logic [86], as the formula

$$[p \vee q](\neg B_b p \wedge \hat{B}_a B_b p).$$

The operator $[p \vee q]$ is a dynamic modality interpreted by model restriction. The belief (or knowledge) modalities bound by it are not interpreted in the initial model, but rather in its restriction to those worlds where $p \vee q$ is true, as seen in Fig. 6.1(PAL).

Let us now assume that announcements are still publicly sent, but individually received. Then, after the announcement $p \vee q$ is made (Fig. 6.1(ii)),




Figure 6.1: Left, public announcement, and right, asynchronous announcement. Right, the announcement $p \vee q$ is sent, after which first Anxélica and then Bertu receives it. What Anxélica and Bertu know is a function of the initial *model* encoding their knowledge and ignorance and the actual *state* in this model (i), and this *history* $(p \vee q).a.b$ of three events. States are labelled with the valuations of p and q , where $\neg p$ stands for \bar{p} and \bar{q} stands for $\neg q$. States that are indistinguishable for an agent are linked with a label for that agent. The greying of states and links represents states becoming inaccessible for the relevant agents.

Anxélica may have received that information $p \vee q$ while Bertu has not (Fig. 6.1(iii)); after this, Bertu receives it too (Fig. 6.1(iv)). Unlike in Figure 6.1(PAL), in (iv) they do not know that the other knows; there is no common knowledge between them of $p \vee q$.

Separating sending from receiving messages permits a notion of asynchronous knowledge in DEL which is a function of the usual modal accessibility but also of uncertainty over the announcements received by other agents. In (iii), after receiving announcement $p \vee q$, Anxélica considers it possible that Bertu knows p : she can conceive that the actual state is $p\bar{q}$ and that Bertu has also received the announcement, as in (iv). For different reasons she also considers it conceivable that Bertu does not know p : for instance if the state is pq (or $\bar{p}q$) and the announcement has been received by Bertu (iv); alternatively, if Bertu has not yet received the announcement (iii). What Anxélica knows in (iii) should be the same as what she knows in (iv).

What does Bertu know? According to the standard usage in distributed computing [58], even when Bertu has not received the announcement $p \vee q$, he can imagine that such a message has been sent and that Anxélica has received it. Therefore, although he is uncertain about p (i), he should consider it possible that announcement $p \vee q$ was made (ii), and that Anxélica has received it (iii), that she therefore considers it possible that he has received it too (iv), and that he therefore now knows that p . This notion of knowledge does not seem to fit well a setting in which the messages are announcements, whose

role is to reduce uncertainty of the value of unchanging facts. It does fit the setting of distributed computing, wherein messages that are broadcast contain novel facts. In PAL the future is predictable: all facts may become known. Thus in our setting, such a notion of knowledge would only allow for weak forms of higher-order knowledge: an agent cannot know that another agent remains ignorant.

We therefore focus on what agents know based on the announcements they have received so far, ignoring possible future announcements. That means that in situations (i), (ii), (iii) above, Bertu ‘knows’ that Anxélica knows that he is uncertain about p , as he has not received the announcement $p \vee q$. In (i) and (ii) this is true, but in (iii) this is no longer true. Bertu’s ‘knowledge’ is then an incorrect belief. Indeed, the asynchronous epistemic notion that we propose is one of asynchronous belief (however, as we will see, of a special kind such that many beliefs will eventually become knowledge). Other defining assumptions of our asynchronous semantics are that agents receive the announcements in the order in which they are made, as is not uncommon in distributed computing [58, 81]; and that announcements are true when sent, as in PAL.

The assumption that announcements are true when sent, results in *partial synchronization*. Let us suppose that Anxélica and Bertu are both uncertain about p . Then, announcement p is followed by announcement $\neg B_b p$. If Anxélica received p and $\neg B_b p$, she should not consider it possible that Bertu has received p before the second announcement was made. In other words, the histories of sending and receiving events that she considers possible include $p.a.(\neg B_b p).b.a$ and $p.(\neg B_b p).a.a.b$, but exclude $p.b.(\neg B_b p).a.a$ and $p.a.b.(\neg B_b p).b.a$ where the first a in the sequence stands for Anxélica receiving the first announcement p , the second a stands for her receiving the second announcement $\neg B_b p$, and similarly for b . Using the terminology of [81], $p.b.(\neg B_b p).a.a$ and $p.a.b.(\neg B_b p).b.a$ would be called *inconsistent cuts*. In order for the second announcement $\neg B_b p$ to be truthful, Bertu must still be uncertain about p , and for Bertu to remain uncertain about p he must not yet have received the first announcement p . We will introduce a so-called ‘agreement’ relation between states and histories, and we then say that a state s in the model does not *agree* with history $p.b.(\neg B_b p).a.a$. Agents only consider histories possible that agree with the states they consider possible. This requires to define a satisfaction relation and an ‘agreement’ relation by simultaneous induction.

Intuitively, in our approach an agent knows/believes ϕ iff ϕ is true: (1) in all states that she considers possible, (2) for all prefixes of announcement

sequences that other agents may have received, (3) taking into account that the announcements she received were true when sent, (4) while ignoring that other agents may have received more announcements than herself.

This chapter is structured as follows: Section 6.1 defines the syntax and Section 6.2 defines the semantics of *asynchronous announcement logic* AA. Section 6.3 provides an axiomatization of AA on the class of models with empty histories. Section 6.4 obtains results for the model class $\mathcal{S5}$ and elaborates on the difference between knowledge and belief. A comparison between this framework and that of [67] is found in Section 6.5.

6.1 Syntax

Let us establish the language of Asynchronous Announcement Logic.

Definition 6.2 (Language of AA). Let Prop be a countable set of propositional variables (denoted p, q, \dots) and A be a finite set of agents (denoted a, b, \dots). The language \mathcal{L}_{AA} of asynchronous announcement logic is defined as follows:

$$\phi, \psi := p \mid \perp \mid \neg\phi \mid (\phi \vee \psi) \mid B_a\phi \mid [\phi]\psi \mid [a]\phi$$

We will follow the standard rules for omission of the parentheses. Without the $[\phi]$ and $[a]$ modal operators we obtain the language \mathcal{L}_{EL} of multi-agent epistemic logic. The *positive* fragment $\mathcal{L}_{\text{EL}}^+$ of \mathcal{L}_{EL} is defined as follows:

$$\phi, \psi := p \mid \neg p \mid \perp \mid \top \mid (\phi \vee \psi) \mid (\phi \wedge \psi) \mid B_a\phi$$

We will use the standard abbreviations for the Boolean operators. We will also use the following abbreviations: $\langle a \rangle\phi := \neg[a]\neg\phi$, $\langle \phi \rangle\psi := \neg[\phi]\neg\psi$ and $\hat{B}_a\phi := \neg B_a\neg\phi$. $B_a\phi$ reads “agent a believes/knows ϕ ,” $[\phi]\psi$ reads “after public announcement ϕ has been sent/made, ψ ,” and $[a]\phi$ reads “after agent a receives/reads the next announcement, ϕ .” For all formulas η, ψ and for all atoms p , we denote by $\eta(p/\psi)$ the uniform substitution of the occurrences of p in η by ψ .

Definition 6.3 (Words). We shall consider the set $(A \cup \mathcal{L}_{\text{AA}})^*$ of *words* over an alphabet using agents $a \in A$ and formulas $\phi \in \mathcal{L}_{\text{AA}}$ as letters; this is simply the set of finite sequences of $A \cup \mathcal{L}_{\text{AA}}$. We use α, β, \dots as variables for these words; the empty word is denoted ϵ .

Rather than using the more standard notation for sequences $\alpha = (x_1, \dots, x_n)$, we will write a word simply as its ‘letters’ separated by dots; for instance, if $a, b \in A$, and $p, q \in \text{Prop}$, then

$$\alpha = p \vee q.a.\neg q.a.b$$

is a word, representing the sequence of formulas and agents $(p \vee q, a, \neg q, a, b)$.

Given a word $\alpha \in (A \cup \mathcal{L}_{AA})^*$, $|\alpha|$ is its length, $|\alpha|_a$ is the number of its a 's for each $a \in A$, $|\alpha|_!$ is the number of its formula occurrences, $\alpha \upharpoonright_!$ is the projection of α to \mathcal{L}_{AA} , and $\alpha \upharpoonright_{!a}$ is the restriction of $\alpha \upharpoonright_!$ to the first $|\alpha|_a$ formulas. These notions have obvious inductive definitions. We say that a word β is a *prefix* of a word α (in symbols $\beta \sqsubseteq \alpha$) if β is an initial sequence of α .

Given words α and β , $\alpha.\beta$ denotes their concatenation.

Let $A = \{a_1, \dots, a_n\}$ be the set of agents; it is clear that

$$|\alpha| = |\alpha|_{a_1} + \dots + |\alpha|_{a_n} + |\alpha|_!.$$

We shall use this notion of ‘words’, in the present and subsequent Chapters, to formalise sequences of announced messages and agent’s readings. In an informal example, in a case in which $a, b \in A$ and $p, q \in \mathbf{Prop}$, the word $\alpha = p \vee q.a.b.\neg q.a$ should represent a situation in which (1st) the message $p \vee q$ is announced, (2nd) agent a receives it, (3rd) agent b receives it, (4th) the message $\neg q$ is announced, and (5th) a receives this last message (the informal intuition here is that agent a should know p to be true after this sequence).

Not every word has the power to represent a ‘valid’ sequence of announcements and readings; for instance if $\beta = a.p \vee q$, we have an undesirable situation in which agent a receives a message before any have been sent. Words that represent ‘acceptable’ sequences will be called *histories*. Formally,

Definition 6.4 (Histories). A word α in the language $A \cup \mathcal{L}_{AA}$ is a *history* if for all prefixes $\beta \sqsubseteq \alpha$ and for all $a \in A$, $|\beta|_! \geq |\beta|_a$.

Obviously, if β is a prefix of a history α , then β is a history too. In the definition of histories, the requirement “for all $a \in A$, $|\beta|_! \geq |\beta|_a$ ” means that, for all $n \in \mathbb{N}$, if there is an n -th occurrence of agent a in α , then there is a prior n -th formula in α , the intuition being that this will be the announcement then received by agent a .

We now formalise the notion that, given history α , history β is an epistemic alternative to α for agent a :

Definition 6.5 (View relation). Let α, β be histories and $a \in A$. We define: $\alpha \triangleright_a \beta$ iff $\beta \upharpoonright_! = \beta \upharpoonright_{!a} = \alpha \upharpoonright_{!a}$. The set $\mathbf{view}_a(\alpha) := \{\beta \mid \alpha \triangleright_a \beta\}$ is the *view* of a given α .

Observe that, for all $a \in A$, if $\alpha \triangleright_a \beta$ then $\beta \upharpoonright_1$ is a prefix of $\alpha \upharpoonright_1$. Informally, the view of agent a given history α consists of all the different ways in which a can receive the announcements she receives in α . In other words, the view of a given α consists of the histories a considers possible but without taking into account neither the meaning of the announcements in the history (which, as we will see, results in a further restriction) nor the possibility of there being unread announcements. In Section 6.4 we will present an alternative for the view relation, without the requirement that $|\beta|_1 = |\alpha|_a$.

Example 6.6. Let us have two agents, $A = \{a, b\}$, and let the history be $\alpha = (p \vee q).a$. Then $\text{view}_a(\alpha)$, the set of all histories β such that $\alpha \triangleright_a \beta$, is

$$\{(p \vee q).a.b, (p \vee q).b.a, (p \vee q).a\},$$

whereas $\text{view}_b(\alpha) = \{\epsilon\}$.

Let now $\alpha' = (p \vee q).a.b$. Then $\text{view}_a(\alpha') = \{(p \vee q).a.b, (p \vee q).b.a, (p \vee q).a\}$ and $\text{view}_b(\alpha') = \{(p \vee q).a.b, (p \vee q).b.a, (p \vee q).b\}$.

The following alternative characterizations of the \triangleright_a relation will be useful. They follow from simply unpacking the definitions; the proof is left to the reader.

Lemma 6.7. *Let α, β be histories and $a \in A$. The following conditions are equivalent:*

1. $|\beta|_a = |\alpha|_a$, $\beta \upharpoonright_{1a} = \alpha \upharpoonright_{1a}$ and $|\beta|_1 = |\alpha|_a$;
2. $|\beta|_a = |\alpha|_a$ and $\beta \upharpoonright_1 = \alpha \upharpoonright_{1a}$;
3. $\beta \upharpoonright_1 = \beta \upharpoonright_{1a} = \alpha \upharpoonright_{1a}$.

We can introduce modalities for histories by abbreviation, using reception and announcement modalities of the form $[a]$ and $[\psi]$.

Definition 6.8. For all words α over $A \cup \mathcal{L}_{AA}$, the modality $[\alpha]$ is inductively defined as:

$$[\epsilon]\phi := \phi, [\alpha.a]\phi := [\alpha][a]\phi, [\alpha.\psi]\phi := [\alpha][\psi]\phi;$$

whereas its dual $\langle \alpha \rangle \phi$ is defined by abbreviation as $\neg[\alpha]\neg\phi$.

For instance, if ϕ is a formula in \mathcal{L}_{AA} and $\alpha = (p \wedge \neg q).a.b.B_b p.a$, then $\langle \alpha \rangle \phi$ is an abbreviation for

$$\langle p \wedge \neg q \rangle \langle a \rangle \langle b \rangle \langle B_b p \rangle \langle a \rangle \phi.$$

We will read $[\alpha]\phi$ as “if the sequence α of events can be executed then ϕ holds after its execution,” whereas we will read $\langle\alpha\rangle\phi$ as “the sequence α of events can be executed and ϕ holds after its execution”.

Note that every formula in \mathcal{L}_{AA} can be written in the form $[\alpha]\phi$.

6.1.1 Results for histories

We continue with some basic results for histories that will be used later.

Lemma 6.9. *Let α, β be histories. For all $a, b \in A$, if $\alpha \triangleright_a \beta$ then $|\beta|_b \leq |\alpha|_a$.*

Proof. Let $a, b \in A$ be such that $\alpha \triangleright_a \beta$. Hence, $|\beta|_1 = |\alpha|_a$. Since β is a history, we have $|\beta|_1 \geq |\beta|_b$. And since $|\beta|_1 = |\alpha|_a$, we have $|\beta|_b \leq |\alpha|_a$. ■

Lemma 6.10. *Let α, β be histories, and $a \in A$.*

1. *If $\epsilon \triangleright_a \alpha$ then $\alpha = \epsilon$,*
2. *if $\alpha \triangleright_a \beta$ then $\beta \triangleright_a \alpha$.*

Proof.

1. Suppose $\epsilon \triangleright_a \alpha$. Hence, $|\alpha|_1 = |\epsilon|_a = 0$. Since α is a history, for all $b \in A$, $|\alpha|_1 \geq |\alpha|_b$, and thus $|\alpha|_b = 0$. Consequently, $\alpha = \epsilon$.
2. Suppose $\alpha \triangleright_a \beta$. Hence, $|\beta|_a = |\alpha|_a$ and $|\beta|_1 = |\alpha|_a$. Thus, $|\beta|_1 = |\beta|_a$. Consequently, $\beta \triangleright_a \alpha$. ■

Given history α and agent a , we recursively define a word α_a as follows:

- $\epsilon_a = \epsilon$;
- $(\alpha.\phi)_a = \alpha_a$;
- $(\alpha.b)_a = \alpha_a$ for each $b \in A \setminus \{a\}$;
- for all $n > 0$, $(\alpha.b.a^n)_a = (\alpha.a^n)_a$ for each $b \in A \setminus \{a\}$;
- for all $n > 0$, if $|\alpha.\phi.a^n|_1 = |\alpha.\phi.a^n|_a$ then $(\alpha.\phi.a^n)_a = \alpha.\phi.a^n$, otherwise $(\alpha.\phi.a^n)_a = (\alpha.a^n)_a$.

Informally, α_a is obtained by taking the prefix γ of α until the $|\alpha|_a$ -th occurrence of a formula in α and adding at the end $|\alpha|_a - |\gamma|_a$ times the letter a .

Lemma 6.11. *For all histories α and for all agents a , α_a is a history such that $\alpha \triangleright_a \alpha_a$*

Proof. Induction on the length of α . ■

Proposition 6.12 (The view relation is serial, transitive, and Euclidean). *Let α , β and γ be histories. For all agents a ,*

1. *there is a history δ such that $\alpha \triangleright_a \delta$,*
2. *if $\alpha \triangleright_a \beta$ and $\beta \triangleright_a \gamma$ then $\alpha \triangleright_a \gamma$,*
3. *if $\alpha \triangleright_a \beta$ and $\alpha \triangleright_a \gamma$ then $\beta \triangleright_a \gamma$.*

Proof.

1. By Lemma 6.11: we can take δ to be α_a .

2. Suppose $\alpha \triangleright_a \beta$ and $\beta \triangleright_a \gamma$.

Then $\alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a} = \beta \upharpoonright_{!}$, and $\beta \upharpoonright_{!a} = \gamma \upharpoonright_{!a} = \gamma \upharpoonright_{!}$, whence $\alpha \upharpoonright_{!a} = \gamma \upharpoonright_{!a} = \gamma \upharpoonright_{!}$.

3. Suppose $\alpha \triangleright_a \beta$ and $\alpha \triangleright_a \gamma$.

Then $\alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a} = \beta \upharpoonright_{!}$ and $\alpha \upharpoonright_{!a} = \gamma \upharpoonright_{!a} = \gamma \upharpoonright_{!}$, whence $\beta \upharpoonright_{!a} = \gamma \upharpoonright_{!a} = \gamma \upharpoonright_{!}$. ■

From Proposition 6.12, we obtain:

Corollary 6.13. *For all histories α, β and for all agents a , if $\alpha \triangleright_a \beta$ then $\alpha_a \triangleright_a \beta$ and $\beta \triangleright_a \alpha_a$.*

Lemma 6.14. *Let α, β be histories. In the single-agent case, if $\alpha \triangleright_a \beta$ then $|\beta| = 2|\alpha|_a$. Otherwise, in the multi-agent case, if $\alpha \triangleright_a \beta$ then*

$$2|\alpha|_a \leq |\beta| \leq (|A| + 1)|\alpha|_a.$$

Proof. In the single-agent case, suppose $\alpha \triangleright_a \beta$. Hence, $|\beta|_a = |\alpha|_a$ and $|\beta|_{!} = |\alpha|_{!}$. Thus, $|\beta| = 2|\alpha|_a$. In the multi-agent case, suppose $\alpha \triangleright_a \beta$. Hence, by Lemma 6.9, for all $b \in A \setminus \{a\}$, $|\beta|_b \leq |\alpha|_a$. Moreover, $|\beta|_a = |\alpha|_a$ and $|\beta|_{!} = |\alpha|_{!}$. Thus, $2 \cdot |\alpha|_a \leq |\beta| \leq (|A| + 1) \cdot |\alpha|_a$. ■

Since A is finite, by Lemma 6.14, the set $\text{view}_a(\alpha) = \{\beta : \alpha \triangleright_a \beta\}$ is finite. But we can do better. Let X and Y be distinct symbols. A *Dyck word* is a string consisting, for some $n \in \mathbb{N}$, of n X 's and n Y 's such that no prefix of the string has more Y 's than X 's. This matches exactly our histories of

announcements (X) and read actions (Y). The number of *Dyck words* of length $2n$ is C_n where C_n is the n -th *Catalan number*, defined as $C_n := \frac{1}{n+1} \binom{2n}{n}$. This generates the sequence $1, 1, 2, 5, 14, 42, \dots$ ²

Proposition 6.15. *In the single-agent case, for all histories α , $|\text{view}_a(\alpha)| = C_{|\alpha|_a}$.*

Proof. If there is a single agent, then histories can be transformed into Dyck words over the symbols X and Y when one replaces announcements by the symbol X and read actions by the symbol Y . Moreover, $\text{view}_a(\alpha)$ is the set of all histories β such that $\beta \upharpoonright ! = \alpha \upharpoonright !_a$ and $\beta \upharpoonright !_a = \alpha \upharpoonright !_a$. Hence, $|\text{view}_a(\alpha)| = C_{|\alpha|_a}$. ■

However, in the multi-agent case, an agent can receive n announcements in many more than C_n ways. Example 6.6 showed that if there are two agents, an agent can receive one announcement in three different ways instead of one way for one agent.

6.2 Semantics

In the present Section we introduce the semantics of Asynchronous Announcement logic. First we need a way to order pairs (α, ϕ) consisting of a history and a formula.

6.2.1 A well-founded order for the semantics

A well-founded order \ll between (history, formula) pairs will be the basis of our semantics. It uses an auxiliary function $\|\cdot\|$ on formulas and on histories, and an auxiliary function $\text{deg}(\cdot)$ on (history, formula) pairs.

For all $\phi \in \mathcal{L}_{AA}$, let $\|\phi\|$ be the positive integer inductively defined as follows:

$$\begin{array}{lll} \|\mathit{p}\| & = 2 & \|\phi \vee \psi\| = \|\phi\| + \|\psi\| \quad \|\llbracket \phi \rrbracket \psi\| = 2\|\phi\| + \|\psi\| \\ \|\perp\| & = 1 & \|\mathit{B}_a \phi\| = \|\phi\| + 1 \quad \|\llbracket a \rrbracket \phi\| = \|\phi\| + 2 \\ \|\neg \phi\| & = \|\phi\| + 1 & \end{array}$$

and for all words α over $A \cup \mathcal{L}_{AA}$, let $\|\alpha\|$ be the nonnegative integer inductively defined as:

$$\|\epsilon\| = 0 \quad \|\alpha.a\| = \|\alpha\| + 1 \quad \|\alpha.\psi\| = \|\alpha\| + \|\psi\|$$

²For details, see <https://oeis.org/A000108>.

Then, for all formulas ϕ , let $\deg(\phi)$ be the nonnegative integer inductively defined as follows (this is often known as the *modal depth* of a formula, the maximum stack of epistemic modalities potentially occurring in it):

$$\begin{aligned} \deg(p) &= 0 & \deg(B_a\phi) &= \deg(\phi) + 1 \\ \deg(\perp) &= 0 & \deg([\phi]\psi) &= \deg(\phi) + \deg(\psi) \\ \deg(\neg\phi) &= \deg(\phi) & \deg([a]\phi) &= \deg(\phi) \\ \deg(\phi \vee \psi) &= \max\{\deg(\phi), \deg(\psi)\} \end{aligned}$$

Also, given a pair of the form (α, ϕ) where α is a history and $\phi \in \mathcal{L}_{AA}$,

$$\deg(\alpha, \phi) = \deg([\alpha]\phi).$$

Finally, let \ll be the well-founded order between (history, formula) pairs defined as follows:

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

Various results for this order are shown in the next subsection.

6.2.2 Results for the well-founded order \ll

Let us see some results for the order \ll , and the functions $\|\cdot\|$ and $\deg(\cdot)$. First, concerning $\|\cdot\|$, note that, obviously, for all words α over $A \cup \mathcal{L}_{AA}$ and for all $a \in A$, $|\alpha|_a \leq \|\alpha\|$.

The following lemmas are shown by induction on the length of $|\alpha|$:

Lemma 6.16. *For all words α over $A \cup \mathcal{L}_{AA}$, for all $a \in A$ and for all $\phi \in \mathcal{L}_{AA}$, $\|a.\alpha\| = \|\alpha\| + 1$ and $\|\phi.\alpha\| = \|\alpha\| + \|\phi\|$.*

Lemma 6.17. *For all words α over $A \cup \mathcal{L}_{AA}$ and for all $\phi \in \mathcal{L}_{AA}$, $\|[\alpha]\phi\| = 2\|\alpha\| + \|\phi\|$.*

Lemma 6.18. *For all words α over $A \cup \mathcal{L}_{AA}$, for all $k \in \mathbb{N}$, for all $\psi_1, \dots, \psi_k \in \mathcal{L}_{AA}$ and for all $\phi \in \mathcal{L}_{AA}$, if $\alpha \upharpoonright_! = \psi_1 \dots \psi_k$ then $\deg([\alpha]\phi) = \deg(\psi_1) + \dots + \deg(\psi_k) + \deg(\phi)$.*

Lemma 6.19. *Let α be a history. Let k be a nonnegative integer and $\phi_1, \dots, \phi_k \in \mathcal{L}_{AA}$. In the single-agent case, if $\alpha \upharpoonright_! = \phi_1 \dots \phi_k$ then $\|\alpha\| = |\alpha|_a + \|\phi_1\| + \dots + \|\phi_k\|$. Otherwise, in the multi-agent case, considering an enumeration (a_1, \dots, a_n) of A without repetition, if $\alpha \upharpoonright_! = \phi_1 \dots \phi_k$ then $\|\alpha\| = |\alpha|_{a_1} + \dots + |\alpha|_{a_n} + \|\phi_1\| + \dots + \|\phi_k\|$.*

Furthermore, we have:

Lemma 6.20. *Let α, β be histories and $a \in A$. If $\alpha \triangleright_a \beta$ then $\|\beta\| \leq |A| \cdot \|\alpha\|$.*

Proof. Suppose $\alpha \triangleright_a \beta$. Then $|\beta|_a = |\alpha|_a$ and $\beta \upharpoonright_1$ is a prefix of $\alpha \upharpoonright_1$. By Lemma 6.9, for all $b \in A \setminus \{a\}$, $|\beta|_b \leq |\alpha|_b$. Thus, by Lemma 6.19, $\|\beta\| \leq |A| \cdot |\alpha|_a + \|\alpha\| - |\alpha|_a$. Consequently, $\|\beta\| \leq (|A| - 1) \cdot |\alpha|_a + \|\alpha\|$, and thus $\|\beta\| \leq |A| \cdot \|\alpha\|$. ■

Lemma 6.21. *Let α, β be histories and $a \in A$. Let $k, l \in \mathbb{N}$ and $\phi_1, \dots, \phi_k, \psi_1, \dots, \psi_l \in \mathcal{L}_{AA}$ be such that $\alpha \upharpoonright_1 = \phi_1 \dots \phi_k$ and $\beta \upharpoonright_1 = \psi_1 \dots \psi_l$. If $\alpha \triangleright_a \beta$, then $\deg(\phi_1) + \dots + \deg(\phi_k) \geq \deg(\psi_1) + \dots + \deg(\psi_l)$.*

Proof. Suppose $\alpha \triangleright_a \beta$. Since $\alpha \upharpoonright_1 = \phi_1 \dots \phi_k$ and $\beta \upharpoonright_1 = \psi_1 \dots \psi_l$, $\psi_1 \dots \psi_l$ is a prefix of $\phi_1 \dots \phi_k$. Thus, $\deg(\phi_1) + \dots + \deg(\phi_k) \geq \deg(\psi_1) + \dots + \deg(\psi_l)$. ■

Lemma 6.22.

1. $(\alpha, \phi) \ll (\alpha.a, \phi)$,
2. $(\alpha, \psi) \ll (\alpha.\psi, \phi)$ and $(\alpha, \perp) \ll (\alpha.\phi, \perp)$,
3. $(\alpha, \phi) \ll (\alpha, \neg\phi)$ and $(\alpha, \phi) \ll (\alpha.\phi, \perp)$,
4. $(\alpha, \phi) \ll (\alpha, \phi \vee \psi)$ and $(\alpha, \psi) \ll (\alpha, \phi \vee \psi)$,
5. if $\alpha \triangleright_a \beta$ then $(\beta, \phi) \ll (\alpha, B_a\phi)$,
6. $(\alpha, \phi) \ll (\alpha, [\phi]\psi)$ and $(\alpha.\phi, \psi) \ll (\alpha, [\phi]\psi)$,
7. $(\alpha.a, \phi) \ll (\alpha, [a]\phi)$,
8. $(\alpha, \perp) \ll (\alpha, p)$,
9. $(\alpha, \perp) \ll (\alpha, B_a\phi)$,
10. $(\alpha, \perp) \ll (\alpha, \neg\phi)$.

Proof.

1. Note that $\deg(\alpha, \phi) = \deg(\alpha.a, \phi)$. Moreover, since $\|\alpha\| + \|\phi\| < \|\alpha\| + 1 + \|\phi\|$, $(\alpha, \phi) \ll (\alpha.a, \phi)$.
2. We have $\deg(\alpha, \psi) \leq \deg(\alpha.\psi, \phi)$. Moreover, since $\|\alpha\| + \|\psi\| < \|\alpha\| + \|\psi\| + \|\phi\|$, $(\alpha, \psi) \ll (\alpha.\psi, \phi)$. In other respect, $\deg(\alpha, \perp) \leq \deg(\alpha.\phi, \perp)$. Moreover, since $\|\alpha\| + 1 < \|\alpha\| + \|\phi\| + 1$, $(\alpha, \perp) \ll (\alpha.\phi, \perp)$.

3. Note that $\deg(\alpha, \phi) = \deg(\alpha, \neg\phi)$. Moreover, since $\|\alpha\| + \|\phi\| < \|\alpha\| + \|\phi\| + 1$, $(\alpha, \phi) \ll (\alpha, \neg\phi)$. In other respect, remark that $\deg(\alpha, \phi) = \deg(\alpha, \phi, \perp)$. Moreover, since $\|\alpha\| + \|\phi\| < \|\alpha\| + \|\phi\| + 1$, we obtain that $(\alpha, \phi) \ll (\alpha, \phi, \perp)$.
4. Note that $\deg(\alpha, \phi) \leq \deg(\alpha, \phi \vee \psi)$ and $\deg(\alpha, \psi) \leq \deg(\alpha, \phi \vee \psi)$. Moreover, since $\|\alpha\| + \|\phi\| < \|\alpha\| + \|\phi\| + \|\psi\|$ and $\|\alpha\| + \|\psi\| < \|\alpha\| + \|\phi\| + \|\psi\|$, we obtain that $(\alpha, \phi) \ll (\alpha, \phi \vee \psi)$ and $(\alpha, \psi) \ll (\alpha, \phi \vee \psi)$.
5. Suppose $\alpha \triangleright_a \beta$. Let $k, l \in \mathbb{N}$ and $\phi_1, \dots, \phi_k, \psi_1, \dots, \psi_l$ be formulas such that $\alpha \upharpoonright_1 = \phi_1 \dots \phi_k$ and $\beta \upharpoonright_1 = \psi_1 \dots \psi_l$. By Lemma 6.21, we have $\deg(\phi_1) + \dots + \deg(\phi_k) \geq \deg(\psi_1) + \dots + \deg(\psi_l)$. Therefore, $\deg(\beta, \phi) < \deg(\alpha, B_a\phi)$. Consequently, $(\beta, \phi) \ll (\alpha, B_a\phi)$.
6. In this case we have $\deg(\alpha, \phi) \leq \deg(\alpha, [\phi]\psi)$ and $\deg(\alpha, \phi, \psi) = \deg(\alpha, [\phi]\psi)$. Moreover, since $\|\alpha\| + \|\phi\| < \|\alpha\| + 2\|\phi\| + \|\psi\|$ and $\|\alpha\| + \|\phi\| + \|\psi\| < \|\alpha\| + 2\|\phi\| + \|\psi\|$, we obtain that $(\alpha, \phi) \ll (\alpha, [\phi]\psi)$ and $(\alpha, \phi, \psi) \ll (\alpha, [\phi]\psi)$.
7. Note that $\deg(\alpha, a, \phi) = \deg(\alpha, [a]\phi)$. Moreover, since $\|\alpha\| + 1 + \|\phi\| < \|\alpha\| + \|\phi\| + 2$, we obtain that $(\alpha, a, \phi) \ll (\alpha, [a]\phi)$.
8. We have $\deg(\alpha, \perp) = \deg(\alpha, p)$. Moreover, since $\|\alpha\| + 1 < \|\alpha\| + 2$, $(\alpha, \perp) \ll (\alpha, p)$.
9. Let $k \in \mathbb{N}$ and ϕ_1, \dots, ϕ_k be formulas such that $\alpha \upharpoonright_1 = \phi_1 \dots \phi_k$. Remark that $\deg(\alpha, \perp) = \deg(\phi_1) + \dots + \deg(\phi_k)$ and $\deg(\alpha, B_a\phi) = \deg(\phi_1) + \dots + \deg(\phi_k) + \deg(\phi) + 1$. Hence, $\deg(\alpha, \perp) < \deg(\alpha, B_a\phi)$ and $(\alpha, \perp) \ll (\alpha, B_a\phi)$.
10. Note that $\deg(\alpha, \perp) \leq \deg(\alpha, \neg\phi)$. Moreover, since $\|\alpha\| + 1 < \|\alpha\| + 1 + \|\phi\|$, we obtain that $(\alpha, \perp) \ll (\alpha, \neg\phi)$. ■

6.2.3 Semantics of asynchronous announcement logic

Definition 6.23 (Models for AA). A *model* is a triple (W, R, V) , where W is a nonempty set of states, $R = \{R_a : a \in A\}$ is a family of binary accessibility relations on W , and $V : \text{Prop} \rightarrow \mathcal{P}(W)$ maps each propositional variable p to the set $V(p)$ of states in W where p is true.

Definition 6.24 (Semantics). Given a model (W, R, V) , we simultaneously define the *agreement relation* \bowtie between states and histories and the *satisfaction relation* \models between pairs of states and histories, and formulas. The model is left implicit in these relations.

$$\begin{aligned}
s \bowtie \epsilon & \\
s \bowtie \alpha.a & \quad \text{iff } s \bowtie \alpha \text{ and } |\alpha|_a < |\alpha|_1 \\
s \bowtie \alpha.\psi & \quad \text{iff } s \bowtie \alpha \text{ and } s, \alpha \models \psi \\
\\
s, \alpha \models p & \quad \text{iff } s \in V(p) \\
s, \alpha \not\models \perp & \\
s, \alpha \models \neg\phi & \quad \text{iff } s, \alpha \not\models \phi \\
s, \alpha \models \phi \vee \psi & \quad \text{iff } s, \alpha \models \phi \text{ or } s, \alpha \models \psi \\
s, \alpha \models B_a\phi & \quad \text{iff } t, \beta \models \phi \text{ for all } t \in W \text{ and for all histories } \beta \\
& \quad \text{such that } sR_at, \alpha \triangleright_a \beta, \text{ and } t \bowtie \beta \\
s, \alpha \models [\phi]\psi & \quad \text{iff } s, \alpha \models \phi \text{ implies } s, \alpha.\phi \models \psi \\
s, \alpha \models [a]\phi & \quad \text{iff } |\alpha|_a < |\alpha|_1 \text{ implies } s, \alpha.a \models \phi
\end{aligned}$$

A formula ϕ is ϵ -*valid* (or *valid*), notation $\models^\epsilon \phi$ (or $\models \phi$), iff for all models (W, R, V) and for all $s \in W$, $s, \epsilon \models \phi$. The set of validities is called AA^ϵ (or AA), for *asynchronous announcement logic*. A formula ϕ is **-valid* (or *always valid*), notation $\models^* \phi$, iff for all histories α , $\models [\alpha]\phi$; further, ϕ is ϵ -*satisfiable* (or *satisfiable*) iff there are (W, R, V) and $s \in S$ such that $s, \epsilon \models \phi$, and ϕ is **-satisfiable* (or *sometimes satisfiable*) if there is a history α such that $\langle \alpha \rangle \phi$ is ϵ -satisfiable. The set of *-validities is called AA^* .

Thanks to the items 1–7 of Lemma 6.22, the reader may verify that the definitions of the relations “agrees with” and “satisfies” are well-founded.

Importantly, the meaning of $[\phi]\psi$ in AA is different from the meaning of $[\phi]\psi$ in PAL .

As dynamic epistemic logics go, AA is unusual because dynamic modalities do not result in model transformations. Such transformations are implicit in the history. Given a model $M = (W, R, V)$, we could somehow see the clause for announcement as a model transformer: as the truth of $[\phi]\psi$ is conditional on the truth of ϕ , the states in the domain W that “survive” this operation are exactly the ϕ -restriction, as in PAL . However, to interpret the ψ bound by announcement ϕ , we may have to access the model prior to that announcement. In that respect our models are rather like the protocol-generated forests of [25], however with the additional complication of uncertainty of reception of announcements by other agents, which is made precise in the $[a]\phi$ and $B_a\phi$ semantics.

6.2.4 Examples

We continue with examples of validities and non-validities.

Example 6.25. We have $\models [p][a]B_ap$. The formula $[p][a]B_ap$ stands for ‘after announcement of factual information p , and subsequent reception by agent a , agent a knows that p ’. To show that it is valid, is elementary. Take any (W, R, V) , $s \in W$. Then the following conditions are equivalent:

- $s, \epsilon \models [p][a]B_ap$,
- if $s, \epsilon \models p$ then $s, p \models [a]B_ap$,
- if $s, \epsilon \models p$ then $s, p.a \models B_ap$,
- if $s, \epsilon \models p$, then $t, \beta \models p$ for all $t \in W$ and for all histories β such that $sR_at, t \bowtie \beta$ and $p.a \triangleright_a \beta$.

If b is the only other agent, the possible histories β such that $p.a \triangleright_a \beta$ are: $p.a, p.b.a, p.a.b$. As they all contain the announcement p , the above conditions are true.

Example 6.26. On the other hand, $\not\models [p]B_ap$. Let p be true but not known to agent a , as in the model

$$s_{(p)} \xrightarrow{a} t_{(\neg p)}$$

Then $s, \epsilon \models p$, and therefore $s \bowtie p$. But we do not have $s, p \models B_ap$: from $p \triangleright_a \epsilon$, $R_ast, t \bowtie \epsilon$, and $t, \epsilon \models \neg p$ it follows that $s, p \models \hat{B}_a \neg p$. In fact, because $p \triangleright_a \beta$ iff $\epsilon \triangleright_a \beta$, we have that $[p]B_ap$ is equivalent to $p \rightarrow B_ap$.

Example 6.27. Next, $\not\models^* [p][a]B_ap$. Consider a model wherein agent a initially is uncertain about the truth of p and wherein in actual state s variables p and q are true with $p \neq q$, e.g.

$$s_{(p,q)} \xrightarrow{a} t_{(\neg p,q)}$$

We have $s, q \models \langle p \rangle \langle a \rangle \hat{B}_a \neg p$, seeing that $s, \epsilon \models q$, $s, q \models p$, $|q.p|_a < |q.p|_!$ and $s, q.p.a \models \hat{B}_a \neg p$. Differently said, given the history $q.p.a$, in the event wherein a receives “the next announcement,” it receives the information that q contained in the first announcement, not the information that p contained in the second announcement, which agent a will only receive next if she reads again.

Example 6.28. We have however $\models^* [a]\perp \rightarrow [p][a]B_ap$. For any state s and history α , $[a]\perp$ is only true in (s, α) if agent a has received all announcements in the history α (i.e., $|\alpha|_! = |\alpha|_a$). This means that if a further announcement is made, such as p , and a then receives ‘the next announcement’, that must be the announcement of p just made. After that, B_ap is true. Similarly, $\models^* [p][a]([a]\perp \rightarrow B_ap)$. It will also be clear that $\models [a]\perp$, but $\not\models^* [a]\perp$.

Example 6.29. In Figure 6.1(iv), after Anxélica and Bertu have both received the announcement $p \vee q$, they both know $p \vee q$: $B_a(p \vee q) \wedge B_b(p \vee q)$ is now true. For this we can, as usual, write $E_{ab}(p \vee q)$ (everybody knows $p \vee q$). But they do not know that the other knows $p \vee q$. However, after the announcement of $E_{ab}(p \vee q)$ and both receiving it we obtain $E_{ab}^2(p \vee q)$: everybody knows that everybody knows $p \vee q$. And so on. Anxélica and Bertu can achieve any finite approximation of common knowledge, but they cannot get common knowledge of $p \vee q$.

With individually received messages no growth of common knowledge will ever occur, unlike in PAL where reception is synchronous [58, 76]. But we can gradually construct so-called *concurrent common knowledge* [58, 81], as above.

6.2.5 Validities and other results for the semantics

We continue with results relating the satisfaction relation and the agreement relation, and with some fairly general always-validities. In the following result, p being a propositional variable and χ, ϕ being formulas, $\chi[p/\phi]$ will denote the formula obtained from χ by replacing a specific occurrence of p (if any are present) by ϕ ; given a history α , the expression $\alpha[p/\phi]$ will denote the history obtained from α by replacing this specific occurrence by ϕ .

Lemma 6.30. *Let (W, R, V) be a model. Let p be a propositional variable. Let ϕ, ψ be formulas such that for all $s \in W$ and for all histories α , $s, \alpha \models \phi$ iff $s, \alpha \models \psi$. Let χ be a formula possibly containing a specific occurrence of p and γ be a history possibly containing a specific occurrence of p . For all $s \in W$, the following conditions hold:*

- $s \bowtie \gamma[p/\phi]$ iff $s \bowtie \gamma[p/\psi]$,
- $s, \gamma[p/\phi] \models \chi$ iff $s, \gamma[p/\psi] \models \chi$,
- $s, \gamma \models \chi[p/\phi]$ iff $s, \gamma \models \chi[p/\psi]$.

Proof. The proof is by \ll -induction on (γ, χ) . ■

Lemma 6.31. *Let (W, R, V) be a model. Let s be a state and let α be a history. If $s \bowtie \alpha$ and $\beta \sqsubseteq \alpha$, then $s \bowtie \beta$.*

Proof. The proof is by induction on $|\alpha|$. ■

Lemma 6.32. *Let (W, R, V) be a model. Let α be a history and β be a word. For every formula χ and for every world s such that $s \bowtie \alpha$, $s, \alpha \models \langle \beta \rangle \chi$ if and only if (i) the concatenation $\alpha.\beta$ is a history, (ii) $s \bowtie \alpha.\beta$, and (iii) $s, \alpha.\beta \models \chi$.*

Proof. The proof is by induction on $|\beta|$.

The case “ $\beta = \epsilon$ ” is trivial; for the case “ $\beta = a.\beta'$ ”, suppose $s, \alpha \models \langle a \rangle \langle \beta' \rangle \chi$. Hence, $|\alpha|_a < |\alpha|$, and therefore $\alpha.a$ is a history and $s \bowtie \alpha.a$. Moreover, $s, \alpha.a \models \langle \beta' \rangle \chi$. Consequently, by induction hypothesis, $\alpha.a.\beta'$ is a history, $s \bowtie \alpha.a.\beta'$ and $s, \alpha.a.\beta' \models \chi$. Conversely, suppose $\alpha.a.\beta'$ is a history, $s \bowtie \alpha.a.\beta'$ and $s, \alpha.a.\beta' \models \chi$. Hence, by induction hypothesis $s, \alpha.a \models \langle \beta' \rangle \chi$. Thus $s, \alpha \models \langle a \rangle \langle \beta' \rangle \chi$.

For the case “ $\beta = \psi.\beta'$ ”. Suppose $s, \alpha \models \langle \psi \rangle \langle \beta' \rangle \chi$. Hence, $s, \alpha \models \psi$ and $s, \alpha.\psi \models \langle \beta' \rangle \chi$. Thus, by induction hypothesis, $\alpha.\psi.\beta'$ is a history, $s \bowtie \alpha.\psi.\beta'$ and $s, \alpha.\psi.\beta' \models \chi$. Conversely, suppose $\alpha.\psi.\beta'$ is a history, $s \bowtie \alpha.\psi.\beta'$ and $s, \alpha.\psi.\beta' \models \chi$. Consequently, $s \bowtie \alpha.\psi$ and $s, \alpha \models \psi$. Moreover, by induction hypothesis, $s, \alpha.\psi \models \langle \beta' \rangle \chi$. Hence, $s, \alpha \models \langle \psi \rangle \langle \beta' \rangle \chi$. ■

Lemma 6.33. *Let (W, R, V) be a model, α be a history, β be a word over $A \cup \mathcal{L}_{AA}$ and ϕ be a formula. For all $s \in W$,*

- $s, \alpha \models \langle \beta \rangle \phi$ iff $s, \alpha \not\models [\beta] \neg \phi$,
- $s, \alpha \models [\beta] \phi$ iff $s, \alpha \not\models \langle \beta \rangle \neg \phi$.

Proof. Induction on $|\beta|$. ■

Lemma 6.34. *Let β be a word over $A \cup \mathcal{L}_{AA}$. For all models (W, R, V) , for all $s \in W$ and for all formulas ϕ ,*

1. $s, \epsilon \models \langle \beta \rangle \phi$ iff β is a history, $s \bowtie \beta$ and $s, \beta \models \phi$,
2. $s, \epsilon \models [\beta] \phi$ iff, if β is a history and $s \bowtie \beta$, then $s, \beta \models \phi$.

Proof. By Lemma 6.32 and Lemma 6.33. ■

Corollary 6.35. *For all $\phi \in \mathcal{L}_{AA}$, $\models^* \phi$ iff for all models (W, R, V) , for all $s \in W$ and for all histories α , if $s \bowtie \alpha$ then $s, \alpha \models \phi$.*

Proof. By Lemma 6.34. ■

We note that the formulation of Corollary 6.35 could well have served as an alternative *definition* of $*$ -validity, instead of “for all histories α , $\models [\alpha]\phi$.”

Lemma 6.36. *Let (W, R, V) be a model. For all histories α and for all states s , the following conditions are equivalent:*

1. $s \bowtie \alpha$,
2. for all histories β , for all words γ and for all formulas ϕ , if $\alpha = \beta.\phi.\gamma$ then $s, \beta \models \phi$.

Proof. Let α be a history and s be a state.

Suppose $s \bowtie \alpha$. Let β be a history, γ be a word and ϕ be a formula such that $\alpha = \beta.\phi.\gamma$. Since $s \bowtie \alpha$, $s \bowtie \beta.\phi$. Hence, $s \bowtie \beta$ and $s, \beta \models \phi$.

Conversely, suppose for all histories β , for all words γ and for all formulas ϕ , if $\alpha = \beta.\phi.\gamma$ then $s, \beta \models \phi$. We prove that $s \bowtie \alpha$ by $<$ -induction on $|\alpha|$.

Case “ $\alpha = \epsilon$ ”. Then $s \bowtie \alpha$.

Case “ $\alpha = \alpha'.a$ ”. Since α is a history, α' is a history such that $|\alpha'|_! \geq |\alpha'|_a$. Moreover, for all histories β , for all words γ and for all formulas ϕ , if $\alpha' = \beta.\phi.\gamma$ then $\alpha = \beta.\phi.\gamma.a$ and, by our hypothesis, $s, \beta \models \phi$. Thus, by induction hypothesis, $s \bowtie \alpha'$. Since $|\alpha'|_! \geq |\alpha'|_a$, $s \bowtie \alpha$.

Case “ $\alpha = \alpha'.\psi$ ”. Since α is a history, α' is a history. Moreover, by our hypothesis, $s, \alpha' \models \psi$. Consequently, $s \bowtie \alpha$. ■

We continue with some results for always-validity \models^* .

Proposition 6.37. *Let $\phi \in \mathcal{L}_{AA}$. If $\models^* \phi$ then $\models \phi$.*

Proof. Suppose $\models^* \phi$. Hence, $\models [\epsilon]\phi$. Thus, $\models \phi$. ■

Proposition 6.38. *Let $\phi, \psi \in \mathcal{L}_{AA}$ and $a \in A$. Then:*

1. $\models^* \phi$ implies $\models^* B_a\phi$,
2. $\models^* B_a(\phi \rightarrow \psi) \rightarrow (B_a\phi \rightarrow B_a\psi)$.

Proof.

1. Suppose $\not\models^* B_a\phi$. Hence, by Corollary 6.35, let (W, R, V) be a model, α be a history and $s \in W$ be such that $s \bowtie \alpha$ and $s, \alpha \not\models B_a\phi$. Let $t \in W$ and β be a history such that $sRat$, $\alpha \triangleright_a \beta$, $t \bowtie \beta$ and $t, \beta \not\models \phi$. Thus, by Lemma 6.34, $t, \epsilon \not\models [\beta]\phi$. Consequently, $\not\models^* \phi$.

2. Suppose $\not\models^* B_a(\phi \rightarrow \psi) \rightarrow (B_a\phi \rightarrow B_a\psi)$. Hence, by Corollary 6.35, let (W, R, V) be a model, α be a history and $s \in W$ be such that $s \bowtie \alpha$ and $s, \alpha \not\models B_a(\phi \rightarrow \psi) \rightarrow (B_a\phi \rightarrow B_a\psi)$. Thus, $s, \alpha \models B_a(\phi \rightarrow \psi)$, $s, \alpha \models B_a\phi$ and $s, \alpha \not\models B_a\psi$. Let $t \in W$ and β be a history such that sR_at , $\alpha \triangleright_a \beta$, $t \bowtie \beta$ and $t, \beta \not\models \psi$. Since $s, \alpha \models B_a(\phi \rightarrow \psi)$ and $s, \alpha \models B_a\phi$, we obtain that $t, \beta \models \phi \rightarrow \psi$ and $t, \beta \models \phi$. Thus, $t, \beta \models \psi$: a contradiction. ■

Lemma 6.39. *Let ϕ, ψ be formulas such that $\models^* \phi \leftrightarrow \psi$. Let $M = (W, R, V)$ be a model. Let α be a history. Let β be a word such that $\alpha.\phi.\beta$ and $\alpha.\psi.\beta$ are histories and let χ be a formula. For all states $s \in W$, $s \bowtie \alpha.\phi.\beta$ iff $s \bowtie \alpha.\psi.\beta$, and $s, \alpha.\phi.\beta \models \chi$ iff $s, \alpha.\psi.\beta \models \chi$.*

Proof. By \ll -induction on (β, χ) . ■

Proposition 6.40 (Substitution of equivalents). *Let ϕ, ψ be formulas such that $\models^* \phi \leftrightarrow \psi$. For all formulas χ and for all atoms p , $\models^* \chi(p/\phi) \leftrightarrow \chi(p/\psi)$.*

Proof. The proof is by induction on χ .

Cases “ χ is an atom”, “ $\chi = \perp$ ”, “ $\chi = \neg\chi'$ ” and “ $\chi = \chi_1 \vee \chi_2$ ”. Left to the reader.

Case “ $\chi = [\eta]\chi'$ ”. Let (W, R, V) be a model, s be a state and α be a history such that $s \bowtie \alpha$. We have: $s, \alpha \models [\eta(p/\phi)]\chi'(p/\phi)$ iff $s, \alpha \models \eta(p/\phi)$ implies $s, \alpha.\eta(p/\phi) \models \chi'(p/\phi)$ iff, by induction hypothesis and using Lemma 6.39, $s, \alpha \models \eta(p/\psi)$ implies $s, \alpha.\eta(p/\psi) \models \chi'(p/\psi)$ iff $s, \alpha \models [\eta(p/\psi)]\chi'(p/\psi)$. Since (W, R, V) , s and α were arbitrary, $\models^* \chi(p/\phi) \leftrightarrow \chi(p/\psi)$.

Case “ $\chi = B_a\chi'$ ”. Let (W, R, V) be a model, s be a state and α be a history such that $s \bowtie \alpha$. We have: $s, \alpha \models B_a\chi'(p/\phi)$ iff for all states t and for all histories β , if sR_at , $\alpha \triangleright_a \beta$ and $t \bowtie \beta$ then $t, \beta \models \chi'(p/\phi)$ iff, by induction hypothesis, for all states t and for all histories β , if sR_at , $\alpha \triangleright_a \beta$ and $t \bowtie \beta$ then $t, \beta \models \chi'(p/\psi)$ iff $s, \alpha \models B_a\chi'(p/\psi)$. Since (W, R, V) , s and α were arbitrary, $\models^* \chi(p/\phi) \leftrightarrow \chi(p/\psi)$. ■

Lemma 6.41. *Let (W, R, V) be a model, s be a state and a be an agent. For all histories α, β , if $\alpha \triangleright_a \beta$, $s \bowtie \alpha$ and $s \bowtie \beta$ then $s, \alpha \models B_a\phi$ iff $s, \beta \models B_a\phi$.*

Proof. By Proposition 6.12. ■

In particular, it follows that $s, \alpha \models B_a \phi$ iff $s, \alpha_a \models B_a \phi$. Let us now see some results concerning the positive fragment $\mathcal{L}_{\text{EL}}^+$. For this we will define a preorder \preceq on histories as follows:

Definition 6.42. $\alpha \preceq \beta$ if and only if:

- $\alpha \upharpoonright_! \sqsubseteq \beta \upharpoonright_!$;
- for all $a \in A$, $|\alpha|_a \leq |\beta|_a$;
- for every model (W, R, V) and every state s , $s \bowtie \beta$ implies $s \bowtie \alpha$.

It is easy to see that \preceq is a reflexive and transitive relation between histories.

Lemma 6.43. *Let histories α, β and $a \in A$ be given.*

1. $\alpha \sqsubseteq \beta$ implies $\alpha \preceq \beta$,
2. $\alpha_a \preceq \alpha$.

Proof. (1). Suppose $\alpha \sqsubseteq \beta$. Hence, $\alpha \upharpoonright_! \sqsubseteq \beta \upharpoonright_!$ and for all $a \in A$, $|\alpha|_a \leq |\beta|_a$. Now, let (W, R, V) be a model and s be a state such that $s \bowtie \beta$. Thus, by Lemma 6.31, $s \bowtie \alpha$. Since (W, R, V) and s were arbitrary, $\alpha \preceq \beta$.

(2). From the construction of α_a (see Section 6.1.1) it follows that $\alpha_a = \gamma.a^n$ for some history γ such that $\gamma \sqsubseteq \alpha$ and $n = |\gamma|_! - |\gamma|_a$. Moreover, $\alpha \upharpoonright_! = \alpha_a \upharpoonright_!$, $|\alpha|_a = |\alpha_a|_a$. Now, let (W, R, V) be a model and s be a state such that $s \bowtie \alpha$. By Lemma 6.31, we have $s \bowtie \gamma$. Since $\alpha_a = \gamma.a^n$ and $n = |\gamma|_! - |\gamma|_a$, $s \bowtie \alpha_a$. Since $\alpha \upharpoonright_! = \alpha_a \upharpoonright_!$ and $|\alpha|_a = |\alpha_a|_a$, we have that $\alpha_a \preceq \alpha$. ■

Lemma 6.44. *Let α, β and γ be histories. If $\gamma \preceq \alpha$ and $\alpha \triangleright_a \beta$ then there exists a history δ such that $\gamma \triangleright_a \delta$ and $\delta \preceq \beta$. In other words, $(\preceq \circ \triangleright_a) \subseteq (\triangleright_a \circ \preceq)$.*

Proof. Suppose $\gamma \preceq \alpha$ and $\alpha \triangleright_a \beta$. Hence, $\gamma \upharpoonright_! \sqsubseteq \alpha \upharpoonright_!$ and $|\gamma|_a \leq |\alpha|_a$. Moreover, $|\beta|_a = |\alpha|_a$, $\beta \upharpoonright_{!a} = \alpha \upharpoonright_{!a}$ and $|\beta|_! = |\alpha|_!$. Thus, $|\gamma|_a \leq |\beta|_a$. Let β' be the initial segment of β up until the $|\gamma|_a$ -th occurrence of a . Consequently, $|\beta'|_a = |\gamma|_a$. Moreover, $\beta' \sqsubseteq \beta$. Hence, by Lemma 6.43, $\beta' \preceq \beta$. Let $\delta = \beta'_a$. Thus, by Lemma 6.11, $\beta' \upharpoonright_{!a} = \delta \upharpoonright_{!a} = \delta \upharpoonright_!$ and $|\delta|_a = |\beta'|_a$. Moreover, by Lemma 6.43, $\delta \preceq \beta'$. Since $\beta' \preceq \beta$, $\delta \preceq \beta$. Since $|\beta'|_a = |\gamma|_a$ and $|\delta|_a = |\beta'|_a$, we obtain that $|\delta|_a = |\gamma|_a$. From all this, it follows that $\gamma \triangleright_a \delta$ and $\delta \preceq \beta$. ■

Lemma 6.45. *Let (W, R, V) be a model. Let $\phi \in \mathcal{L}_{\text{EL}}^+$. For all states s and for all histories α', α , if $\alpha' \preceq \alpha$, $s \bowtie \alpha$ and $s, \alpha' \models \phi$ then $s, \alpha \models \phi$.*

Proof. The proof is by induction on ϕ .

Cases “ $\phi = p$ ”, “ $\phi = \neg p$ ”, “ $\phi = \perp$ ”, “ $\phi = \top$ ”, “ $\phi = \psi \vee \chi$ ” and “ $\phi = \psi \wedge \chi$ ”. Left to the reader.

Case “ $\phi = B_a \psi$ ”. Let s be a state and α', α be histories such that $\alpha' \preceq \alpha$, $s \bowtie \alpha$, $s, \alpha' \models B_a \psi$ and $s, \alpha \not\models B_a \psi$. Let t be a state and β be a history such that $s R_a t$, $\alpha \triangleright_a \beta$, $t \bowtie \beta$ and $t, \beta \not\models \psi$. Since $\alpha' \preceq \alpha$, by Lemma 6.44, let β' be a history such that $\alpha' \triangleright_a \beta'$ and $\beta' \preceq \beta$. Since $t \bowtie \beta$, $t \bowtie \beta'$. Since $s, \alpha' \models B_a \psi$, $s R_a t$ and $\alpha' \triangleright_a \beta'$, we obtain that $t, \beta' \models \psi$. Since $\beta' \preceq \beta$ and $t \bowtie \beta$, by induction hypothesis, $t, \beta \models \psi$: a contradiction. ■

With this lemma in hand, we can now easily demonstrate that:

Proposition 6.46 (Positive is preserved). *For all $\phi \in \mathcal{L}_{\text{EL}}^+$ and words α , $\models^* \phi \rightarrow [\alpha]\phi$.*

Proof. Let $\phi \in \mathcal{L}_{\text{EL}}^+$ and α be a word such that $\not\models^* \phi \rightarrow [\alpha]\phi$. Hence, by Corollary 6.35, let (W, R, V) be a model, s be a state and β be a history such that $s \bowtie \beta$ and $s, \beta \not\models \phi \rightarrow [\alpha]\phi$. Thus, $s, \beta \models \phi$ and $s, \beta \not\models [\alpha]\phi$. Consequently, by Lemma 6.32 and Lemma 6.33, the concatenation $\beta\alpha$ is a history, $s \bowtie \beta\alpha$, and $s, \beta\alpha \not\models \phi$. Hence, $\beta \subseteq \beta\alpha$. Thus, by Lemma 6.43, $\beta \preceq \beta\alpha$. Since $s \bowtie \beta\alpha$ and $s, \beta \models \phi$, by Lemma 6.45, $s, \beta\alpha \models \phi$: a contradiction. ■

After these diverse reflections on the semantics of Asynchronous Announcement Logic, we continue with its axiomatization.

6.3 The logic AA

In Subsection 6.3.1 we axiomatize AA^ϵ , the logic of ϵ -validities. This is a reduction system eliminating reception and announcement modalities. In Subsection 6.3.2 we determine $*$ -validities that are reduction axioms. However, we do not axiomatize AA^* .

6.3.1 Axiomatization of AA

In this section, we present an axiomatization of AA on the class of all models with empty histories. We prove its completeness by showing that for all formulas $\phi \in \mathcal{L}_{\text{AA}}$, there exists an announcement-free and reading-free formula $\psi \in \mathcal{L}_{\text{EL}}$ such that $\phi \leftrightarrow \psi$ is valid in the class of all models with empty histories. In other words, the dynamic modalities $[a]$ and $[\phi]$ can be eliminated from the language, as far as one is concerned with ϵ -validity. We will do this

by using an truth preserving transformation tr . The completeness proof therefore consists in showing that \mathcal{L}_{AA} is equally expressive as \mathcal{L}_{EL} on the class of all models with empty histories. Similar results are well-known for PAL, but one might find them surprising for its asynchronous version.

Definition 6.47. For all words α over $A \cup \mathcal{L}_{AA}$ and for all \mathcal{L}_{AA} -formulas ϕ , we inductively define the \mathcal{L}_{EL} -formula $tr(\alpha, \phi)$ as follows:

- $tr(\epsilon, \perp) = \perp$,
- $tr(\alpha.a, \perp) = tr(\alpha, \perp)$ if $|\alpha|_a < |\alpha|_!$,
- $tr(\alpha.a, \perp) = \top$ if $|\alpha|_a \geq |\alpha|_!$,
- $tr(\alpha.\phi, \perp) = tr(\alpha, \phi) \rightarrow tr(\alpha, \perp)$,
- $tr(\alpha, p) = tr(\alpha, \perp) \vee p$,
- $tr(\alpha, \neg\phi) = tr(\alpha, \phi) \rightarrow tr(\alpha, \perp)$,
- $tr(\alpha, \phi \vee \psi) = tr(\alpha, \phi) \vee tr(\alpha, \psi)$,
- $tr(\alpha, B_a\phi) = tr(\alpha, \perp) \vee \bigwedge \{B_a tr(\beta, \phi) \mid \alpha \triangleright_a \beta\}$,
- $tr(\alpha, [a]\phi) = tr(\alpha.a, \phi)$,
- $tr(\alpha, [\phi]\psi) = tr(\alpha.\phi, \psi)$.

Let us remark that the above definition of the truth preserving translation tr is indeed inductive, namely with respect to the well-founded order \ll between (history, formula) pairs defined in Section 6.2.1 (Lemma 6.22).

Lemma 6.48. *Let (W, R, V) be a model and $s \in W$. For all words α over $A \cup \mathcal{L}_{AA}$ and for all formulas ϕ , $s, \epsilon \models tr(\alpha, \phi)$ iff $s, \epsilon \models [\alpha]\phi$.*

Proof. The proof is done by \ll -induction on (α, ϕ) . In all cases we only show the proof direction from left to right, as the other direction can be shown in a similar way.

Case $\alpha = \epsilon$ and $\phi = \perp$. Obviously, $s, \epsilon \not\models tr(\epsilon, \perp)$ and $s, \epsilon \not\models [\epsilon]\perp$.

Case $\alpha = \beta.a$ and $\phi = \perp$. Suppose $s, \epsilon \models tr(\beta.a, \perp)$. Hence, $s, \epsilon \models tr(\beta, \perp)$ and $|\beta|_a < |\beta|_!$, or $|\beta|_a \geq |\beta|_!$. In the former case, by induction hypothesis, $s, \epsilon \models [\beta]\perp$. Thus, by Lemma 6.34, $s \not\models \beta$. Consequently, $s \not\models \beta.a$. Hence, by Lemma 6.34 again, $s, \epsilon \models [\beta.a]\perp$. In the latter case, $s \not\models \beta.a$. Thus, by Lemma 6.34, $s, \epsilon \models [\beta.a]\perp$.

Case $\alpha = \beta.\psi$ and $\phi = \perp$. Suppose $s, \epsilon \models tr(\beta\psi, \perp)$. Hence, $s, \epsilon \not\models tr(\beta, \psi)$ or $s, \epsilon \models tr(\beta, \perp)$. In the former case, by induction hypothesis, $s, \epsilon \not\models [\beta]\psi$. Thus, by Lemma 6.34, $s \bowtie \beta$ and $s, \beta \not\models \psi$. Consequently, $s \not\bowtie \beta\psi$. Hence, by Lemma 6.34 again, $s, \epsilon \models [\beta\psi]\perp$. In the latter case, by induction hypothesis, $s, \epsilon \models [\beta]\perp$. Thus, by Lemma 6.34, $s \not\bowtie \beta$. Consequently, $s \not\bowtie \beta\psi$. Hence, by Lemma 6.34 again, $s, \epsilon \models [\beta\psi]\perp$.

Case $\phi = p$. Suppose $s, \epsilon \models tr(\alpha, p)$. Hence, $s, \epsilon \models tr(\alpha, \perp)$, or $s, \epsilon \not\models tr(\alpha, \perp)$ and $s, \epsilon \models p$. In the former case, by induction hypothesis, $s, \epsilon \models [\alpha]\perp$. Thus, by Lemma 6.34, $s \not\bowtie \alpha$. Consequently, by Lemma 6.34 again, $s, \epsilon \models [\alpha]p$. In the latter case, by induction hypothesis, $s, \epsilon \not\models [\alpha]\perp$. Moreover, $s \in V(p)$. Hence, by Lemma 6.34, $s \bowtie \alpha$. Moreover, $s, \alpha \models p$. Thus, by Lemma 6.34 again, $s, \epsilon \models [\alpha]p$.

Case $\phi = \neg\psi$. Suppose $s, \epsilon \models tr(\alpha, \neg\psi)$. Hence, $s, \epsilon \not\models tr(\alpha, \psi)$ or $s, \epsilon \models tr(\alpha, \perp)$. In the former case, by induction hypothesis, $s, \epsilon \not\models [\alpha]\psi$. Thus, by Lemma 6.34, $s \bowtie \alpha$ and $s, \alpha \not\models \psi$. Consequently, by Lemma 6.34 again, $s, \epsilon \models [\alpha]\neg\psi$. In the latter case, by induction hypothesis, $s, \epsilon \models [\alpha]\perp$. Hence, by Lemma 6.34, $s \not\bowtie \alpha$. Thus, by Lemma 6.34 again, $s, \epsilon \models [\alpha]\neg\psi$.

Case $\phi = \psi \vee \chi$. Suppose $s, \epsilon \models tr(\alpha, \psi \vee \chi)$. Hence, $s, \epsilon \models tr(\alpha, \psi)$ or $s, \epsilon \models tr(\alpha, \chi)$. In the former case, by induction hypothesis, $s, \epsilon \models [\alpha]\psi$. Thus, by Lemma 6.34, $s \not\bowtie \alpha$ or $s, \alpha \models \psi$. Consequently, $s \not\bowtie \alpha$ or $s, \alpha \models \psi \vee \chi$. Hence, by Lemma 6.34 again, $s, \epsilon \models [\alpha](\psi \vee \chi)$. The latter case is similarly treated.

Case $\phi = B_a\psi$. Suppose $s, \epsilon \models tr(\alpha, B_a\psi)$. Hence, $s, \epsilon \models tr(\alpha, \perp)$ or $s, \epsilon \models \bigwedge\{B_a tr(\beta, \psi) \mid \alpha \triangleright_a \beta\}$. In the former case, by induction hypothesis, $s, \epsilon \models [\alpha]\perp$. Thus, by Lemma 6.34, $s \not\bowtie \alpha$. Consequently, by Lemma 6.34 again, $s, \epsilon \models [\alpha]B_a\psi$. In the latter case, for the sake of the contradiction, suppose $s, \epsilon \not\models [\alpha]B_a\psi$. Thus, by Lemma 6.34, $s \bowtie \alpha$ and $s, \alpha \not\models B_a\psi$. Let $t \in W$ and γ be a history such that sR_at , $\alpha \triangleright_a \gamma$, $t \bowtie \gamma$ and $t, \gamma \not\models \psi$. Since $s, \epsilon \models \bigwedge\{B_a tr(\beta, \psi) \mid \alpha \triangleright_a \beta\}$, we obtain that $s, \epsilon \models B_a tr(\gamma, \psi)$. We recall the reader that sR_at . Moreover, obviously, $\epsilon \triangleright_a \epsilon$ and $t \bowtie \epsilon$. Consequently, $t, \epsilon \models tr(\gamma, \psi)$. Hence, by induction hypothesis, $t, \epsilon \models [\gamma]\psi$. Thus, by Lemma 6.34, $t \not\bowtie \gamma$ or $t, \gamma \models \psi$: a contradiction. Consequently, $s, \epsilon \models [\alpha]B_a\psi$.

Case $\phi = [a]\psi$. Suppose $s, \epsilon \models tr(\alpha, [a]\psi)$. Hence, $s, \epsilon \models tr(\alpha a, \psi)$. Thus, by induction hypothesis, $s, \epsilon \models [\alpha a]\psi$. Consequently, $s, \epsilon \models [\alpha][a]\psi$.

Case $\phi = [\psi]\chi$. Suppose $s, \epsilon \models tr(\alpha, [\psi]\chi)$. Hence, $s, \epsilon \models tr(\alpha\psi, \chi)$. Thus, by induction hypothesis, $s, \epsilon \models [\alpha\psi]\chi$. Consequently, $s, \epsilon \models [\alpha][\psi]\chi$. ■

In particular, for $\alpha = \epsilon$ and given that $[\epsilon]\phi = \phi$, we obtain that for all models, states and formulas: $s, \epsilon \models tr(\epsilon, \phi)$ iff $s, \epsilon \models \phi$, so that $\phi \mapsto tr(\epsilon, \phi)$ therefore defines a truth (value) preserving translation from \mathcal{L}_{AA} to \mathcal{L}_{EL} .

Corollary 6.49 (Elimination of dynamic modalities).

For all $\phi \in \mathcal{L}_{AA}$ there is a $\psi \in \mathcal{L}_{EL}$ (namely $\psi = tr(\epsilon, \phi)$) such that $\models \phi \leftrightarrow \psi$.

With these results in hand we will now present the axiomatization AA. The axioms of AA exactly follow the pattern of the translation function tr of Def. 6.47.

Definition 6.50 (Axiomatization AA). Let AA be the axiomatization given by the following axioms and inference rules:

- the tautologies in the language \mathcal{L}_{AA} ,
- the theorems of the least normal modal logic in the language \mathcal{L}_{EL} ,
- the following axioms:
 - (A1): $[\alpha]p \leftrightarrow [\alpha]\perp \vee p$,
 - (A2): $[\alpha.a]\perp \leftrightarrow [\alpha]\perp$ if $|\alpha|_a < |\alpha|_1$,
 - (A3): $[\alpha.a]\perp$ if $|\alpha|_a \geq |\alpha|_1$,
 - (A4): $[\alpha.\phi]\perp \leftrightarrow [\alpha]\perp \vee \neg[\alpha]\phi$,
 - (A5): $[\alpha]\neg\phi \leftrightarrow [\alpha]\perp \vee \neg[\alpha]\phi$,
 - (A6): $[\alpha](\phi \vee \psi) \leftrightarrow [\alpha]\phi \vee [\alpha]\psi$,
 - (A7): $[\alpha]B_a\phi \leftrightarrow [\alpha]\perp \vee \bigwedge\{B_a[\beta]\phi \mid \alpha \triangleright_a \beta\}$.
- the following inference rules:
 - (MP): from ϕ and $\phi \rightarrow \psi$ infer ψ ,
 - (R2): from $\phi \leftrightarrow \psi$ infer $B_a\phi \leftrightarrow B_a\psi$,

The notion of AA-proof being defined as usual, we will say that a formula ϕ is AA-derivable (denoted $\vdash \phi$) iff there exists a proof of ϕ from the above axiomatization.

Lemma 6.51. For all words α over $A \cup \mathcal{L}_{AA}$ and for all formulas ϕ , $\vdash [\alpha]\phi \leftrightarrow tr(\alpha, \phi)$.

Proof. The proof is by \ll -induction on (α, ϕ) . ■

Also note the following:

Lemma 6.52. *Let $\psi \in \mathcal{L}_{\text{EL}}$ be an announcement-free formula, let $M = (W, R, V)$ be a model and let $s \in W$.*

Then $M, s \models \psi$ with the usual relational semantics if and only if $s, \epsilon \models \psi$ with the AA semantics.

Theorem 6.53 (AA is sound and complete).

For all $\phi \in \mathcal{L}_{\text{AA}}$, $\vdash \phi$ iff $\models \phi$.

Proof. The soundness of AA ($\vdash \phi$ implies $\models \phi$) follows from Lemma 6.48, wherein it is shown that the translation tr is truth preserving, and thus that all the axioms are sound.

We now show the completeness ($\models \phi$ implies $\vdash \phi$). Suppose $\not\vdash \phi$. Let $\psi = tr(\epsilon, \phi)$. Since $\not\vdash \phi$, by Lemma 6.51, $\not\vdash \psi$. Since ψ is a formula in \mathcal{L}_{EL} , by the standard completeness of the least normal modal logic in the language \mathcal{L}_{EL} , there is some model M and some world in s such that $s \not\models \psi$. Hence, by Lemma 6.52, $s, \epsilon \not\models \psi$. Thus, by Lemma 6.48, $\not\models \phi$. ■

Let us remark that, as for \mathcal{L}_{PAL} , we now have for \mathcal{L}_{AA} an effective way to determine whether a given ϕ is ϵ -valid (for the class of models with arbitrary relations): if $tr(\epsilon, \phi)$ is a theorem in the minimal modal logic \mathbf{K} , ϕ is ϵ -valid; otherwise, ϕ is not ϵ -valid. This makes it fairly easy to prove the decidability of AA.

Proposition 6.54. *AA has the finite model property.*

Proof. Suppose ϕ is satisfiable. Let $M = (W, R, V)$ and state $s \in W$ be such that $s, \epsilon \models \phi$. By Lemma 6.51, this means that $s, \epsilon \models tr(\epsilon, \phi)$. Since $tr(\epsilon, \phi) \in \mathcal{L}_{\text{EL}}$, by Lemma 6.52, this gives $M, s \models tr(\epsilon, \phi)$ in the usual relational semantics. As the minimal modal logic \mathbf{K} has the finite model property, there exists a finite model M^f and a world v in M^f such that $M^f, v \models tr(\epsilon, \phi)$. By the same reasoning, this means that in M^f we have $v, \epsilon \models tr(\epsilon, \phi)$ according to the AA semantics, and thus, again by Lemma 6.51, $v, \epsilon \models \phi$. ■

Since AA has a finitary axiomatization and the finite model property we directly obtain decidability.

Corollary 6.55. *AA is decidable.*

6.3.2 Reduction axioms for AA*

In this section we determine always-validities ($*$ -validities) that have the shape of reduction axioms for announcement. This is instructive, because they resemble the reduction axioms of PAL. However, these reductions cannot provide

an complete axiomatization as in the previous section. Although we can eliminate the dynamic modalities from ϵ -validities, as formulated in Corollary 6.49, we cannot eliminate dynamic modalities from $*$ -validities.

The proof of the next result uses the following Lemma:

Lemma 6.56. *Let $\psi \in \mathcal{L}_{\text{EL}}$ be a formula without announcement or reading modalities. Then the following are equivalent:*

- i. ψ is $*$ -valid;*
- ii. ψ is ϵ -valid;*
- iii. ψ is valid in the standard relational semantics.*

Proof. The implication (i) \Rightarrow (ii) is trivial, whereas the equivalence (ii) \Leftrightarrow (iii) follows from the fact that, if $\epsilon \triangleright_a \beta$, then $\epsilon = \beta$.

For (ii) \Rightarrow (i), let ψ be an ϵ -valid formula, α be a history and w be a world in some model $M = (W, R, V)$ such that $w \triangleright_a \alpha$. Let us construct the model $M' = (W', \mathbf{R}, V')$ as follows:

- $W' = \{(t, \beta) : t \in W, \beta \in \mathcal{H}, t \triangleright_a \beta\}$;
- $(t, \beta) \mathbf{R}_a(s, \gamma)$ iff $t R_a s$ and $\beta \triangleright_a \gamma$;
- $V'(p) = \{(t, \beta) \in W' : t \in V(p)\}$.

This is a standard relational model and thus, on account of the equivalence (ii) \Leftrightarrow (iii), we have that $M', (w, \alpha) \models \psi$ via the standard relational semantics.

The fact that this implies that $M, w, \alpha \models \psi$ in the AA semantics is demonstrated via an induction on ψ which is left to the reader. ■

With this:

Proposition 6.57 (Failure of elimination of dynamic modalities).

There are $\phi \in \mathcal{L}_{\text{AA}}$ such that for all $\psi \in \mathcal{L}_{\text{EL}} \not\models^ \phi \leftrightarrow \psi$.*

Proof. Consider the formula $[a]\perp \in \mathcal{L}_{\text{AA}}$. Suppose towards a contradiction that there is a $\psi \in \mathcal{L}_{\text{EL}}$ such that $\models^* [a]\perp \leftrightarrow \psi$. Then in particular we have that $\models [a]\perp \leftrightarrow \psi$. As $\models [a]\perp$, it follows that $\models \psi$: ψ is ϵ -valid. As $\psi \in \mathcal{L}_{\text{EL}}$, by Lemma 6.56 we obtain $\models^* \psi$: ψ is $*$ -valid.

Now take any $M = (W, R, V)$ and $s \in W$. From $\models \psi$ then follows $s, \epsilon \models \psi$. Also, obviously, $s \triangleright_a \top$ and $s, \top \models \neg[a]\perp$. From that and $\models^* [a]\perp \leftrightarrow \psi$ then follows that $s, \top \models \neg\psi$. However, from $\models^* \psi$ we obtain $s, \top \models \psi$. We have the required contradiction. ■

6.4 Equivalence relations

In this section we give results for the class $\mathcal{S5}$ of models where all accessibility relations are equivalence relations, we consider an alternative for the ‘view’-relation resulting in asynchronous *knowledge* instead of asynchronous belief, we motivate the belief semantics by a detailed example involving belief as acknowledgement, we relate the AA semantics to history-based semantics, we present the results of the related asynchronous broadcast logic, and we compare our histories containing announcements and receptions to the cuts of distributed computing.

6.4.1 Asynchronous announcement logic on the class $\mathcal{S5}$

In this section we restrict the models (W, R, V) to those where all accessibility relations R_a are equivalence relations. We call the class of such models $\mathcal{S5}$, and in that case $B_a\phi$ stands for ‘the agent knows ϕ ’, and, in standard Kripke semantics \models^K , the operator then satisfies the so-called *properties of knowledge* $B_a\phi \rightarrow \phi$ (T, factivity), $B_a\phi \rightarrow B_aB_a\phi$ (4, positive introspection), and $\neg B_a\phi \rightarrow B_a\neg B_a\phi$ (5, negative introspection). These properties correspond to, respectively, the facts that the accessibility relation R_a is reflexive, transitive and Euclidean.

The *properties of belief* (also known as *introspective belief*) are as the properties of knowledge, except that $B_a\phi \rightarrow \phi$ is replaced by $B_a\phi \rightarrow \hat{B}_a\phi$ (D, consistency), which corresponds to seriality of underlying frames. The models with serial, transitive and Euclidean relations are known as $\mathcal{KD45}$ models.

We very straightforwardly see that the properties of knowledge for B_a are ϵ -valid in the class $\mathcal{S5}$:

Proposition 6.58. *Let $\phi \in \mathcal{L}_{AA}$. Then*

- $\mathcal{S5} \models B_a\phi \rightarrow \phi$
- $\mathcal{S5} \models B_a\phi \rightarrow B_aB_a\phi$
- $\mathcal{S5} \models \neg B_a\phi \rightarrow B_a\neg B_a\phi$

Proof. Let (W, R, V) and $s \in W$ be given. Then $s, \epsilon \models B_a\phi$ iff $t, \beta \models \phi$ for all t, β such that sR_at , $\epsilon \triangleright_a \beta$, and $t \bowtie \beta$. As $\text{view}_a(\epsilon) = \{\epsilon\}$, and $t \bowtie \epsilon$ holds by definition, we get that: $s, \epsilon \models B_a\phi$ iff $t, \epsilon \models \phi$ for all t such that sR_at . As R_a is an equivalence relation, B_a therefore satisfies the three properties of knowledge. \blacksquare

In asynchronous announcement logic interpreted on $\mathcal{S5}$ models with history, the B_a operator does not satisfy all the properties of knowledge (and for this reason we write B_a and not K_a for this modality). For example, if agents a, b are initially both uncertain about p and this is common knowledge, as in the $\mathcal{S5}$ model

$$s_{(p)} \xrightarrow{a, b} t_{(\neg p)}$$

and the announcement of p is made and received by a but not yet by b , then $s, p.a \models B_a p \wedge B_b \neg B_a p$: the beliefs of agent b are incorrect. In general, whenever $|\alpha|_! > |\alpha|_a$, then agent a has not yet received all announcements and may therefore hold incorrect beliefs. If $|\alpha|_! > |\alpha|_a$ then it is not the case that $\alpha \triangleright_a \alpha$: the view relation \triangleright_a is not reflexive.

However, all the other properties of introspective belief hold for asynchronous announcement logic interpreted on $\mathcal{S5}$ models.

Proposition 6.59. *Let $\phi \in \mathcal{L}_{AA}$. Then:*

- $\mathcal{S5} \models^* B_a \phi \rightarrow \neg B_a \neg \phi$
- $\mathcal{S5} \models^* B_a \phi \rightarrow B_a B_a \phi$
- $\mathcal{S5} \models^* \neg B_a \phi \rightarrow B_a \neg B_a \phi$

Proof. We recall that $s, \alpha \models B_a \phi$ iff $t, \beta \models \phi$ for all t, β such that $sR_a t$, $\alpha \triangleright_a \beta$, and $t \bowtie \beta$, and that the view relation is defined as $\alpha \triangleright_a \beta$ iff $|\beta|_a = |\alpha|_a$, $\beta \upharpoonright_{!a} = \alpha \upharpoonright_{!a}$, and $|\beta|_! = |\alpha|_a$.

We show that the relation \mathbf{R}_a defined on the set of pairs (s, α) with $s \bowtie \alpha$ as follows:

$$(s, \alpha) \mathbf{R}_a (t, \beta) \quad \text{iff} \quad sR_a t, \alpha \triangleright_a \beta, \text{ and } t \bowtie \beta$$

is introspective (i.e., transitive, Euclidean, and serial), so that B_a satisfies the three properties of belief.

Transitivity of \mathbf{R}_a follows from the transitivity of R_a and \triangleright_a (the parts involving \bowtie merely result in a restriction).

$$\begin{array}{l} \text{if} \quad sR_a t, \quad \alpha \triangleright_a \beta, \quad t \bowtie \beta, \\ \text{and} \quad tR_a u, \quad \beta \triangleright_a \gamma, \quad u \bowtie \gamma, \\ \text{then} \quad sR_a u, \quad \alpha \triangleright_a \gamma, \quad u \bowtie \gamma \end{array}$$

Since R_a and \triangleright_a are Euclidean, \mathbf{R}_a is Euclidean.

$$\begin{array}{l} \text{if} \quad sR_a t, \quad \alpha \triangleright_a \beta, \quad t \bowtie \beta, \\ \text{and} \quad sR_a u, \quad \alpha \triangleright_a \gamma, \quad u \bowtie \gamma, \\ \text{then} \quad tR_a u, \quad \beta \triangleright_a \gamma, \quad u \bowtie \gamma \end{array}$$

Seriality of \mathbf{R}_a follows from the reflexivity of R_a and the seriality of \triangleright_a . For the latter, Proposition 6.12.1 showed that for any history α , $\alpha \triangleright_a \alpha_a$. The proof of Lemma 6.43.2 that $\alpha_a \preceq \alpha$ demonstrated that for any state s with $s \bowtie \alpha$ we also have $s \bowtie \alpha_a$. From $sR_a s$, $\alpha \triangleright_a \alpha_a$, and $s \bowtie \alpha_a$ we get that $(s, \alpha) \mathbf{R}_a (s, \alpha_a)$. ■

As $\mathcal{S5} \not\models^* B_a \phi \rightarrow \phi$, $*$ -satisfiability in $\mathcal{S5}$ does not entail Kripke satisfiability in $\mathcal{S5}$. The typical counterexample is the one at the beginning of this section: $p, p.a \models B_a p \wedge B_b \neg B_a p$. But $B_a p \wedge B_b \neg B_a p$ is not satisfiable in $\mathcal{S5}$. Beliefs are not factual in asynchronous announcement logic, nor would one want them to be.

Despite this, in the $\mathcal{S5} \models^*$ semantics some beliefs are, after all, correct, and thus knowledge. This may also come as a surprise:

Proposition 6.60 (Positive beliefs are correct). *Let $\phi \in \mathcal{L}_{\text{EL}}^+$. Then $\mathcal{S5} \models^* B_a \phi \rightarrow \phi$.*

Proof. Let s and α be a world and a history such that $s \bowtie \alpha$ and $s, \alpha \models B_a \phi$. From that and Lemma 6.41 it follows that $s, \alpha_a \models B_a \phi$. From $s \bowtie \alpha$ and Lemma 6.43.2 it follows $s \bowtie \alpha_a$. Then, from $s, \alpha_a \models B_a \phi$, $sR_a s$, $\alpha_a \triangleright_a \alpha_a$, and $s \bowtie \alpha_a$ it follows that $s, \alpha_a \models \phi$. Given that $\phi \in \mathcal{L}_{\text{EL}}^+$ and that $\alpha_a \preceq \alpha$, it follows from Lemma 6.45 $s, \alpha \models \phi$. ■

A possible interpretation of this result is the following: eventually all uncertainty about positive formulas may be resolved. At some stage an agent may well incorrectly believe that another agent is ignorant, notably when the other agent has already received some information unbeknownst to her, but eventually the first agent will also receive those messages and then change her incorrect beliefs into correct and stable beliefs: knowledge of positive formulas.

Let now $\text{AA}_{\mathcal{S5}}$ be the axiomatization formed by extending the axiomatization AA of Asynchronous Announcement Logic with the $\mathcal{S5}$ axioms (T), (4) and (5). Recalling the soundness and completeness of AA (Theorem 6.53), in view of Proposition 6.58 we immediately obtain:

Corollary 6.61 ($\text{AA}_{\mathcal{S5}}$ is sound and complete). *For all $\phi \in \mathcal{L}_{\text{AA}}$, $\text{AA}_{\mathcal{S5}} \vdash \phi$ iff $\mathcal{S5} \models \phi$.*

We conclude this section with yet another observation on the relation between knowledge and belief. Although $\mathcal{S5} \models^*$ satisfies the properties of belief, $\mathcal{KD45} \models^*$ does **not** satisfy the properties of belief. For a simple counterexample, consider the single-agent two-state $\mathcal{KD45}$ model with $R_a = \{(s, t), (t, t)\}$

and where p is only true in s , visualized as $p \xrightarrow{a} \bar{p}$. Then $s, \epsilon \models B_a \neg p$. After the truthful announcement of p and the reception of it by a , the beliefs of agent a are inconsistent (meaning that the agent can no longer access other epistemic states and thus believes \perp), and therefore $s, p.a \not\models B_a p \rightarrow \hat{B}_a p$. This is a well-known problem of $\mathcal{KD}45$ updates in $\mathcal{KD}45$ models [16].

6.4.2 Knowledge or belief?

We recall the definition of the view relation as

$$\alpha \triangleright_a \beta \text{ iff } \alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a} = \beta \upharpoonright_{!}.$$

The restriction $\beta \upharpoonright_{!a} = \beta \upharpoonright_{!}$ rules out that the agent considers other agents having received more announcements than herself. If we remove that constraint, we get

$$\alpha \equiv_a \beta \text{ iff } \alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a}.$$

The relation \equiv_a is an equivalence relation.

The interpretation of B_a is defined as $s, \alpha \models B_a \phi$ iff $t, \beta \models \phi$ for all t, β such that $R_a s t$, $\alpha \triangleright_a \beta$, $t \bowtie \beta$. If R_a is an equivalence relation, and if we replace $\alpha \triangleright_a \beta$ by $\alpha \equiv_a \beta$ in the above definition, then the agreement relation \bowtie is no longer well-founded: note for example that it is always the case that $\alpha \equiv_a \alpha.B_a \phi$. Consequently, in order to determine whether $s, \alpha \models B_a \phi$, given that $R_a s s$ and $\alpha \equiv_a \alpha.B_a \phi$, we have to determine whether $s \bowtie \alpha.B_a \phi$, for which we have to determine whether $s, \alpha \models B_a \phi$: a vicious circle. Or at least vicious on first sight, without alternative modelling solutions such as fixpoints.

We may need a novel way to give a semantics to the epistemic modality. However, any such modality will clearly be interpreted by an *equivalence* relation. Instead of B_a having the properties of belief, it would then have the properties of *knowledge*; and one might as well write K_a for it, as we will do from here on. In the temporal epistemic logics for interpreted systems the epistemic modality is indeed such a knowledge modality, and the view relation in such works always is an equivalence relation [58, 88, 78]. This is also the approach followed in [67].

Given the history of asynchronous knowledge in distributed computing, one would preferably have such a notion of knowledge also in a dynamic epistemic logic. This, AA cannot offer at this stage. Clearly, the generalization of the present framework to these epistemic notions constitutes obligatory further research. This is not to say that the belief semantics is somehow a second choice. Both the knowledge and belief semantics have their advantages, and

ideally one would have a logic wherein both epistemic and doxastic operators appear, and which can be tailored to the needs of the modeller. In the remainder of this subsection, let us more precisely focus on the differences between asynchronous knowledge and asynchronous belief, and on possible modelling advantages of asynchronous belief.

Knowledge of novel propositions. In dynamic epistemic logics, the messages sent do not contain novel relevant propositions but are *updates* on the uncertainty about the currently relevant propositions, that are a given and that have a fixed unchangeable value. The goal of such sequences of updates is to finally determine their value, and the interesting phenomena are those wherein some agents reveal their uncertainty about the beliefs of other agents and thus acquire hard information about such facts.

If facts also change value, for example if messages sent and received are recorded by making fresh variables (atoms) true, even knowledge of atoms can change and $K_a p$ may be true now but $K_a \neg p$ may be true later. This is the common scenario in distributed computing.

Belief in positive formulas is correct. As shown in Proposition 6.46, beliefs in positive formulas are stable. And, as shown in Proposition 6.60, for the class $\mathcal{S}5$ of initial models, such beliefs are correct, and thus knowledge. As explained there, this can be interpreted as all belief eventually becoming knowledge.

Decidability. The knowledge semantics reasons over all possible future updates of the current model, and therefore over all possible model restrictions. In other words, it quantifies over all announcements. *Arbitrary Public Announcement Logic* (APAL) is a logic with a modality for quantifying over announcements and it is known to be undecidable [47]. The asynchronous variant of APAL, which will be introduced in Chapter 8, is also undecidable. One might therefore conjecture a logic of announcements with asynchronous knowledge to be also undecidable, although this is not known. However, the logic AA with the belief semantics is decidable (Corollary 6.55).

Should knowledge of ignorance be unsatisfiable? We now continue to explore somewhat informally the above knowledge semantics with \equiv_a . In this semantics, it seems that an agent can never know that another agent is ignorant.

It is not hard to see why $K_a \neg (K_b p \vee K_b \neg p)$ is unsatisfiable for an atom p : given a state s and a history α that is executable in s , atom p is necessarily either true or false in s and, consequently, either p or $\neg p$ can be announced. In the first case, consider the history $\alpha.p.b^n$ (where we append enough b 's at the end so that b has read the last announcement and a has not). It holds that $s, \alpha.p.b^n \not\models \hat{K}_b \neg p$. In the second case, similarly, b no longer considers any state possible wherein p is true: $s, \alpha.\neg p.b^n \not\models \hat{K}_b p$. It follows that $K_a \neg (K_b p \vee K_b \neg p)$ is unsatisfiable.

Similarly, this argument holds for any Boolean formula instead of an atom: for any Boolean formula ϕ , $K_a \neg (K_b \phi \vee K_b \neg \phi)$ is unsatisfiable.

We conjecture that it is also impossible to know that other agents are ignorant for arbitrary formulas ϕ , but this is even harder to make precise given that we have only informally considered the knowledge semantics. Given a finite model, an *a priori* argument is that we can always announce the characteristic formula of the current state and have this announcement be received by agent b , after which any formula ϕ is either true or false, and, as we conjecture, even known by or knowable to b . From $K_b \phi \vee K_b \neg \phi$ we then obtain $\hat{K}_a (K_b \phi \vee K_b \neg \phi)$, negating the above.

Now consider the belief semantics. Here, it is obvious that formulas of shape $B_a \neg (B_b \phi \vee B_b \neg \phi)$ are satisfiable — and such beliefs may even be correct. For a very basic example, consider an initial model consisting of a p -state s and a $\neg p$ -state t that are indistinguishable for two agents a, b . Obviously, $s, \epsilon \models B_a \neg (B_b p \vee B_b \neg p)$. Of course, also $s, p.b \models B_a \neg (B_b p \vee B_b \neg p)$ even though $s, p.b \models B_b p$. This is trivial. Let us proceed with the non-trivial: belief as acknowledgement.

Belief as acknowledgement Continuing the analysis of this basic example, it is however non-trivial that a may signal to b that she has not yet received novel information, by announcing $B_a \neg (B_b p \vee B_b \neg p)$ — here, we use a truthful public announcement of $B_a \phi$ to represent a truthful announcement by a of ϕ . Then, e.g., history $p.b.(B_a \neg (B_b p \vee B_b \neg p)).b$ reveals to b that he received the first announcement p before a . Such announcements $B_a \neg (B_b p \vee B_b \neg p)$ are more *acknowledgements* by a than *beliefs* of a .

This allows, for instance, for an asynchronous analysis of the well-known *Muddy Children* puzzle [76, 41] which uses the semantics for asynchronous belief, wherein agents gain factual knowledge by *acknowledging* the ignorance of others. This analysis would not be possible using a knowledge semantics where one cannot know about the ignorance of others.

This asynchronous solution to the puzzle, which is presented in detail in

[15], runs into a problem: the fact that there is an implicit public announcement in this scenario, namely in the fact that no child steps forward when the father calls upon the muddy children, is not really usable in a setting in which each child does not know whether the others received Father’s message. Instead, the announcement by each child to her siblings is formalised via the acknowledgement of having received the previous message appended to the information she has gained from this announcement.

6.5 Comparison with Asynchronous Broadcast Logic

Asynchronous Broadcast Logic is developed in [67] and the aforementioned related works as an asynchronous framework for sending and receiving announcements. The epistemic notion is one of knowledge and not one of belief, unlike the one presented here.

The logical language is the same employed here, except with a knowledge modality K_a instead of our B_a and a different notation for the diamond form of the reception modality: \bigcirc_a instead of $\langle a \rangle$. We shall stick to our notation $\langle a \rangle$ in this section.

This language is interpreted on structures called *asynchronous pre-models*, with domain elements that are triples (s, σ, c) where:

- s is a state from some initial Kripke model;
- σ is a sequence of formulas taken from a *protocol*, i.e., a set of allowed sequences of formulas, and
- $c : A \rightarrow \{0, \dots, |\sigma|\}$ is a *cut*, i.e. a specification of how many announcements in the sequence σ each agent has already received.

Just as in the framework presented in this text, announcements must be true when made and are individually received by the agents in the order in which they were sent. Unlike this text, the epistemic modality K_a is interpreted over histories of arbitrary length, thus guaranteeing that knowledge is correct: $K_a\phi \rightarrow \phi$ is always valid. When interpreting knowledge on pre-models, not all triples (s, σ, c) of the pre-model are taken into account but only those where all formulas of σ could have been truthfully announced. They call this the requirement of *consistency* of σ with s (this is comparable to our *agreement* relation \bowtie). Their semantics is then based on mutual recursion of ‘truth’ and ‘consistency’, similar to ours.

As their epistemic notion is one of knowledge, also taking into account announcements that have not yet been received, they face the already mentioned issue of circularity, to which they provide two well-founded solutions. Their first solution is to restrict the structures to (initial) models that are finite point-generated trees, and where the model transformations are relative to the root of the model; this solution is reminiscent of [74]. Their second solution is to restrict the language to the so-called *existential fragment*, in which negations are only allowed of atoms and modalities are only allowed in their ‘diamond’ form $\hat{K}_a\phi$, $\langle a \rangle\phi$, and $\langle \psi \rangle\phi$. They present some validities for their logical semantics, such as $\langle a \rangle\langle b \rangle\phi \leftrightarrow \langle b \rangle\langle a \rangle\phi$: without intervening announcements, the order of reception does not matter. They do not provide an axiomatization.

One observation pertaining the similarities between the framework of AA and ABL has to do with the fact that a pair (σ, c) of a sequence and a cut is related to our notion of *history*, albeit it provides less information. To illustrate this, we shall take a small detour through an alternative semantics for Asynchronous Announcement Logic based on equivalence classes.

6.5.1 Equivalence-based AA

Definition 6.62. Let \equiv be the equivalence relation defined on the set of histories by:

$$\alpha \equiv \beta \text{ iff } \alpha \upharpoonright_! = \beta \upharpoonright_! \ \& \ |\alpha|_a = |\beta|_a \ \forall a \in A.$$

We shall use capital Greek letters to refer to the equivalence classes of this relation, $\alpha \in A, \beta \in B, \gamma \in \Gamma, \dots$. Specifically, we will use E to refer to the class containing only the empty history, $E = \{\epsilon\}$.

Given a class A , and any representative $\alpha \in A$, let us define:

- $A.a$ as the equivalence class of $\alpha.a$;
- $A.\phi$ as the equivalence class of $\alpha.\phi$;
- $A \upharpoonright_!, A \upharpoonright_{!a}, |A| \upharpoonright_!, |A|_a$ and $|A|$ respectively as $\alpha \upharpoonright_!, \alpha \upharpoonright_{!a}, |\alpha| \upharpoonright_!, |\alpha|_a$ and $|\alpha|$. (The last notation is quite nonstandard, but we will never care about the actual cardinality of A).

Note that this is well-defined and does not depend on the choice of $\alpha \in A$. Note as well that, if $\alpha, \alpha' \in A$ and $\beta, \beta' \in B$, it holds that $\alpha \triangleright_a \beta$ iff $\alpha' \triangleright_a \beta'$. We can thus define a relation $A \triangleright_a B$ in the obvious way.

Later in this text we will define a semantics for AA with respect to a pair of a world w and an equivalence class A . It is of note that we can reach the

same class of histories through different paths. For instance $E.p.a.\neg B_a p = E.p.\neg B_a p.a$. The fact that $p.a.\neg B_a p$ is not a compatible history will be taken care of by the semantics below, so that we can have for instance that $w, E \models \langle p.\neg B_a p.a \rangle \top$, but never $w, E \models \langle p.a.\neg B_a p \rangle \top$.

Given a finite sequence of formulas σ , let H_σ be the set of histories α such that $\alpha|_1 = \sigma$. The following is rather obvious:

Lemma 6.63. *Given a finite sequence of formulas σ , there is a 1-to-1 correspondence between the equivalence classes of \equiv in H_σ and the maps $c : A \rightarrow \{0, \dots, |\sigma|\}$, given by*

$$A \mapsto c_A,$$

where $c_A(a) = |A|_a$.

As a consequence, every equivalence class A uniquely determines a pair (σ, c) , namely $(A|_1, c_A)$. The following fact is perhaps not so obvious:

Lemma 6.64. *Let $\alpha \equiv \alpha'$, let (W, R, V) be a model and $s \in W$. If $s \bowtie \alpha$ and $s \bowtie \alpha'$, then for any formula $\phi \in \mathcal{L}_{AA}$,*

$$s, \alpha \models \phi \text{ iff } s, \alpha' \models \phi.$$

Proof. The proof is by induction on ϕ . The cases $\phi = p, \perp, \psi_1 \vee \psi_2, \neg\psi$ are left to the reader.

Case belief: $s, \alpha \models B_a \phi$, iff $t, \gamma \models \phi$ for all (t, γ) such that $R_a s t$, $\alpha \triangleright_a \gamma$, and $t \bowtie \gamma$, iff (by the above observation) $t, \gamma \models \phi$ for all (t, γ) such that $R_a s t$, $\alpha' \triangleright_a \gamma$, and $t \bowtie \gamma$, iff $s, \alpha' \models B_a \phi$.

Case reception: $s, \alpha \models [a]\psi$, iff $|\alpha|_a < |\alpha|_1$ implies $s, \alpha.a \models \psi$. Noting that $\alpha.a \equiv \alpha'.a$ we have by induction hypothesis that $s, \alpha'.a \models \psi$ and therefore $s, \alpha' \models [a]\psi$.

Case announcement: $s, \alpha \models [\phi]\psi$, iff $s, \alpha \models \phi$ implies $s, \alpha.\phi \models \psi$, iff (twice induction plus the fact that $\alpha.\phi \equiv \alpha'.\phi$) $s, \alpha' \models \phi$ implies $s, \alpha'.\phi \models \psi$, iff $s, \alpha' \models [\phi]\psi$. ■

We are now in a position to define a semantics for \mathcal{L}_{AA} in terms of these equivalence classes.

Definition 6.65 (Semantics of Eq-AA). Let w be a world and A be an equivalence class of the \equiv relation. We define by double \ll -recursion³ an agreement relation \bowtie and satisfaction relation \models as follows:

³We omit the details of this but we point out that one can determine a well-founded order \ll between pairs (w, A) of worlds and equivalence classes in a way almost identical to the one in page 88, making the semantics well-defined.

$w \bowtie E$ always;
 $w \bowtie A$ iff
 either $A = A'.a$ and $w \bowtie A'$,
 or $A = A'.\psi$, A' represents a class of histories^(*),
 $w \bowtie A'$ and $w, A' \models \psi$.

$w, A \models p$ iff $w \in V(p)$;
 $w, A \models \neg\phi$ iff $w, A \not\models \phi$;
 $w, A \models \phi \vee \psi$ iff $w, A \models \phi$ or $w, A \models \psi$;
 $w, A \models [\psi]\phi$ iff $w, A \models \psi$ implies $w, A.\psi \models \phi$;
 $w, A \models [a]\phi$ iff $|A|_a < |A|_!$ implies $w, A.a \models \phi$;
 $w, A \models B_a\phi$ iff $t, B \models \phi$ for all t, B with $R_a w t$, $A \triangleright_a B$, and $t \bowtie B$.

(*) Note that this need not be the case: if A is the equivalence class of $p.q.a.a$, then $A = [p.a.a]_{\equiv}.g$, but of course $p.a.a$ is not a history.

We say that a formula is *E-valid* if $w, E \models \phi$ for every world w in every model, where $E = \{\epsilon\}$; let us call this framework *equivalence-based AA* (Eq-AA) and its logic of *E-validities* Eq-AA^E.

Now, we have:

Proposition 6.66. 1. $w \bowtie A$ (in the above sense) iff there exists $\alpha \in A$ such that $w \bowtie \alpha$ (in the AA sense);

2. Assuming $w \bowtie A$, we have: $w, A \models \phi$ iff for all $\alpha \in A$ such that $w \bowtie \alpha$, it holds that $w, \alpha \models \phi$ in the AA sense. (Equivalently, by Lemma 6.64, iff for some $\alpha \in A$ it holds that $w \bowtie \alpha$ and $w, \alpha \models \phi$.)

Proof. 1. By induction on $|A|$. Clearly, $w \bowtie E$ iff $w \bowtie \epsilon$. Now suppose $\alpha \neq \epsilon$ and $w \bowtie \alpha \in A$. Since $\alpha = \alpha'.x$ for some $x \in \mathcal{L} \cup \text{Agt}$, we have that $w \bowtie \alpha'$ and thus (by IH) that $w \bowtie A'$, this being the equivalence class of α' . But then $w \bowtie A'.x = A$. (If x is a formula ψ , we have to use the induction hypothesis again to claim that $w, A' \models \psi$.) Conversely, suppose $w \bowtie A$ and $A \neq E$. By the previous definition, there exists some equivalence class A' and some $x \in \mathcal{L} \cup \text{Agt}$ such that $A = A'.x$ and $w \bowtie A'$. Then by induction hypothesis $w \bowtie \alpha'$ for some $\alpha' \in A'$ and thus $w \bowtie \alpha'.x \in A'.x = A$. (In the case where $x = \psi$, we have to use the induction hypothesis twice to also show that $w, \alpha' \models \psi$.)

2. By induction on ϕ : Obvious if $\phi = p, \psi_1 \vee \psi_2$.

For the case $\phi = \neg\psi$, we have: $w, A \models \neg\psi$ iff $w, A \not\models \psi$ iff (by IH) for all $\alpha \in A$ such that $w \bowtie \alpha$, we have $w, \alpha \not\models \psi$, iff for all $\alpha \in A$ such that $w \bowtie \alpha$ we have $w, \alpha \models \neg\psi$.

For the case $\phi = [\theta]\psi$, suppose $w, A \models [\theta]\psi$: $w, A \models \theta$ imply $w, A.\theta \models \psi$. Now, if $w, A \models \theta$, by induction hypothesis there is $\alpha \in A$ with $w, \alpha \models \theta$. Note that $\alpha.\theta \in A.\theta$, and thus by induction hypothesis again $w, \alpha.\theta \models \psi$. Conversely, suppose $w, \alpha \models [\theta]\psi$: then $w, \alpha \models \theta$ implies $w, \alpha.\theta \models \psi$. By induction hypothesis, this gives $w, A \models \theta$ implies $w, A.\theta \models \psi$.

The case $\phi = [a]\psi$ is analogous to the previous case.

For the case $\phi = B_a\psi$, we have: $w, \alpha \models B_a\phi$ iff for all some t, β such that (1) $R_a w t$, (2) $\alpha \triangleright_a \beta$, and (3) $t \bowtie \beta$ it holds that (4) $t, \beta \models \psi$. But (2) is equivalent to $A \triangleright_a B$ by the definition of \triangleright_a on equivalence classes, and (3)+(4) are equivalent to $t \bowtie B$ and $t, B \models \psi$ by induction hypothesis. Therefore, this is equivalent to $w, A \models B_a\psi$. ■

Corollary 6.67. *The E-validities of Eq-AA are exactly the ϵ -validities of AA.*

Proof. Follows from the above Proposition and Lemma 6.64. ■

6.5.2 A doxastic variant of ABL

We showcase the semantics for Asynchronous Broadcast Logic presented in [67]. We recall that they interpret formulas with respect to triples (w, σ, c) where σ is a finite sequence of formulas (from a set of such acceptable sequences) and $c : A \rightarrow \{0, \dots, |\sigma|\}$ is a *cut*; we recall likewise that these pairs (σ, c) correspond exactly to our equivalence classes above (Lemma 6.63).

We introduce some notation on cuts: given a cut c ,

- the cut c^{a+1} is defined as: $c^{a+1}(a) = c(a) + 1$ and $c^{a+1}(b) = c(b)$ for $b \neq a$;
- $c' \leq c$ iff $c'(a) \leq c(a)$ for all $a \in A$ and $c' < c$ iff $c' \leq c$ and $c' \neq c$.

Let (σ, c) be a (sequence, cut) pair and let A be the equivalence class corresponding to it (i.e. $\sigma = A|_1$, and $c(a) = |A|_a$ for all a).

In [67], one has the following consistency conditions:

- $(w, \epsilon, \mathbf{0}) \models \checkmark$ always (where $\mathbf{0}$ is the zero function).
(Note that this corresponds to our condition: $w \bowtie E$ always).
- $(w, \sigma, c) \models \checkmark$ iff
either there is $c' < c$ s.t. $(w, \sigma, c') \models \checkmark$
(note that this corresponds to: “either $A = A'.a$ with $w \bowtie A'$ ”
or $\sigma = \sigma'.\psi$, $c \leq |\sigma'|$, $(w, \sigma', c) \models \checkmark$ and $(w, \sigma', c) \models \psi$)

(note this corresponds exactly to: “or $A = A'.\psi$, A' represents a class of histories, $w \bowtie A'$ and $w, A' \models \psi$ ”).

And the following semantics for announcements and readings:

- $(w, \sigma, c) \models \langle \psi \rangle \phi$ iff $\sigma.\psi$ is announceable (note that we do not have any restrictions on announceability), $(w, \sigma, c) \models \psi$ and $(w, \sigma.\psi, c) \models \phi$.
(This corresponds, minus announceability restrictions, to: $w, A \models \psi$ and $w, A.\psi \models \phi$.)
- $(w, \sigma, c) \models \langle a \rangle \phi$ iff $c(a) < |\sigma|$ and $(w, \sigma, c^{a+1}) \models \phi$.
(This corresponds to: $|A|! > |A|_a$ and $w, A.a \models \phi$.)

Now, their semantics in [67] the knowledge modality are as follows:

- $(w, \sigma, c) \models K_a \phi$ iff $(w', \sigma', c') \models \phi$ for all (w', σ', c') such that $R_a w w'$ and $\sigma|_{c(a)} = \sigma'|_{c'(a)}$.

This runs into the aforementioned circularity issues, which are sidestepped in the text by limiting the conditions for announceability (i.e., the admissible σ 's) and imposing restrictions on the Kripke models (for example, finite and point-generated).

Let us now define instead a doxastic modality on the ABL framework, making its semantics correspond to our Eq-AA doxastic modality above:

- $(w, \sigma, c) \models B_a \phi$ iff $(w', \sigma|_a, c') \models \phi$ for all w', c' such that $R_a w w'$, $c'(a) = c(a)$ and $(w', \sigma|_a, c') \models \checkmark$ (where $\sigma|_a$ is the prefix of σ of length $c(a)$).

Note that we have indeed defined this so that $(w, \sigma, c) \models B_a \phi$ iff $w, A \models B_a \phi$, where A is the corresponding equivalence class.

Definition 6.68. Let us call the framework of ABL minus the knowledge modality and the restrictions on announceable sequences, plus the doxastic modality defined above, *Doxastic Unrestricted Asynchronous Broadcast Logic* (DU-ABL).

A formula ϕ is *zero-valid* in DU-ABL if $(w, \epsilon, \mathbf{0}) \models \phi$ for every world w in every model. Let us call the logic of zero-validities DU-ABL⁰.

From all the above observations we can extract the following conclusion:

Corollary 6.69. *The zero-validities of DU-ABL are exactly the E -validities of Eq-AA, which are in turn the ϵ -validities of AA.*

In other words, $\text{DU-ABL}^0 = \text{Eq-AA}^E = \text{AA}^\epsilon$.

And therefore:

Corollary 6.70. *DU-ABL^0 is decidable.*

6.5.3 Cuts or histories?

Considering a doxastic variant of Asynchronous Broadcast Logic gives us completeness and decidability. This is done by putting it in relation with the logic AA. A good direction for future work would be the inverse: how could one translate the techniques used in [67] to avoid circularity into a ‘knowledge’ variant of AA?

The notion of a (σ, c) pair is weaker than that of a history, in that it does not specify the order in which each agent reads the announcements. The above analysis highlights, arguably, some of the advantages of considering a history-based semantics of AA over a cut-based one like ABL. We have arrived at an axiomatisation of DU-ABL by using the results obtained for AA in the previous sections of this chapter. It is not clear that these results could be as simple if one started off with a cut-based semantics. To give an example, our Lemma 6.32, which states:

For every formula ϕ and for every world s such that $s \bowtie \alpha$, $s, \alpha \models \langle \beta \rangle \chi$ if and only if (i) the concatenation $\alpha.\beta$ is a history, (ii) $s \bowtie \alpha.\beta$, and (iii) $s, \alpha.\beta \models \phi$.

does not seem to have an easy translation to the DU-ABL cut-based semantics.

We note a major difference between the notion of consistency “ \checkmark ” in DU-ABL (which corresponds exactly to the notion “ \bowtie ” in Eq-AA) and the notion of executability in AA:

Lemma 6.71. *Let $\alpha \in A$. Then $w \bowtie \alpha$ in the AA sense iff, for every prefix $\beta \sqsubseteq \alpha$, $w \bowtie B$ in the Eq-AA sense, where B is the equivalence class of β .*

We leave the proof of this result to the reader. As a consequence of this, a very simple result like

$w, \epsilon \models \langle \alpha \rangle \phi$ iff $w \bowtie \alpha$ and $w, \alpha \models \phi$

would look quite more involved in the DU-ABL setting:

$w, \epsilon, \mathbf{0} \models \langle \alpha \rangle \phi$ iff $w, \alpha \upharpoonright_1, c_\alpha \models \phi$ and $w, \beta \upharpoonright_1, c_\beta \models \checkmark$ for all $\beta \sqsubseteq \alpha$
 (where we define $c_\gamma(a) := |\gamma|_a$ for all a).

Moreover one cannot easily escape the notion of histories in the cut-based framework: for any satisfiable formula of the shape $\langle \alpha \rangle \phi$, α must be a history in the AA sense. The reduction axiom for the belief modality in DU-ABL (which is the same as AA) has the shape

$$[\alpha]B_a\phi \leftrightarrow [\alpha]\perp \vee \bigwedge_{\alpha \triangleright_a \beta} B_a[\beta]\phi,$$

which requires, on account of the big conjunction, a definition of a “ \triangleright_a ” relation on the set of histories, as well as one on the set of pairs (σ, c) to define the semantics.

Arguably, having a semantics for asynchronous announcements which puts histories front and center has a few formal advantages over one which uses cuts.

Discussion. This chapter has presented asynchronous announcement logic AA, a logic of epistemic change due to announcements, with separate modalities for sending and for receiving such messages where the epistemic modality is one of belief, not one of knowledge. The axiomatisation for this logic is a reduction system: every formula is equivalent to a formula without announcement and reception modalities. The logic AA is therefore also decidable. We determined results for special formulas and for special model classes: the positive formulas are preserved after update, and on the model class $\mathcal{S5}$, belief of positive formulas is correct and thus knowledge.

The complexity of model checking and of satisfiability of AA is left for further research. This work admits numerous generalisations, such as for subgroups synchronously receiving announcements, for non-public actions, and instead of belief for knowledge, wherein one also reasons about the future. Some of these are explored in further chapters, whereas some would constitute very interesting directions for future research.



Partially synchronised Announcement Logic

In this chapter we present a variation of Asynchronous Announcement Logic in which the reception of announcements is done by a group of agents as opposed to an individual. Both PAL and Asynchronous Announcement Logic are special cases in this framework. A history-based semantics for our ‘Partially Synchronised Announcement Logic’, is offered proposing three different interpretations of the notion of asynchronicity in this setting. We 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



What if phones, but too much?

Mallory Ortberg, in a review of the TV show Black Mirror published in The Toast in 2016.

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 .

We will model this, not unlike the previous chapter, with the use of *histories*.

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

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

¹This chapter is based on the paper

Philippe Balbiani, Hans Van Ditmarsch, and Saúl Fernández González. From Public Announcements to Asynchronous Announcements. In *ECAI 2020: 24th European Conference on Artificial Intelligence*, 2020.

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 puts us in the situation of Asynchronous Announcement Logic, discussed in Chapter 6. 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. This chapter provides several intuitive solutions to address that. Similar problems have of course been widely discussed in the temporal epistemic (and distributed computing) literature [32, 78, 58, 88, 82].

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.

What all three scenarios have in common is indeed a notion of *public* announcement, in the sense of a 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. Like in the previous chapter, 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*.

But unlike the previous chapter, this one 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 [86] and AA. 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 [81]: 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 [76], 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 [17, 42].

The remainder of this chapter is structured as follows: in Sections 7.1, 7.2 and 7.3 we present the technicalities of Partially Synchronised Announcement Logic (PSAL). In Sections 7.4 and 7.5 we show how we can emulate AA and PAL, respectively, within this framework. In Section 7.6 we give a proposal for a notion of Common Belief in this framework, of which we provide some validities.

7.1 Syntax and histories

Throughout this chapter, let \mathbf{Prop} 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 $\mathcal{L}_{\text{PSAL}}$, defined as:

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

where $p \in \mathbf{Prop}$ 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 $[G]\phi = \neg\langle G\rangle\neg\phi$.²

We will consider words of the shape of β in the introductory example: finite sequences made out of formulas and subsets of A , i.e. words $\alpha \in (\mathcal{L}_{\text{PSAL}} \cup \mathcal{G})^*$. For cleanliness in presentation, when writing down these histories explicitly announcements and readings will be separated with dots and the agents in a group G will be written as a concatenation rather than as a set (as β in the introduction).

Much like in the previous chapter, 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 , i.e.:

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

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

Let us order the elements appearing in α by \sqsubset_α and let us use $\alpha|_A$ and $\alpha|_!$ to denote the projection of α to, respectively, \mathcal{G} and $\mathcal{L}_{\text{PSAL}}$. We use $|\alpha|_!$ to denote the length of $\alpha|_!$, 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

²It would perhaps be more natural to define the language as \mathcal{L}_{AA} was defined in Chapter 6, by making B_a and $[\cdot]$ the primitive modalities and \hat{B}_a , $\langle\cdot\rangle$ the defined duals; I find that using the language as described here, however, simplifies the proofs by induction by some margin.

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, $\alpha|_a$, will be the first $|\alpha|_a$ announcements in $\alpha|_!$. That is, if $\alpha|_! = \phi_1 \dots \phi_n$, we have

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

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

In order for a word to be a *history*, an agent should not be able, at any point along the history, to read more announcements than those which have been made; in addition, 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 previously. That is, a word $\alpha \in (\mathcal{L}_{\text{PSAL}} \cup \mathcal{G})^*$ 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,

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

Example 7.1. Let

$$\alpha_1 = p. \neg B_a p. abc. B_b B_a p. bc. a.$$

The formulas p and $\neg B_a p$ are announced in succession, after which the agents read the first announcement p together. After this occurs, the formula $B_b B_a p$ is now true, and it is announced. Afterwards, b and c read the announcement $\neg B_a p$ 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 these announcements were, but also they communicate to 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 $\alpha \upharpoonright_A$, which we can define as the reflexive and transitive closure of

$$\bigcup_{a \in A} \xrightarrow{a},$$

where $G \xrightarrow{a} G'$ is the relation representing “ G' is the next group, after G , in which a appears”: formally, $G \xrightarrow{a} G'$ 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 group reading of α including a_i .

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

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

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,

$$\alpha \upharpoonright_a := \downarrow_{\leq_\alpha^*} \text{last}_a \alpha = \langle \{G \in \alpha \upharpoonright_A : G \leq_\alpha^* \text{last}_a \alpha\}, \leq_\alpha^* \rangle.$$

Note that, for any group G , the set $\downarrow_{\leq_\alpha^*} G$, defined as the tuple

$$\langle \{G' \in \alpha \upharpoonright_A : G' \leq_\alpha^* G\}, \leq_\alpha^* \rangle,$$

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 $\downarrow_{\leq^*} G$, or rather the *height* of G in the poset. For $G \in \alpha \upharpoonright_A$,

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

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 $\alpha \upharpoonright_1$, i.e. $\alpha \upharpoonright_{!a} = \phi_1 \dots \phi_{|\alpha|_a}$.

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 $|\beta|_a \leq |\alpha|_a$.

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,

$$\alpha \triangleright_a \beta \text{ iff } \begin{cases} \alpha \upharpoonright_a = \beta \upharpoonright_a \text{ and} \\ \alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a} = \beta \upharpoonright_1. \end{cases}$$

Example 7.2. Let

$$\alpha_2 = p. \neg B_a p. ab. B_b B_a p. bc. B_c p. ca.$$

Here, p and $\neg B_a p$ are announced, after which a and b together read p . Then, $B_b B_a p$ is announced and b and c read together the second announcement $\neg B_a p$. Now c is also aware of p because b has communicated to c the first announcement. After this, $B_c p$ 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, $B_b B_a p$, 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 Carmiña, who had previously watched the first one by herself.

For this we define \leq_α^* , $\text{last}_a \alpha$ and $\text{h}(G)$ as above. Under this interpretation, a only reads as many announcements as times it occurs in α . We then define $\alpha \upharpoonright_a$ 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 $\alpha \upharpoonright_a$. 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 $\alpha \upharpoonright_{!a} = \phi_{n_1^a} \dots \phi_{n_{|\alpha|_a}^a}$, where $n_k^a = \text{h}(\alpha_k^a)$.

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 $\beta \sqsubseteq \alpha$, $\text{h}(\text{last}_a \beta) \leq |\beta|_a$.

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 7.3. 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 7.2 for formal details on executability).

While each of these interpretations gives rise to different notions of executability (\boxtimes) and indistinguishability (\triangleright_a), and thus to different semantics (see Section 7.2), 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 7.3 for details).

Unless stated otherwise, the results in the remainder of this chapter 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.

7.2 Semantics

We read formulas of $\mathcal{L}_{\text{PSAL}}$ 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 $w \bowtie \alpha$. Similar to the previous chapter, 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 7.4. We define $\text{deg } \phi$ and $\|\phi\|$ recursively:

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

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

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

Finally, for pairs (α, ϕ) we set:

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

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

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

Definition 7.5 (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 7.1.

$w \bowtie \epsilon$	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 7.1: Semantics of PSAL

Let us show a last example before moving on:

Example 7.6. Let us see why α_3 from Example 7.3 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 $w \not\models \alpha_3$. Indeed, suppose $\beta \triangleright_c \gamma$. We have that $\beta \upharpoonright_c = \downarrow_{\leq^*} \text{last}_c \beta = \langle \{ab, bc\}, \leq^* \rangle$, 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 $w, \beta \models B_c p$.

7.3 The logic PSAL

We maintain the notions of validity from the preceding chapter: 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 we will call it $*$ -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 7.7. *Given a model (W, R, V) and $w \in W$, for any formula $\phi \in \mathcal{L}_{\text{EL}}$ 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, $\phi \in \mathcal{L}_{\text{EL}}$ 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 \mathbf{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 are the counterparts of Lemmas 6.31 and 6.32 in the previous chapter and can both be easily proven by induction on the length of α .

Lemma 7.8. *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 7.9. *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 7.2.

- i. All the axioms and rules of the minimal modal logic \mathbf{K} for each of the B_a modalities;
- ii. the following reduction axioms (where $\alpha \in (\mathcal{L}_{\text{PSAL}} \cup \mathcal{G})^*$):
 - $(R_{\top 1})$ $\langle \alpha.G \rangle \top \leftrightarrow \langle \alpha \rangle \top$ if $\alpha.G$ is a history;
 - $(R_{\top 2})$ $\langle \alpha.G \rangle \top \leftrightarrow \perp$ otherwise;
 - $(R_{\top 3})$ $\langle \alpha.\phi \rangle \top \leftrightarrow \langle \alpha \rangle \phi$;
 - (R_p) $\langle \alpha \rangle p \leftrightarrow (\langle \alpha \rangle \top \wedge p)$;
 - (R_{\neg}) $\langle \alpha \rangle \neg \phi \leftrightarrow (\langle \alpha \rangle \top \wedge \neg \langle \alpha \rangle \phi)$;
 - (R_{\vee}) $\langle \alpha \rangle (\phi \vee \psi) \leftrightarrow (\langle \alpha \rangle \phi \vee \langle \alpha \rangle \psi)$;
 - (R_B) $\langle \alpha \rangle \hat{B}_a \phi \leftrightarrow (\langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \phi)$;
- iii. the Modus Ponens rule.

Table 7.2: The logic PSAL

As remarked at the end of Section 7.1, 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. To justify this, let us assume one has fixed one's favourite interpretation.

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

Proposition 7.10. (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 7.9 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 7.9), 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 7.9 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 7.9) $w, \epsilon \models \langle \alpha \rangle \hat{B}_a \phi$. ■

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 of Thm. 6.53. Let us briefly sketch this here:

- i. Similar to Corollary 6.49, we show by \ll -induction that, for every formula $\phi \in \mathcal{L}_{\text{PSAL}}$ there exists an announcement-free formula ψ such that $\phi \leftrightarrow \psi$ is provable in PSAL. In particular, $\psi = s_{\mathbf{A}} \phi$, where the translation $s_{\mathbf{A}}$ is defined by \ll -induction in Table 7.3.
- ii. Since $\phi \leftrightarrow s_{\mathbf{A}} \phi$ is ϵ -valid and $s_{\mathbf{A}} \phi \in \mathcal{L}_{\text{EL}}$, it follows that $w \models s_{\mathbf{A}} \phi$ with the usual Kripke semantics if and only if $w, \epsilon \models \phi$. Thus ϕ is a theorem of PSAL if and only if $s_{\mathbf{A}} \phi$ is a theorem of \mathbf{K} . Completeness follows.

Therefore, we have:

Theorem 7.11. PSAL is a sound and complete axiomatisation of the logic of ϵ -validities for Partially Synchronised Announcements.

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

Table 7.3: The map $s_A : \mathcal{L}_{\text{PSAL}} \rightarrow \mathcal{L}_{\text{EL}}$, defined by \llcorner -recursion on (α, ϕ) . In the last four rows of the table we assume that $\alpha \neq \epsilon$.

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.

7.4 PSAL generalises AA

Recall from the previous chapter that the language of Asynchronous Announcement Logic is \mathcal{L}_{AA} , defined as follows:

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

with $p \in \text{Prop}$ and $a \in A$.

Recall as well that the logic AA is very similar to the logic PSAL as it appears in Table 7.2, with the following changes in the reduction axioms:

$$\begin{aligned} (R_{\top 1(AA)}) \quad & \langle\alpha.a\rangle\top \leftrightarrow \langle\alpha\rangle\top \text{ if } |\alpha|_a < |\alpha|_!; \\ (R_{\top 2(AA)}) \quad & \langle\alpha.a\rangle\top \leftrightarrow \perp \text{ otherwise;} \\ (R_{B(AA)}) \quad & \langle\alpha\rangle\hat{B}_a\phi \leftrightarrow ((\alpha)\top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a\langle\beta\rangle\phi). \end{aligned}$$

PSAL generalises Asynchronous Announcement Logic in a very straightforward way: we can say that a formula in the language of AA 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 AA.

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


7.5 PSAL generalises PAL

Let us emulate³ 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 7.12. 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 $Rts, p \triangleright_a \epsilon$ 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

³This word, *emulate*, perhaps captures what we are trying to do in this section better than the word *generalise*, which is used rather liberally in the title: obviously, since both PSAL and PAL can be reduced to regular epistemic logic, they are both equally expressive and thus one cannot, formally speaking, talk of one as a ‘generalisation’ of the other. This section shows how one can easily *emulate* a public announcement in the PSAL setting.

the only pair (x, γ) such that Rwx , $\alpha' \triangleright_a \gamma$ and $\gamma \triangleright \alpha'$ 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 7.13. Let $\tau : \mathcal{L}_{\text{PAL}} \rightarrow \mathcal{L}_{\text{PSAL}}$ be defined, by recursion on the length of ϕ , as follows:

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

The following holds for all three interpretations.

Theorem 7.14. *For every model (W, R, V) , for all $w \in W$ and all $\phi \in \mathcal{L}_{\text{PAL}}$, 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 $\mathbf{t} : \mathcal{L}_{\text{PAL}} \rightarrow \mathcal{L}_{\text{PSAL}}$, defined as: $\mathbf{t} \top = \top$, $\mathbf{t} p = p$, $\mathbf{t} (\neg \phi) = \neg \mathbf{t} \phi$, $\mathbf{t} (\phi \wedge \psi) = \mathbf{t} \phi \wedge \mathbf{t} \psi$, $\mathbf{t} (B_a \phi) = B_a \mathbf{t} \phi$, $\mathbf{t} (\langle \phi \rangle \psi) = \langle \mathbf{t} \phi . A \rangle \mathbf{t} \psi$. We have:

Lemma 7.15. *If ϕ is an announcement-free formula in the language of PAL, then it holds that $\phi = \mathbf{t} \phi$ and $w \models_{\text{PAL}} \phi$ iff $w, \epsilon \models_{\text{PSAL}} \phi$.*

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 $w \models_{\text{PAL}} \hat{B}_a \psi$ if and only if $t \models_{\text{PAL}} \psi$ for some t such that $R_a w t$ if and only if (by induction hypothesis) $t, \epsilon \models_{\text{PSAL}} \psi$ for some t with $R_a w t$ if and only if (given that $\epsilon \triangleright_a \beta$ implies $\beta = \epsilon$) $w, \epsilon \models_{\text{PSAL}} \hat{B}_a \psi$. ■

For the following result, we will say that a formula $\phi \in \mathcal{L}_{\text{PAL}}$ 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:

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

 Table 7.4: The map SP defined by \ll -recursion.

Lemma 7.16. *i. The map $\mathsf{s}_A : \mathcal{L}_{\text{PSAL}} \rightarrow \mathcal{L}_{\text{EL}}$ defined in Table 7.3 satisfies: for every formula ϕ and every model (W, R, V) we have that $\mathsf{s}_A \phi$ is announcement-free and $w, \epsilon \models \phi$ iff $w, \epsilon \models \mathsf{s}_A \phi$.*

ii. There exists a translation map $\mathsf{SP} : \mathcal{L}_{\text{PAL}} \rightarrow \mathcal{L}_{\text{EL}}$ such that for every formula ϕ in standard form and every model (W, R_a, V) we have that $\mathsf{SP} \phi$ is announcement-free and $w \models \phi$ iff $w \models \mathsf{SP} \phi$. This map is defined by \ll -recursion⁴ in Table 7.4.

With this:

Lemma 7.17. *For any formula ϕ in the language of PAL in standard form, $\mathsf{SP} \phi = \mathsf{s}_A \mathfrak{t} \phi$.*

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 $\mathsf{s}_A \mathfrak{t} \phi = \mathsf{s}_A \langle \mathfrak{t} \chi . A \rangle \top = \mathsf{s}_A \langle \mathfrak{t} \chi \rangle \top = \mathsf{s}_A \mathfrak{t} \chi = \mathsf{SP} \chi = \mathsf{SP} \langle \chi \rangle \top$, the fourth equality given by the induction hypothesis and the rest directly by Table 7.4.

If $\phi = \langle \chi \rangle p$, then $\mathsf{s}_A \mathfrak{t} \phi = \mathsf{s}_A \langle \mathfrak{t} \chi . A \rangle p = \mathsf{s}_A \langle \mathfrak{t} \chi . A \rangle \top \wedge p$. The first conjunct of this last formula equals $\mathsf{SP} \chi$ by the previous step, therefore $\mathsf{s}_A \mathfrak{t} \langle \chi \rangle p = \mathsf{SP} \chi \wedge p = \mathsf{SP} \langle \chi \rangle p$.

If $\phi = \langle \chi \rangle \neg \psi$ we have that $\mathsf{s}_A \mathfrak{t} \phi = \mathsf{s}_A \langle \mathfrak{t} \chi . A \rangle \top \wedge \neg \mathsf{s}_A \langle \mathfrak{t} r \chi . A \rangle \mathfrak{t} \psi$. The first conjunct is again equal to $\mathsf{SP} \chi$ (as before) and the second equals $\neg \mathsf{s}_A \mathfrak{t} \langle \chi \rangle \psi$ which, by induction hypothesis, equals $\mathsf{SP} \langle \chi \rangle \psi$. Therefore $\mathsf{s}_A \mathfrak{t} \phi = \mathsf{SP} \chi \wedge \neg \mathsf{SP} \langle \chi \rangle \psi = \mathsf{SP} \phi$.

⁴Slight abuse of notation: for formulas in PAL, $\phi \ll \psi$ iff $(\text{deg } \phi < \text{deg } \psi)$ or $(\text{deg } \phi = \text{deg } \psi \text{ and } \|\phi\| < \|\psi\|)$, where deg and $\|\cdot\|$ are as in Def. 9.18.

If $\phi = \langle \chi \rangle \hat{B}_a \psi$, then $\mathbf{s}_A \mathbf{t} \phi = \mathbf{s}_A \langle \mathbf{t} \chi . A \rangle \hat{B}_a \mathbf{t} \psi$. Now note that if $\mathbf{t} \chi . A \triangleright_a \beta$, then $\beta = \mathbf{t} \chi . A$ and therefore $\mathbf{s}_A \mathbf{t} \phi = \mathbf{s}_A \langle \mathbf{t} \chi . A \rangle \top \wedge \hat{B}_a \mathbf{s}_A \langle \mathbf{t} \chi . A \rangle \mathbf{t} \psi$. By induction hypothesis the first conjunct equals $\mathbf{s}_P \chi$ and the second one equals $\hat{B}_a \mathbf{s}_P (\langle \chi \rangle \psi)$, thus their conjunction equals $\mathbf{s}_P \phi$. ■

And now:

Corollary 7.18. *If $\phi \in \mathcal{L}_{\text{PAL}}$ is in standard form, then $w \models_{\text{PAL}} \phi$ iff $w, \epsilon \models_{\text{PSAL}} \mathbf{t} \phi$.*

Proof. $w \models_{\text{PAL}} \phi$ iff (by Lemma 7.16) $w \models_{\text{PAL}} \mathbf{s}_P \phi$ iff (by Lemma 7.15 and the fact that $\mathbf{s}_P \phi$ is announcement-free) $w, \epsilon \models_{\text{PSAL}} \mathbf{s}_P \phi$ iff (by the previous Lemma) $w, \epsilon \models \mathbf{s}_A \mathbf{t} \phi$ iff $w, \epsilon \models \mathbf{t} \phi$. ■

Now, let \mathbf{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 $\langle \top \rangle \phi \leftrightarrow \phi$. More explicitly,

$$\begin{aligned} \mathbf{sf} \top &= \top, & \mathbf{sf} p &= p, & \mathbf{sf} \neg \phi &= \neg \mathbf{sf} \phi, & \mathbf{sf} \hat{B}_a \phi &= \hat{B}_a \mathbf{sf} \phi, \\ \mathbf{sf} \langle \phi \rangle \psi &= \langle \mathbf{sf} \phi \rangle \mathbf{sf} \psi & (\text{if } \psi \text{ is not of the form } \langle \chi_1 \rangle \chi_2), \\ \mathbf{sf} \langle \phi \rangle \langle \psi \rangle \chi &= \langle \langle \mathbf{sf} \phi \rangle \mathbf{sf} \psi \rangle \mathbf{sf} \chi. \end{aligned}$$

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

Proof of Thm. 7.14. Simply note that $\tau = \mathbf{t} \mathbf{sf}$. We have: $w \models \phi$ iff $w \models \mathbf{sf} \phi$ iff (by the previous corollary) $w, \epsilon \models \mathbf{t} \mathbf{sf} \phi$. ■

7.6 Common belief

As mentioned above, a notion of common belief (or knowledge) does not make much sense in the setting of AA, wherein 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 $t \bowtie \beta$. We can read $CB_G \phi$ in terms

of the transitive closure of the union of these \mathbf{R}_a 's, $\mathbf{R}_G = (\bigcup_{a \in G} \mathbf{R}_a)^T$,⁵ so that $w, \alpha \models CB_G\phi$ iff $(w, \alpha)\mathbf{R}_G(t, \beta)$ implies $(t, \beta) \models \phi$. Equivalently,

Definition 7.19. $w, \alpha \models CB_G\phi$ 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 7.20. *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 7.21. Let

$$\alpha_2 = p. \neg B_a p. ab. B_b B_a p. bc. B_c p. ca,$$

as in Example 7.2, 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

⁵Here, R^T denotes the transitive closure of a relation R .

$\neg B_a p$, and the readings ab and bc . Therefore, $t \not\vdash \beta_i$ 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

$$\beta_1 = p. \neg B_a p. ab. bc, \quad \beta_2 = p. ab.$$

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 $R_c w t$, $\beta_2 \triangleright_c \epsilon$, $t \not\vdash \epsilon$ and $t, \epsilon \not\models p$. Therefore, $w, \alpha_2 \not\models CB_{ab} B_c p$.

Discussion. This chapter has 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; three intuitive interpretations of the ‘view’ of an agent have been provided, along with the sound and complete logic of each of them (for the class of models with empty histories). A proposal for common belief in this framework has been given.

This 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 if a sends p to b (announcement $B_a p$) and if after b eventually receives this announcement, b acknowledges this by sending $B_a p$ to a (announcement of $B_b B_a p$). In this way, we can model as diverse systems as: the internet protocol TCP guaranteeing correctness of initial sequences of packages [60, 99] (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 [65, 5] (an example of full synchronization for all agents after simultaneous partial synchronization for subsets of size two, namely for 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) [62, 54].

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.

An appealing endeavour of future research is 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 the non-reflexive \triangleright_a relation for belief presented here. 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 Chapter 6.

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

Quantification over Asynchronous Information Change

We propose a logic AAA for Arbitrary Asynchronous Announcements. In this logic, the sending and receiving of messages that are announcements are separated and represented by distinct modalities. Additionally, the logic has a modality that represents quantification over information change in the shape of sequences of sending and receiving events, called histories. We present a complete however infinitary axiomatisation, and various results for the logical semantics, wherein we consider both how the logic is different from asynchronous announcement logic AA and how the logic is different from arbitrary public announcement logic APAL. We also address the expressivity and we demonstrate the preservation of an extended fragment of positive formulas (wherein negations do not occur before epistemic modalities). Finally, we present work in progress on the logic AAM of Asynchronous Action Models and the logic AAAM of Arbitrary Asynchronous Action Models.¹

— All right, Marge, I'll tell you, but first you have to promise you will not get mad.

— I promise you that, no matter what you're about to say, I will get mad, because I always do when you make me promise I won't.

Conversation between Homer and Marge Simpson, in The Simpsons S3E14, written by Jay Kogen and Wallace Wolodarsky.

THE PRESENT CHAPTER showcases a generalisation of Asynchronous Announcement Logic in several ways: primarily, to the logic AA of asynchronous announcements we now add a quantifier $[!]\phi$ for ‘after any sequence of events, ϕ holds’. This is motivated by a similar quantifier in the logic APAL [7], that stands for ‘after any (arbitrary) announcement, ϕ ’. Clearly,

¹This chapter is based on the paper

Philippe Balbiani, Hans van Ditmarsch, and Saúl Fernández González. Quantifying over asynchronous information change. In *Advances in Modal Logic*, 2020.

in the asynchronous setting we cannot have it merely quantifying over unreceived announcements, as this would not affect the beliefs or knowledge of agents. As the order of reception of announcements may vary greatly and may take place much later after an announcement, and possibly many subsequent announcements, have been sent, the natural form of quantification is therefore over arbitrary sequences of such sending and receiving events. We show that the resulting logic AAA has a complete axiomatisation, and varies in crucial respects from the motivating precedent APAL [7]. Such a logic of arbitrary asynchronous announcement may be useful for diverse tasks such as asynchronous epistemic planning, the formalisation of epistemic protocols in distributed computing.

One particular further generalisation is also presented in some detail: namely, a similar quantification over asynchronous non-public events, in the sense of events that are not known to be eventually received by all agents (such as an agent a privately receiving information on a proposition p , while an agent b receives the information that a is privately receiving p although not simultaneous with a). This utilises an asynchronous notion of *action models* [17]. It is known from the works of Hales and collaborators [56, 57] that quantification over action models behaves much better than quantification over announcements: it is decidable, the quantifier can be eliminated from the language, and given $\langle ! \rangle \phi$, for ‘there is an action model after which ϕ ’, an action model can be synthesised that if executable always results in ϕ . One would conjecture similar results for quantifying over asynchronous action models. In particular, asynchronous synthesis seems a highly desirable future goal.

The structure of this chapter is the following: Section 8.1 introduces Arbitrary Asynchronous Announcement logic AAA and presents some semantics results. Section 8.2 addresses expressivity, and Section 8.3 studies the preservation (after history extension) of the fragment of *positive formulas*. Section 8.4 presents a complete infinitary axiomatisation. Finally, Section 8.5 addresses the generalisation of our results to a logic for quantification over asynchronous action models.

8.1 The logic AAA

Syntax. Let A be a finite set of epistemic agents and Prop a countable set of propositional variables. We consider the following language \mathcal{L}_{AAA} :

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

where $p \in \mathbf{Prop}$, $a \in A$.

We define duals $B_a\phi = \neg\hat{B}_a\neg\phi$, $[a]\phi = \neg\langle a\rangle\neg\phi$, $[\psi]\phi = \neg\langle\psi\rangle\neg\phi$, $[!]\phi = \neg\langle!\rangle\neg\phi$.

Note that \mathcal{L}_{AA} (see Section 6.1) is the fragment of this language without the $\langle!\rangle$ modality.

Given a word $\alpha \in (A \cup \mathcal{L}_{AAA})^*$, we define the notions $|\alpha|_!$, $|\alpha|_a$, $\alpha|_!$, $\alpha|_a$ exactly as in Section 6.1, and we say α is a *history* whenever $|\beta|_! \geq |\beta|_a$ for all $a \in A$ and $\beta \sqsubseteq \alpha$.

We denote by \mathcal{H} the set of histories. We define a *view relation* in \mathcal{H} as in AA (see Def. 6.5):

$$\alpha \triangleright_a \beta \text{ iff } \alpha|_a = \beta|_a = \beta|_!.$$

Semantics. To define the semantics we will use the following well-founded preorder. First, we define $\deg_B \phi$, $\deg_! \phi$ and $\|\phi\|$ recursively: for $k = !, B$,

$$\begin{array}{ll} \deg_k p = 0 & \|p\| = 2 \\ \deg_k \top = 0 & \|\top\| = 1 \\ \deg_k (\neg\phi) = \deg_k \phi & \|\neg\phi\| = \|\phi\| + 1 \\ \deg_k (\phi \wedge \psi) = \max\{\deg_k \phi, \deg_k \psi\} & \|\phi \wedge \psi\| = \|\phi\| + \|\psi\| \\ \deg_k (\langle a \rangle \phi) = \deg_k \phi & \|\langle a \rangle \phi\| = \|\phi\| + 2 \\ \deg_k (\langle \phi \rangle \psi) = \deg_k \phi + \deg_k \psi & \|\langle \phi \rangle \psi\| = 2\|\phi\| + \|\psi\| \\ \deg_B (\hat{B}_a \phi) = \deg_B \phi + 1 & \|\hat{B}_a \phi\| = \|\phi\| + 1 \\ \deg_! (\hat{B}_a \phi) = \deg_! \phi & \\ \deg_B (\langle ! \rangle \phi) = \deg_B \phi & \|\langle ! \rangle \phi\| = \|\phi\| + 1 \\ \deg_! (\langle ! \rangle \phi) = \deg_! \phi + 1 & \end{array}$$

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

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

Finally, for pairs (α, ϕ) we set: $\deg_k(\alpha, \phi) = \deg_k \alpha + \deg_k \phi$, 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

$$\left\{ \begin{array}{l} \deg_!(\alpha, \phi) < \deg_!(\beta, \psi), \text{ or} \\ \deg_!(\alpha, \phi) = \deg_!(\beta, \psi) \ \& \ \deg_B(\alpha, \phi) < \deg_B(\beta, \psi), \text{ or} \\ \deg_!(\alpha, \phi) = \deg_!(\beta, \psi) \ \& \ \deg_B(\alpha, \phi) = \deg_B(\beta, \psi) \ \& \ \|(\alpha, \phi)\| < \|(\beta, \psi)\|. \end{array} \right.$$

We interpret formulas on models (W, R, V) (where $R : A \rightarrow \mathcal{P}(W^2)$ and $V : \mathbf{Prop} \rightarrow \mathcal{P}(W)$) with respect to pairs (w, α) where $w \in W$ and $\alpha \in \mathcal{H}$.

$w \bowtie \epsilon$	always;
$w \bowtie \alpha.\phi$	iff $w \bowtie \alpha$ and $w, \alpha \models \phi$;
$w \bowtie \alpha.a$	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 $w, \alpha \models \phi_i, i = 1, 2$;
$w, \alpha \models \langle a \rangle \phi$	iff $ \alpha _a < \alpha _!$ and $w, \alpha.a \models \phi$;
$w, \alpha \models \langle \psi \rangle \phi$	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 $R_a w t, \alpha \triangleright_a \beta, t \bowtie \beta$
$w, \alpha \models \langle ! \rangle \phi$	iff $w, \alpha \models \langle \beta \rangle \phi$ for some $\beta \in (\mathcal{L}_{AA} \cup A)^*$.

Table 8.1: Semantics of AAA

We define the relations “ w agrees with α ” ($w \bowtie \alpha$) and “ (w, α) satisfies ϕ ” ($w, \alpha \models \phi$) by \ll -induction in the same way as the semantics of AA but with an extra clause for the arbitrary announcement modality: see Table 8.1.

Note that the $\langle ! \rangle$ modality only quantifies over words wherein $\langle ! \rangle$ does not occur. This is to avoid a circular definition. The dual of $\langle ! \rangle$ is read as follows: $w, \alpha \models [!] \phi$ if and only if, for every possible sequence $\beta \in (\mathcal{L}_{AA} \cup A)^*$, it is the case that $w, \alpha \models [\beta] \phi$.

We use the same notions of validity as in Def. 6.24: ϕ is ϵ -valid if $M, w, \epsilon \models \phi$ for every model M and every world w , and $*$ -valid if $M, w, \alpha \models \phi$ for every model, world and history such that $w \bowtie \alpha$.

Bisimulation. The notion of bisimulation in this framework is, perhaps surprisingly, the usual notion of bisimulation between Kripke models: given (W, R, V) and (W', R', V') , a *bisimulation* is a relation $Z \subseteq W \times W'$ such that, if $w Z w'$:

- i. $w \in V(p)$ iff $w' \in V'(p)$;
- ii. if $R_a w v$, then there exists $v' \in W'$ such that $R'_a w' v'$ and $v Z v'$;
- iii. if $R'_a w' v'$, then there exists $v \in W$ such that $R_a w v$ and $v Z v'$.

As one might expect, we have the following:

Proposition 8.1. *Let Z be a bisimulation such that wZw' , and let $(\alpha, \phi) \in \mathcal{H} \times \mathcal{L}_{AAA}$. We have:*

$$w, \epsilon \models \langle \alpha \rangle \phi \text{ iff } w', \epsilon \models \langle \alpha \rangle \phi.$$

Proof. By \ll induction on (α, ϕ) . Trivial for the cases where $(\alpha, \phi) = (\epsilon, \top)$ and (ϵ, p) . For the case where $(\alpha, \phi) = (\beta.a, \top)$, we note that $w \bowtie \beta.a$ iff $w \bowtie \beta$ and $w, \beta.a \models \top$ iff $w, \beta \models \top$, and thus we can apply induction hypothesis, for $(\beta, \top) \ll (\beta.a, \top)$. For the case $(\alpha, \phi) = (\beta.\psi, \top)$, we note that $(\beta, \psi) \ll (\beta.\psi, \top)$.

For the cases $(\alpha, \phi) = (\alpha, \neg\psi)$ and $(\alpha, \psi) = (\alpha, \psi_1 \wedge \psi_2)$, we note that $(\alpha, \psi) \ll (\alpha, \neg\psi)$ and $(\alpha, \psi_i) \ll (\alpha, \psi_1 \wedge \psi_2)$.

For the case $(\alpha, \phi) = (\alpha, \hat{B}_a\psi)$, we have: $w \bowtie \alpha$ iff $w' \bowtie \alpha$ by induction hypothesis applied to (α, \top) . If $w, \alpha \models \hat{B}_a\psi$, then there is some $v \in W$ and some history β such that R_avv , $\alpha \triangleright_a \beta$, $v \bowtie \beta$ and $v, \beta \models \psi$. But then there is some $v' \in W'$ with vZv' and $R_av'v'$ and thus, by induction hypothesis applied to $(\beta, \psi) \ll (\alpha, \hat{B}_a\psi)$, we have $v' \bowtie \beta$, $v', \beta \models \psi$ and thus $w', \alpha \models \hat{B}_a\psi$. The converse is analogous.

For the cases $(\alpha, \psi) = (\alpha, \langle a \rangle \psi)$ and $(\alpha, \psi) = (\alpha, \langle \theta \rangle \psi)$, we note that $(\alpha.a, \psi) \ll (\alpha, \langle a \rangle \psi)$ and $(\alpha.\theta, \psi) \ll (\alpha, \langle \theta \rangle \psi)$.

For the case $(\alpha, \phi) = (\alpha, \langle ! \rangle \psi)$, we have: on the one hand, $w \bowtie \alpha$ iff $w' \bowtie \alpha$ by induction hypothesis applied to (α, \top) . On the other hand, suppose $w, \alpha \models \langle ! \rangle \psi$. Then $w, \alpha \models \langle \beta \rangle \psi$ for some history β which does not contain any occurrences of $\langle ! \rangle$. Therefore $\text{deg}_! \langle \beta \rangle \psi < \text{deg}_! \langle ! \rangle \psi$, and thus by induction hypothesis $w', \alpha \models \langle \beta \rangle \psi$, which entails $w', \alpha \models \langle ! \rangle \psi$. The converse is analogous. \blacksquare

Under certain constraints, if two states satisfy the same formulas, they are bisimilar. Indeed:

Proposition 8.2. *Let (W, R, V) and (W', R', V') be two models such that $R_a[w]$ and $R'_a[w']$ are finite sets for all $w \in W$, $w' \in W'$. Set wZw' iff, for all $(\alpha, \phi) \in \mathcal{H} \times \mathcal{L}$, $w, \epsilon \models \langle \alpha \rangle \phi$ iff $w', \epsilon \models \langle \alpha \rangle \phi$. Then Z is a bisimulation.*

Proof. It is obvious that condition i. is satisfied. Now, suppose condition ii. fails. That is, for some $v \in W$, we have R_avv but for all (the finitely many) $v' \in W'$ such that $R_av'v'$ it is not the case that vZv' . Let $R'_a[w'] = \{v'_1, \dots, v'_n\}$. For each v'_i there exists some pair (α_i, ϕ_i) such that either $v, \epsilon \models \langle \alpha_i \rangle \phi_i$ but $v'_i, \epsilon \not\models \langle \alpha_i \rangle \phi_i$, or $v, \epsilon \not\models \langle \alpha_i \rangle \phi_i$ but $v'_i, \epsilon \models \langle \alpha_i \rangle \phi_i$. Let $\theta_i = \langle \alpha_i \rangle \psi_i$ in the former case and $\theta_i = \neg \langle \alpha_i \rangle \psi_i$ in the latter, and call $\psi = \bigwedge_{i=1}^n \theta_i$. Note that $v, \epsilon \models \psi$ and thus $w, \epsilon \models \hat{B}_a\psi$. But then by the definition of Z we have

that $w', \epsilon \models \hat{B}_a \psi$, and thus w' has a successor satisfying each formula θ_i : contradiction. Condition iii. is proven similarly. ■

Properties of belief. Similarly to Asynchronous Announcement Logic, $B_a \phi \rightarrow \phi$ is ϵ -valid, (as long as the relation R_a is reflexive) but it is not $*$ -valid. Other properties of our doxastic modality, however, are $*$ -valid. Let $\mathcal{S5}$ denote the class of models where the relations R_a are equivalence relations. Similarly to AA (and with a virtually identical proof), we have:

Proposition 8.3. *Let $\phi \in \mathcal{L}$. Then:*

- i. $\mathcal{S5} \models^* B_a \phi \rightarrow \neg B_a \neg \phi$
- ii. $\mathcal{S5} \models^* B_a \phi \rightarrow B_a B_a \phi$
- iii. $\mathcal{S5} \models^* \neg B_a \phi \rightarrow B_a \neg B_a \phi$

Belief before and after update. If an agent will believe ϕ after a certain sequence of events then the agent believes that there is a sequence of events after which ϕ holds, but not the other way around. Indeed:

Proposition 8.4. $\models^\epsilon \langle ! \rangle \hat{B}_a \phi \rightarrow \hat{B}_a \langle ! \rangle \phi$, whereas $\not\models^\epsilon \hat{B}_a \langle ! \rangle \phi \rightarrow \langle ! \rangle \hat{B}_a \phi$.

Proof. Let model $M = (W, R, V)$ and $s \in W$ be given, and let $\alpha \in (\mathcal{L}_{AA} \cup A)^*$ be such that $s, \epsilon \models \langle \alpha \rangle \hat{B}_a \phi$. Therefore α is a history, $s \triangleright \alpha$ and $s, \alpha \models \hat{B}_a \phi$, so that there are t, β such that $R_a s t$, $\alpha \triangleright_a \beta$, $t \triangleright \beta$, and $t, \beta \models \phi$. As $t \triangleright \beta$ and $t, \beta \models \phi$, it follows that $t, \epsilon \models \langle \beta \rangle \phi$. It therefore follows that $t, \epsilon \models \langle ! \rangle \phi$. Finally, as $R_a s t$, $\epsilon \triangleright_a \epsilon$ and $t \triangleright \epsilon$ we conclude $s, \epsilon \models \hat{B}_a \langle ! \rangle \phi$.

On the other hand, $\hat{B}_a \langle ! \rangle \phi \rightarrow \langle ! \rangle \hat{B}_a \phi$ is not ϵ -valid. Consider the model $M = (W, R, V)$ for a single agent a and atom p and where $W = \{s, t\}$, $R_a = W^2$, and $V(p) = \{s\}$. We then have that $s, \epsilon \models \hat{B}_a \langle ! \rangle B_a \neg p$, because $s, \epsilon \models \hat{B}_a \langle \neg p.a \rangle B_a \neg p$ (because $t, \epsilon \models \langle \neg p.a \rangle B_a \neg p$), whereas clearly $s, \epsilon \not\models \langle ! \rangle \hat{B}_a B_a \neg p$. ■

Church-Rosser and McKinsey Let us see that neither of the formulas

$$\text{(CR)} \quad \langle ! \rangle [!] \phi \rightarrow [!] \langle ! \rangle \phi \quad \text{(McK)} \quad [!] \langle ! \rangle \rightarrow \langle ! \rangle [!] \phi$$

are valid. It is known from APAL that these properties are valid for arbitrary announcement on the class of $\mathcal{S5}$ models (where all accessibility relations are equivalence relations) [7]. We address the case $\mathcal{S5}$ at the end of this paragraph.

First let us see (McK) is not sound. Let $\phi = [a] \perp$. Then ϕ will be satisfied at a pair w, α if and only if $|\alpha|_a = |\alpha|$. For any history β it is the case that

$|\beta|_a \leq |\beta|_!$: let $\delta_\beta = a\dots a$ be the concatenation of $|\beta|_! - |\beta|_a$ times the letter a . Then, for every β there exists a word δ_β such that $w, \epsilon \models [\beta]\langle\delta_\beta\rangle[a]\perp$. However, $\langle!\rangle[!][a]\perp$ is never satisfied: indeed, for any history β , let δ_β be a concatenation of the formula \top enough times so that $|\beta\delta_\beta|_! > |\beta|_a$. Then we have $w, \epsilon \not\models \langle\beta\rangle[\delta_\beta][a]\perp$.

Let us now see a counterexample for (CR).² Consider the following one-agent model:

Let $W = \{w_1, w_2, w_3\}$, $R_a = \{(w_1, w_2), (w_2, w_2), (w_2, w_3)\}$ and $V(p) = \{w_1, w_2\}$. We have that $w_1, \epsilon \models \langle!\rangle[!]\hat{B}_a\top$. Indeed, consider the history $p.a.[a]\perp.a$. We can easily prove the following by induction on ϕ :

If β is a history having $p.a.[a]\perp.a$ as a prefix, then for all ϕ , $w_1, \beta \models \phi$ iff $w_2, \beta \models \phi$.

In particular, any β having $p.a.[a]\perp.a$ as a prefix will be executable at w_1 iff it is executable at w_2 . Now, take any sequence γ such that $p.a.[a]\perp.a.\gamma$ is executable at w_1 . There exists a β such that $p.a.[a]\perp.a.\gamma \triangleright_a \beta$ and β is executable at w_1 . Note that β is necessarily of the form $\beta = p.a.[a]\perp.a.\gamma'$ for some γ' . But this means, by the previous remark, that β is executable at w_2 , and thus $w_1, p.a.[a]\perp.a.\gamma \models \hat{B}_a\top$, which means $w_1, \epsilon \models \langle p.a.[a]\perp.a \rangle[!]\hat{B}_a\top$ and thus $w_1, \epsilon \models \langle!\rangle[!]\hat{B}_a\top$.

However, $w_1, \epsilon \not\models [!]\langle!\rangle\hat{B}_a\top$: indeed, consider the sequence $B_ap.a$. It is never the case that $w_1, B_ap.a \models \langle\beta\rangle\hat{B}_a\top$ for any announcement, given that, whenever $B_ap.a.\beta \triangleright_a \gamma$, γ has B_ap as its first formula, and therefore γ cannot be compatible with w_2 , since $w_2, \epsilon \not\models B_ap$.

(CR) is not sound in general in APAL (this can be seen via a similar counterexample), but, as said, only with equivalence relations. Whether (CR) is sound on AAA on the class of $\mathcal{S}5$ models is an open question.

8.2 Expressivity of AAA

We assume the usual terminology to compare the expressivity of logics or logical languages with respect to a semantics and a class of models. Given two languages \mathcal{L}_1 and \mathcal{L}_2 interpreted over the same class \mathcal{C} of models, we say that \mathcal{L}_1 is *at least as expressive as* \mathcal{L}_2 *with respect to* \mathcal{C} iff for all formulas $\phi_2 \in \mathcal{L}_2$, there exists a formula $\phi_1 \in \mathcal{L}_1$ such that for all models $\mathfrak{M} \in \mathcal{C}$, $\mathfrak{M} \models \phi_1$ iff $\mathfrak{M} \models \phi_2$.

If \mathcal{L}_1 is at least as expressive as \mathcal{L}_2 and \mathcal{L}_2 is at least as expressive as \mathcal{L}_1 then \mathcal{L}_1 and \mathcal{L}_2 are *as expressive*. If \mathcal{L}_1 is at least as expressive as \mathcal{L}_2 and


²Thanks to Louwe Kuiper for this counterexample.

This is a contradiction. ■

It seems likely, although we have not proven this, that on the class $\mathcal{S5}$ for a single agent the $\langle ! \rangle$ modality is definable in AA, therefore making AAA as expressive as AA in that particular case. However, without imposing any frame properties single-agent AAA is more expressive than AA, again shown similarly to the previous proposition and [7, Prop. 3.14]

Proposition 8.6. \mathcal{L}_{AAA} is more expressive than \mathcal{L}_{AA} for a single agent.

Proof. For a single agent we consider the formula $\langle ! \rangle(B_a p \wedge B_a \neg B_a p)$ and proceed as in Prop. 8.5, where in this case we observe that in model N' it holds that $s', (p \vee \neg q).a \models B_a p \wedge B_a \neg B_a p$.



■

A logic is called *compact* if a set of formulas is satisfiable whenever any finite subset thereof is satisfiable.

Proposition 8.7. The logic AAA is not compact.

Proof. Using the above expressivity results, this can be shown by considering the set of formulas

$$\{\langle ! \rangle(B_a p \wedge B_a \neg B_b p)\} \cup \{\neg \langle \beta \rangle(B_a p \wedge B_a \neg B_b p) : \beta \in (\mathcal{L}_{AA} \cup A)^*\}.$$

This set is not satisfiable, but any finite subset thereof is satisfiable: indeed, we use that some variable q must necessarily not occur in such a subset. We then consider M, M' as above. The q -less finite subset will be satisfied at s' . ■

8.3 Positive formulas

In modal logic, the fragment of the language where negations do not bind epistemic modalities is known as the *positive* fragment [22, 40, 7]. It corresponds to the universal fragment in first-order logic. It has the property that

it preserves truth under submodels. In AAA, preservation under submodels is formalised by preservation after history extension. A formula $\phi \in \mathcal{L}_{AAA}$ is *preserved* iff $\models^* \phi \rightarrow [!]\phi$. We wish therefore to identify a fragment of the language \mathcal{L} that guarantees preservation.

For AA, it is shown in Prop. 6.46 that the fragment

$$\phi ::= p \mid \neg p \mid \perp \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid B_a \phi,$$

that corresponds in a very direct way to the universal fragment, is preserved.

For AAA we can expand that frontier, in the direction earlier taken in [40] for (synchronous) public announcements, where an extra inductive clause $[\neg\phi]\phi$ is added; this is further expanded in [7] with an inductive clause $[!]\phi$ (where $[!]$ is the APAL quantifier over announcements). We will only define a fairly minimal extension and subsequently present some of the difficulties in obtaining a result analogous to those in [40, 7], and what the desirable final goal seems to be.

The proof uses Lemma 6.44, which we repeat below:

Lemma 6.44 *Let α, β and γ be histories. If $\gamma \preceq \alpha$ and $\alpha \triangleright_a \beta$ then there exists a history δ such that $\gamma \triangleright_a \delta$ and $\delta \preceq \beta$. In other words, $(\preceq \circ \triangleright_a) \subseteq (\triangleright_a \circ \preceq)$.*

Recall that, by Definition 6.42, $\alpha \preceq \beta$ iff (1) $\alpha \upharpoonright_1 \subseteq \beta \upharpoonright_1$, (2) for all $a \in A$ $|\alpha|_a \leq |\beta|_a$, and (3) for any state s in any model $s \bowtie \beta$ implies $s \bowtie \alpha$. Recall as well that $\alpha \sqsubseteq \beta$ implies $\alpha \preceq \beta$, but not the converse.

Consider the following *positive fragment* \mathcal{L}_+ :

$$\phi ::= p \mid \neg p \mid \perp \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid B_a \phi \mid [!]\phi.$$

We show that positive formulas are preserved.

Proposition 8.8 (Positive implies preserved).

Let $\phi \in \mathcal{L}_+$. Then $\models_ \phi \rightarrow [!]\phi$.*

Proof. We need to prove the following proposition:

Let $\phi \in \mathcal{L}_+$. For all models $M = (W, R, V)$ and $s \in W$, and for all histories $\alpha: s, \epsilon \models [\alpha](\phi \rightarrow [!]\phi)$.

This is equivalent to:

Let $\phi \in \mathcal{L}_+$. For all models $M = (W, R, V)$ and $s \in W$, and for all histories α, β such that $\alpha \sqsubseteq \beta$: $s, \epsilon \models [\alpha]\phi$ implies $s, \epsilon \models [\beta]\phi$.

A standard inductive proof on the structure of ϕ fails because in the case $B_a\phi$ we would need that if $\alpha \sqsubseteq \beta$ and $\beta \triangleright_a \delta$, then there is a γ with $\gamma \sqsubseteq \delta$ and $\alpha \triangleright_a \gamma$. Such a γ may not exist, namely if many yet unread announcements in δ precede the a in δ that corresponds to the last a in α . However, we can then still find a γ such that $\gamma \preceq \delta$. Therefore, it suffices to show:

Lemma 8.9. *Let $\phi \in \mathcal{L}_+$. For all models $M = (W, R, V)$ and $s \in W$, and for all histories α, β such that $\alpha \preceq \beta$: $s, \epsilon \models [\alpha]\phi$ implies $s, \epsilon \models [\beta]\phi$.*

Proof. We show the following: Let $\phi \in \mathcal{L}_+$. For all models $M = (W, R, V)$ and $s \in W$, and for all histories α, β with $\alpha \preceq \beta$: if $s \bowtie \alpha$ and $s, \alpha \models \phi$, then if $s \bowtie \beta$ it holds that $s, \beta \models \phi$.

The proof is by induction on the structure of (simple positive) ϕ :

Case \perp . If $s, \epsilon \models [\alpha]\perp$, then $s \not\bowtie \alpha$, and thus $s \not\bowtie \beta$, by definition of \preceq , which means $s, \epsilon \models [\beta]\perp$.

Case atoms. If $s, \epsilon \models [\alpha]p$, then either $s \not\bowtie \alpha$, in which case $s \not\bowtie \beta$ and thus $s, \epsilon \models [\beta]p$, or $s \bowtie \alpha$ and $s \in V(p)$, in which case $s, \epsilon \models [\beta]p$ as well. The case for $\phi = \neg p$ is analogous.

Case conjunction. If $s, \epsilon \models [\alpha](\phi_1 \wedge \phi_2)$, and assuming $s \bowtie \beta$ (for otherwise it is trivial), we have that $s, \alpha \models \phi_i$ for $i = 1, 2$ and thus, by induction hypothesis, $s, \beta \models \phi_i$, whence $s, \epsilon \models [\beta](\phi_1 \wedge \phi_2)$.

Case disjunction. Analogous.

Case belief. Suppose $s, \epsilon \not\models [\beta]B_a\phi$. Then $s, \epsilon \models \langle \beta \rangle \hat{B}_a \neg\phi$, which means there exist t, δ with R_ast , $\beta \triangleright_a \delta$ and $t, \delta \not\models \phi$. By Lemma 6.44, there is a γ with $\alpha \triangleright_a \gamma$ and $\gamma \preceq \delta$, which gives, by induction hypothesis, $t, \delta \not\models \phi$ and thus $s, \alpha \not\models B_a\phi$.

Case $[\!] \phi$. Suppose $s, \epsilon \not\models [\beta][\!] \phi$. This means that $s, \epsilon \models \langle \beta \rangle \langle [\!] \rangle \neg\phi$, i.e. there exists a word $\delta \in (\mathcal{L}_{AA} \cup A)^*$ such that $s \bowtie \beta.\delta$ and $s, \beta.\delta \not\models \phi$. Since $\alpha \preceq \beta.\delta$, this gives that $s \bowtie \alpha$ and $s, \alpha \not\models \phi$, and thus $s, \epsilon \not\models [\alpha][\!] \phi$. ■

This finishes the proof of Prop. 8.8. ■

This definition of preservation does not include a clause for announcement in the inductive definition of positive formulas, which merits some discussion. The obvious analogue of the $[\neg\phi]\psi$ clause from [40] would be $[\neg\phi.A]\psi$ (where ‘ A ’ represents some arbitrary permutation of all agents in A).⁴ Consider for instance a model M for one agent a consisting of four worlds for the four

⁴The possibly strange form of this clause where a negation appears has to do with the semantics of public announcement. In PAL, $M, w \models [\neg\phi]\psi$ iff $(M, w \models \neg\phi$ implies $M^\phi, w \models \psi)$ iff $(M, w \models \phi$ or $M^\phi, w \models \psi)$: in this disjunctive description, the negation disappears.

possible valuations of variables p and q , all four indistinguishable for a . Let w be the world where p and q are true. We now have that $w, \epsilon \models [q.a]B_a q$ whereas $w, p \not\models [q.a]B_a q$, because the a in history ‘ $q.a$ ’ reads announcement q in the first case whereas she reads announcement p in the second case. As long as agent a has not received announcement q , she remains uncertain about the value of q .

According to a clause of the form $[\neg\phi.A]\psi$, $[q.a]B_a q$ should be a positive formula. But this formula is not preserved, for we do not have $\models^* \phi \rightarrow [!]\phi$.

Beyond that, having $[\neg\phi.A]\psi$ be positive would also make $[p.a.p.a]B_a p$ positive in the above example: should we then not want $[p.p.a.a]B_a p$, where both announcements are only read after they have been announced, to be positive as well?

It seems that the above definition of preservation effectively rules out the inclusion of announcements and read modalities in a positive fragment. Now, it could be the case that the positive fragment as defined syntactically characterizes the preserved formulas, analogous to van Benthem’s result for the (usual) positive fragment [22]. This is yet unknown.

One could alternatively define a notion of preservation with respect to ϵ -validity (so, ‘ ϵ -preserved’ could mean $\models^\epsilon \phi \rightarrow [!](\bigwedge_{a \in A}[a]\perp \rightarrow \phi)$). This would allow a more liberal fragment of positive formulas including the above examples. This should be further investigated, with the ultimate goal of a syntactic characterisation of ϵ -preservation.

Before moving on, let us point out another property of the positive fragment: when the believed formula ϕ is positive, and the accessibility relation reflexive, belief becomes factive.

Proposition 8.10. *Let $\phi \in \mathcal{L}_+$. For any model (W, R, V) such that R_a is reflexive, for all $s \in W$ and α such that $s \bowtie \alpha$, we have $s, \alpha \models B_a \phi \rightarrow \phi$. As a consequence, $\mathcal{S5} \models^* B_a \phi \rightarrow \phi$.*

Proof. Suppose $s, \alpha \models B_a \phi$. Consider $\beta = \alpha.\phi.a^k$, as constructed in the proof of Prop. 8.3. We have $R_a s s$, $\alpha \triangleright_a \beta$, and $s \bowtie \beta$, and thus $s, \beta \models \phi$. Moreover, since $\delta.\phi \sqsubseteq \alpha$ and $|\beta|_a = |\alpha|_a$, we have $\beta \preceq \alpha$. By Lemma 8.9, this entails $s, \alpha \models \phi$. ■

8.4 Axiomatisation of AAA

The axiomatisation of ϵ -validities of AAA and its completeness proof is based on the axiomatisation of AA (Thm. 6.53) and on that of APAL [7] and its completeness proof uses the method pioneered in [12, 6].

Given a symbol $\#$ we define a set AF of *admissible forms* as follows:

$$L ::= \#|B_a L|\phi \rightarrow L|\langle\alpha\rangle L,$$

where $\phi \in \mathcal{L}_{AAA}$, $a \in A$, $\alpha \in \mathcal{H}$, $L \in AF$. Given $L \in AF$ and $\phi \in \mathcal{L}_{AAA}$, the formula $L(\phi)$ is the result of substituting the unique occurrence of $\#$ in L by ϕ .

The following holds:

Lemma 8.11. *Let L be an admissible form. For all $M \in AF$ and for all modal formulas ϕ, ψ , if $L([\!]\phi) = M([\!]\psi)$ then $L = M$ and $\phi = \psi$.*

Proof. By induction on L . ■

The logic AAA consists of the following axioms and rules, for $\alpha \in \mathcal{H}$, $p \in \text{Prop}$, $a \in A$, $L(\#) \in AF$

(MP)	If $\vdash \phi$ and $\vdash \phi \rightarrow \psi$, then $\vdash \psi$
(Nec _B)	If $\vdash \phi$, then $\vdash B_a \phi$
(K _B)	$B_a(\phi \rightarrow \psi) \rightarrow (B_a \phi \rightarrow B_a \psi)$
(R _{⊤1})	$\langle\alpha.a\rangle\top \leftrightarrow \langle\alpha\rangle\top$ if $ \alpha _a < \alpha !$
(R _{⊤2})	$\langle\alpha.a\rangle\top \leftrightarrow \perp$ otherwise;
(R _{⊤3})	$\langle\alpha.\phi\rangle\top \leftrightarrow \langle\alpha\rangle\phi$;
(R _p)	$\langle\alpha\rangle p \leftrightarrow (\langle\alpha\rangle\top \wedge p)$;
(R _¬)	$\langle\alpha\rangle\neg\phi \leftrightarrow (\langle\alpha\rangle\top \wedge \neg\langle\alpha\rangle\phi)$;
(R _∨)	$\langle\alpha\rangle(\phi \vee \psi) \leftrightarrow (\langle\alpha\rangle\phi \vee \langle\alpha\rangle\psi)$;
(R _B)	$\langle\alpha\rangle\hat{B}_a\phi \leftrightarrow (\langle\alpha\rangle\top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a\langle\beta\rangle\phi)$;
([\!]-elim)	$L([\!]\phi) \rightarrow L([\beta]\phi)$ (where $\beta \in (\mathcal{L}_{AA} \cup A)^*$);
([\!]-int ^ω)	If $\vdash L([\beta]\phi)$ for all $\beta \in (\mathcal{L}_{AA} \cup A)^*$, then $\vdash L([\!]\phi)$

Remark 8.12. If we remove the last two lines of the above table we obtain the logic AA, defined in Chapter 6 for the language \mathcal{L}_{AA} (except written in its dual form).

Completeness proof. A *theory* is a set of formulas T such that:

- i. $AAA \subseteq T$;
- ii. T is closed under Modus Ponens: if $\phi, \phi \rightarrow \psi \in T$, then $\psi \in T$;
- iii. T is closed under the following rule:
If $L([\beta]\phi) \in T$ for all $\beta \in (\mathcal{L}_{AA} \cup A)^*$, then $L([\!]\phi) \in T$.

A theory is *consistent* if $\perp \notin T$. Note that AAA is the least consistent theory, and \mathcal{L} is the only inconsistent theory.

A consistent theory is *maximal* if no proper superset of T is a consistent theory.

The following holds:

Lemma 8.13. *Given a theory T , a formula ψ , and an agent $a \in A$, the sets $T_{B_a} = \{\phi : B_a\phi \in T\}$ and $T_\psi = \{\phi : \psi \rightarrow \phi \in T\}$ are also theories.*

Moreover, $T \subseteq T_\psi$, $\psi \in T_\psi$ and, if $\neg\psi \notin T$, then T_ψ is consistent.

Proof. Checking the first item is easy: if $\phi \in \text{AAA}$, then $B_a\phi \in \text{AAA}$ (by necessitation) and $\psi \rightarrow \phi \in \text{AAA}$ (by classical propositional logic). Therefore $B_a\phi \in T$ and $\psi \rightarrow \phi \in T$, and thus $\phi \in T_{B_a} \cap T_\psi$.

T_{B_a} is closed under modus ponens because if $\phi \rightarrow \theta \in T_{B_a}$ and $\phi \in T_{B_a}$, then $B_a(\phi \rightarrow \theta), B_a\phi \in T$, which by the K axiom plus modus ponens gives $B_a\theta \in T$ and thus $\theta \in T_{B_a}$. For T_ψ , suppose $\phi \rightarrow \theta, \phi \in T_\psi$. Then $\psi \rightarrow (\phi \rightarrow \theta) \in T$ and $\psi \rightarrow \phi \in T$. But note that the former is logically equivalent to $(\psi \rightarrow \phi) \rightarrow (\psi \rightarrow \theta)$, and, since T is closed under logical equivalence, this means by modus ponens that $\psi \rightarrow \theta \in T$ and thus $\theta \in T_\psi$.

For the third condition, suppose $L([\beta]\phi) \in T_{B_a}$ for all β . Then $B_aL([\beta]\phi) \in T$ for all β and, since $B_aL(\#)$ is an admissible form, then $B_aL([\!]\psi) \in T$, and thus $L([\!]\phi) \in T_{B_a}$. If $L([\beta]\phi) \in T_\psi$ for all β , then $\psi \rightarrow L([\beta]\phi) \in T$ for all β and, again, since $\psi \rightarrow L(\#)$ is an admissible form, this entails $\psi \rightarrow L([\!]\phi) \in T$ and therefore $L([\!]\phi) \in T_\psi$.

With respect to the last statement: $\psi \in T_\psi$ because $\psi \rightarrow \psi$ is a tautology; if $\neg\psi \notin T$, then $\psi \rightarrow \perp \notin T$ thus $\perp \notin T_\psi$, and if $\phi \in T$, then (since $\phi \rightarrow (\psi \rightarrow \phi)$ is a tautology) $\psi \rightarrow \phi \in T$ and thus $\phi \in T_\psi$. ■

We also have:

Proposition 8.14 (Lindenbaum's Lemma). *A consistent theory can be extended to a maximal consistent theory.*

Proof. Let T_0 be a consistent theory. Let $\{\phi_k : k \in \omega\}$ be an enumeration of the formulas in \mathcal{L} where each formula appears infinitely many times. For $k \in \omega$ we will construct a consistent theory T_{k+1} , which is a superset of T_k , as follows:

- i. If $\neg\phi_k \notin T_k$, then $T_{k+1} = (T_k)_{\phi_k}$;
- ii. If $\neg\phi_k \in T_k$ and ϕ_k is of the form $L([\!]\psi)$, then there must exist some $\beta \in (\mathcal{L}_{AA} \cup A)^*$ such that $L([\beta]\psi) \notin T_k$ (for otherwise, by rule iii., we would have that $\phi_k \in T_k$: contradiction). We set $T_{k+1} = (T_k)_{-L([\beta]\psi)}$.

iii. If $\neg\phi_k \in T_k$ and ϕ_k is *not* of the form $L([\!]\psi)$, then $T_{k+1} = T_k$.

Each T_k is consistent due to the last statement in the previous Lemma. Then $T = \bigcup_{k \in \omega} T_k$ is consistent. T is trivially closed under modus ponens. For any formula ϕ_k , either $\neg\phi_k$ was already in the k -th step of the construction, or ϕ_k was added to T_{k+1} ; therefore T cannot have proper consistent supersets closed under modus ponens. Finally suppose $L([\beta]\psi) \in T$ for all β . If $L([\!]\psi) \notin T$, then $\neg L([\!]\psi) \in T$ and thus $\neg L([\!]\psi) \in T_k$ for some k . Let $m > k$ such that $\phi_m = L([\!]\psi)$. By construction there exists a β such that $\neg L([\beta]\psi) \in T_{m+1} \subseteq T$: contradiction. Therefore T is a maximal consistent theory. ■

Now we define a relation between maximal consistent theories as: $TR_a S$ iff, for all ϕ , $B_a\phi \in T$ implies $\phi \in S$ (equivalently, iff $T_{B_a} \subseteq S$).

Proposition 8.15 (Diamond Lemma). *Suppose $\hat{B}_a\phi \in T$. Then there exists a maximal consistent theory S such that $TR_a S$ and $\phi \in S$.*

Proof. Consider the theory $(T_{B_a})_\phi$. First, note that T_{B_a} is a consistent theory, because $\vdash \hat{B}_a\phi \rightarrow \neg B_a\perp$, so $B_a\perp \notin T$ and thus $\perp \notin T_{B_a}$. Moreover, $B_a\neg\phi \notin T$, thus $\neg\phi \notin T_{B_a}$. By Lemma 8.13, we thus have that $T_{B_a} \subseteq (T_{B_a})_\phi$, $\phi \in (T_{B_a})_\phi$ and $(T_{B_a})_\phi$ is consistent. It then suffices to extend $(T_{B_a})_\phi$ by Lindenbaum's lemma to the desired successor. ■

Now we can define our canonical model: let W be the family of maximal consistent theories, let R_a be defined as above and let $V(p) = \{T \in W : p \in T\}$. We have:

Proposition 8.16 (Truth Lemma). *For any history α and formula ϕ , we have: $T, \epsilon \models \langle \alpha \rangle \phi$ iff $\langle \alpha \rangle \phi \in T$.*

Proof. By induction on (α, ϕ) .

The case $(\alpha, \phi) = (\epsilon, \top)$ is trivial. The cases $(\alpha, \phi) = (\alpha'.\psi, \top)$ and $(\alpha'.a, \top)$ follow from the axioms $R_{\top 1}$, $R_{\top 2}$ and $R_{\top 3}$ and the fact that $(\alpha', \psi) \ll (\alpha'.\psi, \top)$ and $(\alpha', \top) \ll (\alpha'.a, \top)$.

The case (α, p) follows from the definition of $V(p)$ and axiom R_p combined with the fact that $(\alpha, \top) \ll (\alpha, p)$.

The cases $(\alpha, \neg\psi)$ and $(\alpha, \psi_1 \vee \psi_2)$ follow from R_{\neg} and R_{\vee} , respectively, plus the fact that $(\alpha, \psi_i) \ll (\alpha, \psi_1 \vee \psi_2)$ (for the first case), and $(\alpha, \psi) \ll (\alpha, \neg\psi)$ (for the second case).

Let us see the case $(\alpha, \hat{B}_a\phi)$: if $T, \epsilon \models \langle \alpha \rangle \hat{B}_a\phi$, then on the one hand we have that $T \bowtie \alpha$ (i.e., $T, \epsilon \models \langle \alpha \rangle \top$, which by induction hypothesis paired with the fact that $(\alpha, \top) \ll (\alpha, \hat{B}_a\psi)$ gives us that $\langle \alpha \rangle \top \in T$), and on the

other hand, $S, \beta \models \phi$ by some S, β such that R_aTS , $\alpha \triangleright_a \beta$ and $S \bowtie \beta$. This means that $S, \epsilon \models \langle \beta \rangle \psi$ and thus (by induction hypothesis due to the fact that $(\beta, \psi) \ll (\alpha, \hat{B}_a \psi)$), we have that $\langle \beta \rangle \psi \in S$. This entails that $\hat{B}_a \langle \beta \rangle \psi \in T$ and thus $\langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \psi \in T$, which by R_B gives $\langle \alpha \rangle \hat{B}_a \psi \in T$. For the converse, we use R_B and the Diamond Lemma.

The cases $(\alpha, \langle a \rangle \psi)$ and $(\alpha, \langle \theta \rangle \psi)$ follow directly from the fact that $(\alpha.x, \psi) \ll (\alpha, \langle x \rangle \psi)$ for $x \in \mathcal{L} \cup A$.

Let us see the case $(\alpha, [!]\psi)$. If $T, \epsilon \models \langle \alpha \rangle [!]\psi$, then $T \bowtie \alpha$ and $T, \alpha \models [!]\psi$, which means that, for all $\beta \in (\mathcal{L}_{AA} \cup A)^*$, $T, \epsilon \models \langle \alpha \rangle [\beta]\psi$. By induction hypothesis, noting that $(\alpha, [\beta]\psi) \ll (\alpha, [!]\psi)$ whenever β does not contain occurrences of $[!]$, we have that $\langle \alpha \rangle [\beta]\psi \in T$ for all β and thus $\langle \alpha \rangle [!]\psi \in T$. Conversely, if $\langle \alpha \rangle [!]\psi \in T$, then $\langle \alpha \rangle \top \in T$ (and thus, by IH, $T, \epsilon \models \langle \alpha \rangle \psi$, which means $T \bowtie \alpha$), and, for all $\beta \in (\mathcal{L}_{AA} \cup A)^*$, $\langle \alpha \rangle [\beta]\psi \in T$, which again by induction hypothesis gives $T, \epsilon \models \langle \alpha \rangle [\beta]\psi$ for all β and thus $T, \alpha \models [!]\psi$, whence $T, \epsilon \models \langle \alpha \rangle [!]\psi$. ■

We will say that a formula ϕ is *consistent* if $\not\vdash \neg\phi$ and that a set of formulas Γ is *consistent* if it can be extended to a consistent theory. Note that ϕ is consistent if and only if the singleton set $\{\phi\}$ is consistent (for if $\neg\phi \notin \text{AAA}$, we can extend $\{\phi\}$ to the consistent theory AAA_ϕ).

We have:

Theorem 8.17. *AAA is strongly complete with respect to Kripke models.*

Proof. Let Γ be a consistent set of formulas. Then there exists a consistent theory $T_0 \supseteq \Gamma$ and, by Lindenbaum's lemma, a maximal consistent theory $T \supseteq T_0$. We construct the canonical model as above and we have that $T, \epsilon \models \phi$ for all $\phi \in \Gamma$. ■

8.5 Asynchronous Action Models

This final section shortly presents two logics for asynchronous reception of partially observed actions, including quantification over such actions. These logics contrast in interesting ways with the logic AA and with the logic AAA, the main subject of this chapter.

8.5.1 Asynchronous Action Model Logic

Action Model Logic was proposed by Baltag, Moss and Solecki in [17]. An *action model* is much like a relational model, but the elements of the domain are called *actions* instead of states, and instead of a valuation a *precondition*

is assigned to each domain element. A public announcement corresponds to a singleton action model where the precondition is the formula of the announcement. Under synchronous conditions, executing an action model on a Kripke model means constructing what is known as the *restricted modal product*. This product encodes the new state of information, after action execution. Under asynchronous conditions we do not construct the product model but calculate the belief consequences of actions from the histories, just as for the particular singleton action model that is the public announcement we do not construct model restrictions in AA but instead use the history.

An asynchronous non-public action is partially observed by the agents, just as in Action Model Logic, but it is unclear when the different agents partially observe the action, just as in AA. An example of an asynchronous partially observed action could consist for instance of two agents, Agripina and Benjamín, who are both ignorant about p , and are informed that Agripina will receive the truth about some proposition p but not Benjamín. Suppose that Agripina is going to receive the information that p (is true). By the time Benjamín learns that Agripina will be informed in this way, he considers it possible that Agripina has already been informed, in which case she now believes p or believes $\neg p$, but he also considers it possible that she has not yet been informed and thus remains ignorant about p . Dually, by the time Agripina learns that p but Benjamín has not yet learnt that Agripina will be informed about p , Benjamín incorrectly believes that Agripina is ignorant about p .

Action model. Formally, an *action model* $\mathcal{E} = (E, S, \text{pre})$ consists of a finite domain E of actions e, f, \dots , a family of *accessibility relations* $S = \{S_a : a \in A\}$, where each $S_a \subseteq E^2$, and a *precondition function* $\text{pre} : E \rightarrow \mathcal{L}$, where \mathcal{L} is a logical language. A pointed action model is a pair (\mathcal{E}, e) where $e \in E$, for which we write \mathcal{E}_e . We abuse the nomenclature and also call a pointed action model an *action*.

Syntax. Similarly to AA we can conceive a modal logical language with $\langle \mathcal{E}_e \rangle \phi$ as an inductive language construct, for action models \mathcal{E} with *finite domains*. The class of finite pointed action models is called \mathcal{AM} .

Histories are words in $(\mathcal{AM} \cup A)^*$. Much like in AA, we will use $\alpha \upharpoonright_!$ to refer to the projection of α to \mathcal{AM} and use $\alpha \upharpoonright_{!a}$, $|\alpha|_!$, $|\alpha|_a$ as usual.

View relation. The definition of the \triangleright_a relation in this setting incorporates the partial observability of action models: given $\alpha \triangleright_a \beta$, we demand that the *action models* appearing in α and β are the same. However, for agent a the

actions in α (points of these action models) may be different from the actions in β . That is, $\alpha \triangleright_a \beta$ iff $|\alpha|_a = |\beta|_a = |\beta|!$, and for all $i \leq |\alpha|_a$, if \mathcal{E}_e is the i -th action of α and \mathcal{F}_f is the i th action of β , then $\mathcal{E} = \mathcal{F}$ and $S_a e f$. This relation \triangleright_a is post-reflexive, transitive and post-symmetric if we are dealing with $\mathcal{S}5$ action models (in which all accessibility relations S_a are equivalence relations). Note that, given a history α , the set $\{\beta : \alpha \triangleright_a \beta\}$ is finite.

Semantics. We define an executibility relation \bowtie as follows:

- $w \bowtie \epsilon$,
- $w \bowtie \alpha.a$ iff $w \bowtie \alpha$,
- $w \bowtie \alpha.\mathcal{E}_e$ iff $w \bowtie \alpha$ and $w, \alpha \models \text{pre}(e)$.

With this, the semantics for belief and action model execution are what one might expect, namely:

$$w, \alpha \models \langle \mathcal{E}_e \rangle \phi \quad \text{iff } w, \alpha \models \text{pre}(e) \text{ and } w, \alpha.\mathcal{E}_e \models \phi.$$

$$w, \alpha \models \hat{B}_a \phi \quad \text{iff } t, \beta \models \phi \text{ for some } (t, \beta) \text{ such that } t \bowtie \beta, R_a w t, \text{ and } \alpha \triangleright_a \beta.$$

We call this *Asynchronous Action Model Logic*, AAM.

Reduction axioms and axiomatisation. We recall that the axiomatisation AAA presented in Section 8.4 consists of the rules and axioms of AA plus an axiom and a rule dedicated to the quantifier (Remark 8.12).

It is straightforward to see that the axiomatisation of AAM is as the axiomatisation of AA where only axiom $R_{\top 3}$ needs to be (analogously) reformulated for action models, whereas the axiom R_B is the same in AA and in AAM, except that, clearly, the relation \triangleright_a used in that axiom now refers to the much more involved view relation for partial observability defined above, where an agent considers all actions possible that are accessible for her given the actual action. These two relevant axioms are:

$$\begin{aligned} (R'_{\top 3}) \quad & \langle \alpha.\mathcal{E}_e \rangle \top \leftrightarrow \langle \alpha \rangle \text{pre}(e); \\ (R'_B) \quad & \langle \alpha \rangle \hat{B}_a \phi \leftrightarrow (\langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \phi). \end{aligned}$$

Just as for AA we can show that this axiomatisation is complete with respect to the class of models with empty histories, and that this is again a reduction system, such that every formula in the logical language is equivalent to a formula without dynamic modalities $\langle \mathcal{E}_e \rangle$ for action execution and $\langle a \rangle$ for receiving that information.

To prove that this system is a complete axiomatisation of AAM, we need to define a total preorder \ll from a complexity measure $|\cdot|$ which takes into consideration the precondition formulas present in action models \mathcal{E}_e . It therefore seems that this demands that

$$\begin{aligned} |(\mathcal{E}, e)| &= \sum_{e' \in E} |\text{pre}(e')| \\ |\alpha| &= \sum \{ |(\mathcal{E}, e)| : (\mathcal{E}, e) \text{ occurs in } \alpha \}. \end{aligned}$$

We wish to investigate this further and show completeness.

8.5.2 Arbitrary Asynchronous Action Model Logic

A further generalisation is the extension of the logical language with a quantifier $\langle \otimes \rangle$ over action models, such that $\langle \otimes \rangle \phi$ means that ϕ is true after the execution of some finite action model in the current (s, α) pair of the given model.

Let $\mathcal{AM}_{-\otimes}$ be the class of finite pointed action models where $\langle \otimes \rangle$ does not occur in the preconditions. We then get that

$$w, \alpha \models \langle \otimes \rangle \phi \text{ iff there exists } \beta \in (\mathcal{AM}_{-\otimes} \cup A)^* \text{ such that } w, \alpha \models \langle \beta \rangle \phi.$$

Let us call the logic with this quantifier AAAM (an extra A, for Arbitrary). Although work on this logic is also very much work in progress, it is illuminating to compare this extension AAAM of AAM with the logic AAA of Section 8.1, wherein we quantify over histories containing announcements. For the synchronous version of arbitrary action model logic, Hales showed in [56] that the restriction to quantifier-free precondition formulas in action models can be relaxed, and that we can prove the property (not the definition) of this logic that

$$w, \alpha \models \langle \otimes \rangle \phi \text{ iff there exists } \beta \in (\mathcal{AM} \cup A)^* \text{ such that } w, \alpha \models \langle \beta \rangle \phi.$$

He also showed that, given ϕ , we can *synthesize* a multi-pointed action model \mathcal{E}_F (where $F \subseteq \mathcal{D}(\mathcal{E})$) from ϕ such that $\langle \otimes \rangle \phi$ is equivalent to $\langle \mathcal{E}_F \rangle \phi$.

If it were possible to prove similar results for the logic AAAM of arbitrary asynchronous action models, that would be of great interest, as this would then show that AAAM is as expressive as AAM (without quantification), by reducing every formula to one without quantifiers, unlike the larger expressivity of quantifying over asynchronous announcements in AAA; and it would also show decidability of AAAM. Even independent from that, synthesis of asynchronous partially observable actions, and the complexity of such tasks, seems of interest to investigate further.

Discussion. This chapter presented the logic AAA of arbitrary asynchronous announcements, which can be used to reason about agents receiving and sending each other information under asynchronous conditions. Some properties of this framework with an arbitrary announcement quantifier were investigated, among which are bisimulation invariance, the larger expressivity of the logical language with the quantifier, and preservation after history extension of the fragment of the positive formulas. A complete infinitary axiomatisation was provided. A further generalisation to quantification over action models was tentatively described.

Some directions for future work have been outlined in the main text of this chapter. These include an axiomatisation of AAM, a thorough study of the synthesis problem for AAAM and, of course, a study of the logic of *-validities for the different frameworks presented here.

Some Dynamic Extensions of Social Epistemic Logic

We propose dynamic extensions to the framework of Social Epistemic Logic [92] discussed in Chapter 5.

One of them introduces a notion of a semi-public announcement made by an aware agent and only transmitted to this agent’s social connections. Another extension along the epistemic dimension of the framework contains separate ‘sending’ and ‘reading’ modalities that allow for a more realistic *asynchronous* spreading of messages in a social network.

Finally, we briefly discuss an extension along the ‘social’ dimension of the framework which introduces an operator to break links within an epistemic social network based on information at the agent’s disposal.

Completeness and decidability results are provided.¹



— My apologies, I genuinely thought you were dead.

— Your assumption is not wrong, it’s just early.

Twitter conversation between users @samatlounge and @CormacMcCrthy (the latter successfully posing as consecrated author Cormac McCarthy), 1st of August 2021.

THIS CHAPTER AIMS TO STUDY a fundamental aspect of social networks (such as Facebook or Twitter) – namely, that of *change*. The change users effect on social networks, be it in the way of posting or establishing links with other users (in the form of ‘friendships’ or ‘follows’), is the reason users keep coming back to them. Conversely, these posts can change the users’ epistemic state, their knowledge and beliefs, oftentimes to headline-worthy extents.

¹This chapter is based on the paper

Saúl Fernández González. Change in Social Networks: Some Dynamic Extensions of Social Epistemic Logic. *Journal of Logic and Computation*, 2022.

To do so, we shall present proposals for modelling some quotidian notions of change which occur in social networks, using as a basis the framework introduced by Seligman, Liu and Girard in *Logic in the Community* [92], which was the object of study of Chapter 5. We recall it briefly in Section 9.1.

Social Epistemic Logic (SEL) is a powerful framework which naturally lends itself to dynamic extensions. Some prescriptive notions of ‘public announcement’ are defined in the original paper [92], and various ideas of ‘epistemic update’ are explored in later literature [94, 104, 105, 89]. In Section 9.2 we start off by proposing a type of epistemic update which is perhaps more realistic with regards to the workings of information flow in social networks such as Facebook or Twitter: a message is sent, semi-publicly and knowingly, by an agent and received exclusively by this agent’s ‘friends’, who then update their epistemic state accordingly.

We continue in Subsection 9.2.4 with a dynamic exploration of the other dimension of this bidimensional framework: the rather uncharted (with the exception of [94]) territory of ‘social updates’. Relying on Arrow Update Logics [68], we propose the modelling of a situation in which an agent decides to ‘unfollow’ another based on the information at her disposal.

By using the tools of Public Announcement Logic [86] to model information flow in this setting, one makes the arguably undesired assumption that messages are received by all agents at the same time they are sent. This does not quite reflect posting a social network such as Twitter, in which a message can be sent by an agent (‘tweeted to her followers’), and not read by the recipients until they check their Twitter timeline hours or days later. Section 9.3, the most substantial section, addresses this assumption. Combining the frameworks of SEL and AA, we propose separate ‘sending’ and ‘receiving’ modalities, we explain how an agent may reason epistemically based on a sequence of announcements and readings, and we provide a sound and complete logic for this extension of the framework.

9.1 Social Epistemic Logic

Recall from Chapter 5 that the logic of SEL models (imposing no constraints on the relations) consists of the axioms and rules depicted in Table 9.1.

Recall moreover that, if one wants to consider the class of models in which the \sim_a are equivalence relations, one needs to add to SEL the S5 axioms for the K modality; if one wants the \simeq_w relations to be symmetric and reflexive (as is the case, e.g., for Facebook friendships), one needs to add

(Taut)	all propositional tautologies	(MP)	from ϕ and $\phi \rightarrow \psi$, infer ψ
(K _K)	$K(\phi \rightarrow \psi) \rightarrow (K\phi \rightarrow K\psi)$	(Nec _K)	from ϕ , infer $K\phi$
(K _F)	$F(\phi \rightarrow \psi) \rightarrow (F\phi \rightarrow F\psi)$	(Nec _F)	from ϕ , infer $F\phi$
(K _@)	$@_n(\phi \rightarrow \psi) \rightarrow (@_n\phi \rightarrow @_n\psi)$	(Nec _@)	from ϕ , infer $@_n\phi$
(Ref)	$@_nn$	(Selfdual)	$\neg @_n\phi \leftrightarrow @_n\neg\phi$
(Elim)	$@_n\phi \rightarrow (n \rightarrow \phi)$	(Agree)	$@_n @_m\phi \rightarrow @_m\phi$
(Back)	$@_n\phi \rightarrow F @_n\phi$	(DCom)	$@_n K @_n\phi \leftrightarrow @_n K\phi$
(Rigid ₌)	$@_nm \rightarrow K @_nm$	(Rigid _≠)	$\neg @_nm \rightarrow K \neg @_nm$
(Name)	From $@_n\phi$ infer ϕ , where n is fresh in ϕ		
(LBG)	From $L(@_n\hat{F}m \rightarrow @_m\phi)$ infer $L(@_nF\phi)$, m fresh in $L(@_nF\phi)$, with $L ::= \# \phi \rightarrow L @_nKL$.		

Table 9.1: The axioms and rules of ‘Basic SEL’

$$\text{(refl)} \quad @_n\hat{F}n, \quad \text{(sym)} \quad @_n\hat{F}m \rightarrow @_m\hat{F}n;$$

and if one wants the logic of “k-y-f” frames in which every agent knows who her own friends are (i.e. if $w \sim_a v$, then $a \succ_w b$ iff $a \succ_v b$), one needs to also add the axiom

$$\text{(kyf)} \quad \hat{F}m \rightarrow K\hat{F}m,$$

Liang [105] shows that the (LBG) inference rule can be removed in favour of the simpler rule:

$$\text{(@Name)} \quad \text{From } @_n\phi \rightarrow m_1 \vee \dots \vee m_k, \text{ infer } \neg @_n\phi,$$

where m_1, \dots, m_k are variables not appearing in ϕ .

Definition 9.1. Unless stated otherwise, in the remainder of this chapter we will just use *models* to refer to named, rigid, k-y-f models wherein the epistemic relations \sim_a are equivalence relations and the social relations \succ_w are symmetric and reflexive.

From now on, we use SEL to refer to the logic of these models, i.e., the axioms and rules in Table 9.1 plus the S5 axioms for K plus (kyf), (sym) and (refl).

9.2 Updates

In this section we shall introduce two notions of ‘update’ in the SEL framework, one of them affecting the epistemic dimension (in the form of a private announcement à la [18]), and one affecting the social dimension.

Since the choice for how to encode these message transmissions relies heavily on the type of situation one is set to model, let us establish some conceptual rules here:

First. We shall be concerned with ‘social networks’ not in a broad sense (agents linked by some sort of social connection), but in the contemporary, online sense of the term – i.e., we will be talking about such things as Facebook or Twitter, in which agents are capable of sending messages and the ‘social’ relations determine who receives these messages.

We thus remain closer to the intuitions behind [92, 94, 104] than those of Liang [105], which presents a more conceptually ambitious framework wherein the social relation corresponds to some sort of ‘social visibility’. Some dynamic notions present in [105], such as that of ‘unaware announcement’, will be of no use in our setting.

Second. The social relations \succsim_w may be symmetric (like a ‘Facebook friendship’) or asymmetric (like a ‘Twitter follow’); most results presented here are true for both types of relations unless stated otherwise. Whether this relation is reflexive or not admits some argument (one is not technically Facebook friends with oneself, but has access to all of one’s own posts) but it is formally irrelevant for our purposes. We keep them reflexive for simplicity.

Third. Knowledge should follow at the very least the axioms of S4: it should be factual and positively introspective.

Fourth. Epistemic change occurs only as a consequence of a post (e.g. a tweet) in one of these networks. We thus think of a post as a sort of ‘public announcement’, made by an agent who is aware of its contents before broadcasting it, who (perhaps controversially) believes it to be true, and should only be received by this agent’s social connections.

We are not using ‘public announcement’ here strictly in the sense of PAL [86], for these announcements have a particular sender and are only received by a subset of the agents as opposed to all of them. We will take care of the second fact when defining the post-announcement update, which will increase the size of the model to add uncertainty for the agents who have not received the message, while shrinking the set of accessible worlds for those who have. The fact that a message is sent by a particular knowing agent and not some external entity as is the standard in Public Announcement Logic is not particularly problematic: for instance, in [42, Example 4.69] it is argued that an adequate way to model the notion “agent a says ϕ ” in PAL is via the announcement “[$K_a\phi!$]”. We shall do something similar in this section.

Fifth. A ‘social update’ corresponds to a change in the composition of the social network, more specifically in the way agents are linked (a unfollows b on Twitter, c and d become Facebook friends, etc.).

We make some undesirable assumptions for the sake of simplicity. Among these is the fact that a message is received simultaneously by all agents (and only those agents) to which it is sent, which does not mirror the workings of a social network. In the next section we give a proposal for an asynchronous network; for the time being, we will sidestep this issue and simply not consider notions such as common knowledge, which is arguably not very useful in a non-synchronous setting.

Recall from Chapter 5 that a *BCN formula* is a Boolean combination of *named formulas* (i.e. formulas of the shape $@_n\phi$), and that the fact that $@_n$ distributes over all the Booleans plus the (Agree) axiom of *SEL* give us Corollary 5.4:

If ϕ is a BCN formula and $n \in \text{Nom}$, then $\vdash \phi \leftrightarrow @_n\phi$.

By this fact, plus the semantics of $@_n$, we obtain:

Lemma 9.2. *If ϕ is a BCN and $a, b \in A$, then $M, w, a \models \phi$ iff $M, w, b \models \phi$.*

For BCN formulas we will use $M, w, * \models \phi$ to abbreviate the fact that $M, w, a \models \phi$ for any $a \in A$.

9.2.1 Epistemic updates in the literature.

In [92], two notions for public announcement are given: one is described as ‘agent n announces ϕ ’ and defined as follows:

$$M, w, a \models [n!\phi]\psi \text{ iff } M, w, \underline{n} \models \phi \text{ implies } M_{\underline{n}, \phi}, w, a \models \psi,$$

where $M_{\underline{n}, \phi}$ is the restriction of M to the worlds $\{v : M, v, * \models @_n\phi\}$. The other is an indexical notion, intended to be read as: ‘after I announce ϕ to my friends, ψ holds’, and defined as:

$$M, w, a \models [F!\phi]\psi \text{ iff } M, w, a \models \phi \text{ implies } M', w, a \models \psi,$$

where M' is a restriction of the model M in which only links between ϕ worlds are kept for friends of a at w , i.e., for all $u, v \in W$ and $b \in A$: $u \sim'_b v$ iff $u \sim_b v$ and either $a \not\prec_w b$ or ($a \prec_w b$ and $M, u, a \models \phi$ and $M, v, a \models \phi$).

Neither of these notions seem to capture well the idea of sending a message by an agent. The former is simply a public announcement to every agent in the network of the fact that ϕ is true for agent n (and it may well be that n herself didn’t know this fact before the announcement), whereas the latter, combined with irreflexivity of the friendship relations in [92], makes it possible




Figure 9.1: Example of a situation wherein $(w, a) \models [F!p]@_a \neg Kp$ according to the semantics of the ‘Announcing to my friends’ operator: a , who does not know p , announces p to her only friend b , after which she still does not know p .

for an agent to remain ignorant that p is true even after having announced p to all her friends.

[105] makes a few proposals for a dynamic framework which are more careful with these considerations; notably, through the use of action model-like structures, one can distinguish notions of *aware* and *unaware* announcements. The awareness considered here is post-hoc, i.e., the agent might not know p *before* an aware announcement of p (e.g. ‘after the wind blew away my hat, everybody, including myself, was aware that I had forgotten to wear my wig’). As noted earlier, [105] employs this framework to model a logic of ‘social visibility’.

Some frameworks which iterate on SEL and deal with epistemic updates do not do so necessarily through PAL-like announcements. These include the Boolean announcements of [104], the PDL-like programs in [94] which effect, among other things, the epistemic state of agents, the AGM[2]-like approach to belief revision taken in [72, 73], or the epistemic updates through action model structures present in [105, 89].

None of these proposals quite adjust to the requirements of our fourth rule above. In the following subsection we propose an announcement that can only be made by an aware agent and will only be received by that agent’s friends. For reasons that will soon become apparent, we drop the requirement for negative introspection, we weaken our axioms of knowledge to **S4** and thus we shall consider (just for the remainder of this section) relations \sim_a which are reflexive and transitive instead of equivalence relations \sim_a .

9.2.2 SEL_!

We add a dynamic operator to the language, $[n!\phi]\psi$, which can be read: ‘after agent n announces ϕ to all of her friends, ψ holds’.

Definition 9.3. Given a model $M = (W, A, \rightsquigarrow, \succ, V)$, a nominal n and a formula ϕ , the *update of M with the announcement $n!\phi$* is the model


$$M^{n!\phi} = (W', A, \rightsquigarrow', \succ', V'),$$

where $W' = \{(w, 1) : M, w, * \models @_n K\phi\} \cup \{(w, 2) : w \in W\}$,

$$(w, i) \rightsquigarrow'_a (v, j) \text{ iff } w \rightsquigarrow_a v \text{ and } \begin{cases} \text{either } i = j, \\ \text{or } i = 1, j = 2 \ \& \ a \not\prec_w \underline{n}, \end{cases}$$

and $\succ'_{(w,i)} = \succ_w$, and $((w, i), a) \in V'(p)$ iff $(w, a) \in V(p)$.

What this update does is create a whole copy of the model which is only accessible for agents who are not friends of \underline{n} , i.e., agents who have not received the announcement ϕ ; agents who have received it can only access worlds in which ϕ is announceable (i.e. worlds in which agent n knows ϕ). This combines the PAL-like ‘model shrinking’ intuitions of [92] and the action-model like updates of [105] into an update which is more suited to our postulates above. Readers familiar with Action Model Logic might see a certain similarity between this and an update of a multi-agent epistemic model via the action



which represents sending a message ϕ to only a subset G of the agents (see e.g. [42, Ch. 6.9]), with the difference that our models remain reflexive after an update.

Note that the submodel of $M^{n!\phi}$ induced by the set $W \times \{2\}$ is a generated submodel of $M^{n!\phi}$ isomorphic to M , and thus:

Lemma 9.4. $M, w, a \models \psi$ iff $M^{n!\phi}, (w, 2), a \models \psi$.

Definition 9.5. The modal operator $[n!\phi]\psi$ is intended to be read as ‘after agent n announces ϕ to her friends, ψ holds’, and it is interpreted on models as follows:

$$M, w, a \models [n!\phi]\psi \text{ iff } M, w, \underline{n} \models K\phi \text{ implies } M^{n!\phi}, (w, 1), a \models \psi.$$

We call this extended language $\mathcal{L}_{SEL!}$.

Note that an update by an agent n does not change the epistemic state of agents who are not friends of n , in the sense that they consider possible whatever they considered possible before the update. Indeed, the following is valid:

$$\neg \hat{F}n \rightarrow (\hat{K}\psi \rightarrow [n!\phi]\hat{K}\psi).$$

This is because, if $a \not\prec_w \underline{n}$ and before the update a could access a world where ψ was true ($w \rightsquigarrow_a v$ with $v, a \models \psi$) then after the update a can access the world $(v, 2)$ from $(w, 1)$ and (due to Lemma 9.4) $M^{n!\phi}, (v, 2), a \models \psi$ and thus $M^{n!\phi}, (w, 1), a \models \hat{K}\psi$.

However, due to the factive nature of knowledge it does not hold that agents know whatever they knew before the update, for an agent could know a true proposition ψ which has, unbeknownst to her, become false after a message was sent to other agents. This is why we drop negative introspection: we do not want the agent to know she has suddenly stopped knowing ψ .

9.2.3 Axiomatisation of SEL₁

We use SEL₅₄ to refer to the logic SEL minus the 5 axiom $\hat{K}\phi \rightarrow K\hat{K}\phi$.

Proposition 9.6. *The following reductions are valid in SEL₁:*

- (R!_⊤) $[n!\phi]\top \leftrightarrow \top$
- (R!_p) $[n!\phi]p \leftrightarrow (@_n K\phi \rightarrow p)$
- (R!_¬) $[n!\phi]\neg\psi \leftrightarrow (@_n K\phi \rightarrow \neg[n!\phi]\psi)$
- (R!_∧) $[n!\phi](\psi_1 \wedge \psi_2) \leftrightarrow [n!\phi]\psi_1 \wedge [n!\phi]\psi_2$
- (R!_@) $[n!\phi]@_m\psi \leftrightarrow @_m[n!\phi]\psi$
- (R!_F) $[n!\phi]F\psi \leftrightarrow F[n!\phi]\psi$
- (R!_K) $[n!\phi]K\psi \leftrightarrow (@_n K\phi \rightarrow (K\psi \vee \hat{F}n) \wedge K[n!\phi]\psi)$
- (RE!) from $\phi \leftrightarrow \psi$, infer $[n!\theta]\phi \leftrightarrow [n!\theta]\psi$

Proof. We focus on the sixth line; the rest are rather straightforward and left to the reader. We show the validity of (R!_K) in its dual form, $\langle n!\phi \rangle \hat{K}\phi \leftrightarrow @_n K\phi \wedge ((\neg \hat{F}n \wedge \hat{K}\psi) \vee \hat{K}\langle n!\phi \rangle \psi)$.

If $M, w, a \models \langle n!\phi \rangle \hat{K}\phi$, then $M, w, a \models @_n K\phi$ and $M^{n!\phi}, (w, 1), a \models \hat{K}\psi$. Therefore there is a successor (v, j) such that $(w, 1) \rightsquigarrow'_a (v, j)$ and $M^{n!\phi}, (v, j), a \models \psi$. If $j = 2$, this means that $a \not\prec_w \underline{n}$ in the original model, (and therefore $M, w, a \models \neg \hat{F}n$), and (by Lemma 9.4) $M, v, a \models \psi$ and therefore $M, v, a \models \hat{K}\psi$. If $j = 1$, then by the above semantics we have that $M, v, a \models \langle n!\phi \rangle \psi$ and thus $M, w, a \models \hat{K}\langle n!\phi \rangle \psi$. The converse is analogous. ■

The preceding reductions can be used to construct, using a rather standard argument (see [42, 86] for details on this proof method), a translation τ from formulas of the extended language into announcement-free \mathcal{L}_{SEL} formulas in such a way that $\models \tau(\phi) \leftrightarrow \phi$. As a result,

Corollary 9.7. *SEL_{S4} plus the above reduction axioms $(R!_p) - (R!_K)$ plus the rule $(RE!)$ is the sound and complete logic of the extended language \mathcal{L}_{SEL} .*

9.2.4 Social updates

In the previous subsections announcement modalities $[n!\phi]$ were proposed. These announcements affect the epistemic state of the agents and provides them with new information. In other words, these announcements represent an epistemic update. New information changes what agents *know* but it does not change the agent's disposition towards one another. This notion of epistemic update, and those explored in [104, 105, 73], makes the framework of *SEL* dynamic along its epistemic dimension.

In this subsection we propose a way to make this framework dynamic along its other dimension as well: we shall add a notion of *social update* to *SEL*. In the same way that the $[n!\phi]$ updates made some worlds inaccessible to some agents by breaking epistemic links, these social updates will break friendships within a world.

This type of update is made to represent changes in opinion (e.g. ‘I wish to only remain friends with people who oppose fascism’, ‘I don’t want to follow my ex-husband on Instagram’, etc.). In order to encode this, we shall use the mechanisms of Arrow Update Logic of [68].

We shall drop, if only for this section, the assumption that the \succsim_w relation must be symmetric, for the ‘follow’ relation does not tend to be. We will be using the notation R_w for our social relations instead of \succsim_w .

Definition 9.8 (Arrow Updates in Epistemic Logic [68]). In a multi-agent epistemic logic, an *arrow update* is a finite set $U = \{(\phi_1, a_1, \psi_1), \dots, (\phi_n, a_n, \psi_n)\}$, where the ϕ_i and ψ_i 's are formulas and the a_i 's are agents, representing the fact that, after the update, agent a_i will only maintain epistemic links that go from worlds satisfying ϕ_i to worlds satisfying ψ_i .

For each such set, an update operation on multi-agent Kripke models is defined, $(W, \{R_a\}_{a \in \text{Agt}}, V) \otimes U = (W, \{R_a^U\}_{a \in \text{Agt}}, V)$, with

$$R_a^U = \{(u, v) \in R_a : \exists(\phi, a, \psi) \in U \text{ s.t. } M, u \models \phi \ \& \ M, v \models \psi\},$$

and a modal operator $[U]\phi$ is added to the language of SEL, which is read as follows:

$$M, w \models [U]\phi \text{ iff } M \otimes U, w \models \phi.$$

Thus, following [68], we shall add a modal operator $[U]\phi$ to our language. In our case, however, this arrow update will be used to model not the erasure of epistemic links between worlds, but the erasure of social relations within a world.

There are several options as to what arrow updates will represent in this setting. We present two of them, and the different types of updates which encode them, although we will focus on the first idea for the technical results in the remainder of this subsection.

An update U could be:

One. A finite set of the shape $U = \{(n_1, \phi_1), \dots, (n_k, \phi_k)\}$ where a pair (n_i, ϕ_i) represents that the agent denoted by n_i will only keep friendships, at each possible world, that she already had and which satisfy the formula ϕ_i . The semantics then would be:

$$M, w, a \models [U]\phi \text{ iff } M^U, w, a \models \phi,$$

where $M^U = (W, A, \sim, R^U, V)$, and

$$R_v^U ab \text{ iff } R_v ab \text{ and } \exists (n, \psi) \in U : \underline{n} = a \ \& \ M, v, b \models \psi.$$

Two. Perhaps more in line with [68], U could be a set of the form $\{(n_1, \psi_1, \phi_1), \dots, (n_k, \psi_k, \phi_k)\}$. Here (n_i, ψ_i, ϕ_i) means that, if agent \underline{n}_i satisfies ψ_i , then she will keep (or break) her friendship with any agent satisfying ϕ_i . Same semantics but:

$$\begin{aligned} R_w^U ab \text{ iff } R_w ab \ \& \ \exists (n, \psi, \phi) \in U \\ \text{such that } (\underline{n} = a \ \& \ M, w, a \models \psi \ \& \ M, w, b \models \phi). \end{aligned}$$

This could allow to characterise a situation such as: ‘only in the worlds where a satisfies $\neg p$ will she stop being friends with agents satisfying p ’.

Let us now see some reduction axioms.

Lemma 9.9. *The following are valid (for both definitions of U):*

- (RU_{\top}) $[U]\top \leftrightarrow \top$;
- (RU_p) $[U]p \leftrightarrow p$;
- (RU_{\neg}) $[U]\neg\phi \leftrightarrow \neg[U]\phi$;
- (RU_{\wedge}) $[U](\phi_1 \wedge \phi_2) \leftrightarrow ([U]\phi_1 \wedge [U]\phi_2)$;
- ($RU_{@}$) $[U]@_n\phi \leftrightarrow @_n[U]\phi$;
- (RU_K) $[U]K\phi \leftrightarrow K[U]\phi$.

Proof. Straightforward unpacking of the semantics; left to the reader. ■

The reduction rule pertaining to the F modality, however, depends on the definition of U we are considering. Let us from now on stick to the first one.

Lemma 9.10. *The following equivalence is sound:*

$$(RU_F) \quad [U]F\phi \leftrightarrow \bigwedge_{(n,\psi) \in U} (n \rightarrow F(\psi \rightarrow [U]\phi)).$$

Proof. Left to right: suppose $M, w, a \models [U]F\phi$. This means that $M^U, w, a \models F\phi$ or, equivalently, that $M^U, w, b \models \phi$ for all b such that $R_w^U ab$. Let $(n, \psi) \in U$. Now, if $M, w, a \models n$ (i.e., if $\underline{n} = a$) let us see that all successors of a satisfy $\psi \rightarrow [U]\phi$. Take b such that $R_w ab$. If $M, w, b \models \psi$, then by definition we have $R_w^U ab$, which entails $M^u, w, b \models \phi$, or, equivalently, $M, w, b \models [U]\phi$, as we wanted to prove.

Right to left: suppose $M, w, a \models \bigwedge_{(n,\psi) \in U} (n \rightarrow F(\psi \rightarrow [U]\phi))$ and take b such that $R_w^U ab$. This means that $R_w ab$ and for some (n, ψ) we have $\underline{n} = a$ and $M, w, b \models \psi$. But since $M, w, a \models n \rightarrow F(\psi \rightarrow [U]\phi)$, this means that $M, w, b \models [U]\phi$, and thus $M^U, w, b \models \phi$, which entails $M^U, w, a \models F\phi$ and thus $M, w, a \models [U]F\phi$. ■

By an argument via reduction, analogous to the one in the previous subsection, one easily arrives to the following result:

Corollary 9.11. *The sound and complete logic of social epistemic models with social updates consists of the axioms and rules of SEL plus all the above reduction axioms.* ■

Comparison to Facebook Logic. In [94], the basic language of SEL is enriched with PDL-like programs, allowing to express complex epistemic and social interactions in a SEL model. Multiple examples of such programs are provided, expressing things such as private communications, announcements about the sender/ receiver, and questions. This framework is called Facebook Logic, no axiomatisation for which is given.

On the social side of things, a program which removes the link among agents n and m is given: $cut_F(n, m)$, which, utilizing the PDL-like notation of [94], is built from the basic programs “ $\phi?$ ” and “ F ” as follows: $cut_F(n, m) = (\neg n?; F) \cup (F; \neg m?)$. In the present framework, the removal of a link from n to m simply corresponds to the arrow update $U = \{(n, \neg m)\}$.

While this simplifies things, it comes at a cost: the PDL tools of [94] allow for more complex social interactions such as ‘adding a friend’, ‘sending/ accepting/ rejecting a friend request’, etc., while the framework presented in this section is unable to capture these notions.

Social updates outside of SEL. Let us briefly allude to some literature which deals with notions of ‘social updates’ in networks outside the setting of SEL. [96] presents a framework wherein a strong difference of opinion among the agents on a certain issue can change their predisposition to one another and might eventually break their friendship. [84] models situations of social influence in which an agent can be persuaded by two social groups with opposing views and has to side with one of them, which may involve breaking social links with the other group. In the social network models of [95], friendships are created based on similarity of opinion.

Incorporating any of these ideas with Social Epistemic Logic would be a worthy endeavour.

9.3 Asynchronous reception of messages

The epistemic update defined in the previous section was designed to behave in certain ways like the sending and reception of a message in a social network such as Twitter or Facebook would. One way in which it does not, however, resides in the implicit assumption carried over from public announcement logics that the reception of these messages by the corresponding agents occurs at the same time they are sent. In a more realistic scenario, once a person publishes a tweet, they simply ‘queue’ it, so that when another agent checks their Twitter timeline, they get to see (‘at once’, let us assume) all the tweets sent by their friends in the lapse from the last time they checked Twitter.

In the present section we introduce a proposal to model such a scenario, employing the framework Asynchronous Announcement Logic, albeit substantially tweaking it for these purposes.

9.3.1 Announcements, readings and histories

Let us consider two dynamic modalities, one of them of the form $[n!\phi]$, indicating the sending of message ϕ by agent n to the queue, and one of the form $[n:r]$, indicating the receiving by agent n of all queued messages sent by her friends.

These modalities will be iterated, so that we will often find sequences of announcements and reading modalities such as

$$[n_1!\phi_1][n_2!\phi_2][n_1:r][n_3!\phi_3][n_2:r]\psi.$$

In these cases we omit the inner square brackets and represent this situation by


$$[n_1!\phi_1. n_2!\phi_2. n_1:r. n_3!\phi_3. n_2:r]\psi.$$

We call these finite sequences of message transmissions and receptions *histories*.

Note that, unlike in previous chapters, we do not impose any restrictions on sequences: any word $\alpha \in ((\text{Nom} \times \mathcal{L}_{\text{SEL}}) \cup (\text{Nom} \times \{r\}))^*$ is a history. As usual, the empty sequence will be denoted ϵ .

We shall evaluate formulas on models with respect to a triple consisting of a world, and agent, and the sequence of events which has been executed (see Def. 9.19 below).

Example 9.12. Recall the model in Figure 5.1, reproduced below.



Suppose n_a, n_b and $n_c \in \text{Nom}$ are names for the three agents (e.g. ‘Alice’, ‘Bob’ and ‘Charlie’). After Bob, at world w , tweets that he is grey (represented by p), the epistemic state of a and c , who have not yet read the tweet, does not change. Likewise for Charlie tweeting p . Note that initially Alice is uncertain about both Bob and Charlie being grey, because she cannot distinguish the world w in which they are and the world w' in which they are not. But after reading the queued tweets in her timeline at world w she receives Bob’s tweet stating that he is grey (seeing as Alice and Bob are friends at w), and she does not receive Charlie’s tweet (for they do not follow each other). The following holds:

$$\begin{aligned} w, a, (n_b!p. n_c!p) &\models \neg K@_{n_b}p \wedge \neg K@_{n_c}p \\ w, a, (n_b!p. n_c!p. n_a:r) &\models K@_{n_b}p \wedge \neg K@_{n_c}p \end{aligned}$$

Note that the sequence of events $\alpha = n_b!p. n_c!p. n_a:r$ can only occur at world w : indeed, we make the assumption that messages must be true (at least when they are sent), and thus agent b could not possibly tweet p at worlds w' and w'' . Similarly to Chapter 6, we represent the fact that α is *executable at w* (and not at w') by $w \bowtie \alpha$ (and $w' \not\bowtie \alpha$). We shall formalise this notion later.

In order to evaluate her own epistemic state, an agent needs to reason about the messages she has received. But it could be (as it is often) that the messages she has received are just a subset of the messages which have been sent. It might well be that agent does not consider the ‘actual history’ in order to form her knowledge.

Let us note as well that an agent lacks knowledge about which of the messages she has received have been received by other agents, as well as about potential unreceived messages. For instance, suppose agent a has received the announcement that p is true for agent c , but b does not know $@_c p$ yet; it can never be the case that agent a knows $@_b \neg K @_c p$: indeed, she may consider it possible that b has not received the announcement p from c , but she can not be certain that b has not received this announcement unbeknownst to her.

Given a history α , let us determine the set of histories that an agent will quantify over at a world in order to form her beliefs.

First, some definitions:

Definition 9.13. Given a history $\alpha \in ((\text{Nom} \times \mathcal{L}_{\text{SEL}}) \cup (\text{Nom} \times \{r\}))^*$, we define:

- i. $\alpha \upharpoonright_1$ as the projection of α to $\text{Nom} \times \mathcal{L}_{\text{SEL}}$ (i.e. the subsequence of α consisting only of announcements of the form ' $n!\phi$ ');
- ii. $\text{nom}(\alpha)$ as the set of nominal variables occurring in α (both as reading modalities, within announcement modalities $n!\phi$ and as subformulas of said ϕ).

Now, given a history α , a model $M = (W, A, \sim, \succ, V)$ and a pair $(w, a) \in W \times A$, we will define $\alpha \upharpoonright_{(w,a)}$ as the list of announcements read by agent a at world w ; we shall do this by splicing the history α into several sequences, using the readings of agent a as cut points; we shall then consider the subsequences of each one consisting only of the announcements made by friends of a .

As an example before introducing the formal definition, let

$$\alpha = n_1!p. k_1!q. n_1:r. m_2!\neg F(p \wedge q). n_2!p. n_1:r. k_1:r. m_1!q. m_2!\hat{F}p. n_2:r,$$

and suppose we have a model in which $W = \{w\}$, $A = \{a, b, c\}$, $\underline{n}_1 = \underline{n}_2 = a$, $\underline{k}_1 = b$, and $\underline{m}_1 = \underline{m}_2 = c$. Suppose that in the only world in this model a and b are the only pair of friends, i.e. $\succ_w = \{(a, b), (b, a)\}$. In this case, a receives messages three times (corresponding to the occurrences of the readings $n_1:r$ and $n_2:r$), and each time only receives messages sent, by herself or her friends, from the last time she checked (i.e., only those sent by n_1, n_2, k_1). We have the following situation:

$$\underline{n}_1!p. \underline{k}_1!q. \mathbf{n}_1:r. m_2!\neg F(p \wedge q). \underline{n}_2!p. \mathbf{n}_1:r. k_1:r. m_1!q. m_2!\hat{F}p. \mathbf{n}_2:r,$$

and thus $\alpha \upharpoonright_{(w,a)} = \langle (n_1!p. k_1!q), (n_2!p), \epsilon \rangle$.

Note that we do not need all the information provided to us by the model in order to define $\alpha \upharpoonright_{(w,a)}$; we only need to know, out of the finitely many nominal

variables occurring in α , which ones name the same agents in the model and which of these agents are friends at world w . Thus we can define a syntactic notion of $\alpha \upharpoonright_{(w,a)}$ having only the following information: (i) a specification of which nominal variables, out of a certain set of nominals, refer to the same agent (via a partition); (ii) a specification of which cell of this partition (if any) refers to the agent we are concerned with, and (iii) a specification (via a binary relation) of which of these agents are ‘friends’. We shall call this a *pseudomodel*:

Definition 9.14. Given a finite set of nominals $N \subset \text{Nom}$, a N -*pseudomodel* is a tuple (A, w, a) where A is a partition of N , $a \in A \cup \{\emptyset\}$, and $w \subseteq A^2$ is a reflexive binary relation. We use $[\cdot]$ to denote the equivalence classes on N induced by A .

Given a N -pseudomodel (A, w, a) , we define $\alpha \upharpoonright_{(w,a)}^A$ to be the sequence $\langle \alpha_1, \dots, \alpha_k \rangle$, where k is the number of readings $n:r$ occurring in α such that $n \in a$, and each α_i is the maximal sequence of announcements $n_1!\phi_1 \dots n_m!\phi_m$ between the $i-1$ th and the i th occurrence of these readings such that $(a, [n_l]) \in w$ for $1 \leq l \leq m$. For a recursive definition:

- $\epsilon \upharpoonright_{(w,a)}^A = \epsilon$.
- $(\alpha. m!\phi) \upharpoonright_{(w,a)}^A = \alpha \upharpoonright_{(w,a)}^A$ for all m ;
- if $m \notin a$, $(\alpha. m:r) \upharpoonright_{(w,a)}^A = \alpha \upharpoonright_{(w,a)}^A$;
- if $m \in a$, $(\alpha. m:r) \upharpoonright_{(w,a)}^A = \langle \beta \upharpoonright_{(w,a)}^A, (n_1!\phi_1 \dots n_k!\phi_k) \rangle$, where $\alpha = (\beta. n':r. \delta)$, $n' \in a$, δ is a history in which no reading actions $n:r$ such that $n \in a$ occur, and $(n_1!\phi_1 \dots n_k!\phi_k)$ is the maximal subsequence of announcements in δ such that $(a, [n_j]) \in w$ for all j .

Going from this syntactic definition to a semantic one is what one might expect:

Definition 9.15. Given a model $M = (W, A, \sim, \succ, V)$ and a pair (w, a) , we define the N -*pseudomodel induced by* (M, w, a) as follows: A_M is the partition induced on N by the equivalence relation $n \equiv m$ iff $\underline{n} = \underline{m}$, a_M is the equivalence class of any $n \in N$ such that $\underline{n} = a$ (or $a_M = \emptyset$ if there is no such n), and $([n], [m]) \in w_M$ iff $\underline{n} \succ_w \underline{m}$.

We set $\alpha \upharpoonright_{(w,a)} := \alpha \upharpoonright_{(w_M, a_M)}^{A_M}$, where (A_M, w_M, a_M) is the $\text{nom}(\alpha)$ -pseudomodel induced by (M, w, a) .

Note that $\alpha \upharpoonright_{(w,a)}^A$ is defined as a finite sequence of finite sequences of announcements. Given such a meta-sequence γ , we let $\sqcup\gamma$ be the concatenation of its sequences (e.g., in the above example $\sqcup\alpha \upharpoonright_{(w,a)} = n_1!p. k_1!q. n_2!p$).

We are now in position to define a binary relation $\triangleright_{(w,a)}$ on the set of histories in such a way that $\alpha \triangleright_{(w,a)} \beta$ if and only if, given an actual sequence of events α , β is one of the alternatives for agent a at world w .

Definition 9.16. Given histories α and β , and given an N -pseudomodel (A, w, a) with $N \supseteq \text{nom}(\alpha)$, we define $\alpha \triangleright_{(w,a)}^A \beta$ if and only if:

1. $\alpha \upharpoonright_{(w,a)}^A = \beta \upharpoonright_{(w,a)}^A$, and
2. $\sqcup\alpha \upharpoonright_{(w,a)}^A = \beta \upharpoonright_{(w,a)}^A$.

For a pair (w, a) in a model M , we define $\triangleright_{(w,a)} = \triangleright_{(w_M, a_M)}^{A_M}$. (Note that the set $\{\beta : \alpha \triangleright_{(w,a)} \beta\}$ is finite.)

In other words, β is a possible sequence of events for agent a at world w given α if and only if agent a has received the same announcements in the same order in both α and β , and furthermore *there are no other announcements in β* . This second requirement is related to the point raised above: an agent need not use the *actual history* to form her beliefs, but rather approximations of it based on the messages she has received. The relation $\triangleright_{(w,a)}$ is not reflexive; the ‘epistemic’ modality we will be constructing from it is thus not a knowledge modality. This is adequate for our purposes: in this setting we would not want an agent to *know* that another has not received a certain information (so for example $@_n(p \wedge K@_m \neg K@_n p)$ should be impossible), and in a knowledge-centered setting like PAL such truths are knowable.

There is a formal reason for this requirement as well: forcing the histories $\{\beta : \alpha \triangleright_{(w,a)} \beta\}$ to not contain any announcements a has not read helps us avoid circularity issues. If a had to consider all potential unread announcements in order to evaluate her knowledge, it could well be that, in order to check whether Kp is true at (w, a) with respect to history α , we would need to evaluate whether Kp is true at (w, a) with respect to α , and so on.

For these reasons, we give our framework a doxastic modality B instead of the epistemic K .

Let us now define our language and semantics:

9.3.2 Language and semantics of ASEL

Definition 9.17 (Language of ASEL). We define a language $\mathcal{L}_{\text{ASEL}}$ by:

$$\phi ::= p|n|\perp|\neg\phi|(\phi \wedge \phi)|F\phi|B\phi|[n!\phi]\phi|[n:r]\phi,$$

with $n \in \text{Nom}$ and $p \in \text{Prop}$. We use $\hat{F}, \hat{B}, \langle n!\phi \rangle, \langle n:r \rangle$ as the duals of the above operators.

As noted above, we will read formulas of $\mathcal{L}_{\text{ASEL}}$ in SEL models with respect to a tuple (w, a, α) , where w is a world, a is an agent and α is a history which is executable at w .

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 9.18. Given a formula ϕ , we define quantities $\text{deg}_B \phi$ and $\|\phi\|$ recursively:

$$\begin{array}{ll} \text{deg}_B p = \text{deg}_B n = 0 & \|p\| = \|n\| = 2 \\ \text{deg}_B \perp = 0 & \|\perp\| = 1 \\ \text{deg}_B (\neg\phi) = \text{deg}_B \phi & \|\neg\phi\| = \|\phi\| + 1 \\ \text{deg}_B (\phi \wedge \psi) = \max\{\text{deg}_B \phi, \text{deg}_B \psi\} & \|\phi \wedge \psi\| = \|\phi\| + \|\psi\| \\ \text{deg}_B ([n!\phi]\psi) = \text{deg}_B \phi + \text{deg}_B \psi & \|[n!\phi]\psi\| = 2\|\phi\| + 2 + \|\psi\| \\ \text{deg}_B ([n:r]\phi) = \text{deg}_B \phi & \|[n:r]\phi\| = \|\phi\| + 2 \\ \text{deg}_B (\Box\phi) = \text{deg}_B \phi + 1 & \|\Box\phi\| = \|\phi\| + 1, \end{array}$$

with $\Box = B, F$ or $@_n$ in the last line. We extend these to histories by setting $\text{deg}_B \alpha := \sum\{\text{deg} \psi : \psi \text{ announcement in } \alpha\}$ and

$$\|\epsilon\| = 0, \|\alpha. n:r\| = \|\alpha\| + 1, \|\alpha. n!\phi\| = \|\alpha\| + \|\phi\| + 1.$$

Finally, for pairs (α, ϕ) we set: $\text{deg}_B(\alpha, \phi) = \text{deg}_B \alpha + \text{deg}_B \phi$ 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

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

Definition 9.19 (Semantics of ASEL). Let (W, A, \sim, \succ, V) be a SEL model. We define by simultaneous \ll -recursion the notions $w \bowtie \alpha$ (“history α is executable at world w ”) and $w, a, \alpha \models \phi$ (“ ϕ holds at (w, a) after history α has been executed”), as follows:

$w \bowtie \epsilon$	always,
$w \bowtie \alpha. n:r$	iff $w \bowtie \alpha$,
$w \bowtie \alpha. n!\phi$	iff $w \bowtie \alpha$ and $w, \underline{n}, \alpha \models B\phi$,
$w, a, \alpha \models p$	iff $(w, a) \in V(p)$,
$w, a, \alpha \models n$	iff $a = \underline{n}$,
$w, a, \alpha \models \neg\phi$	iff $w, a, \alpha \not\models \phi$,
$w, a, \alpha \models \phi_1 \wedge \phi_2$	iff $w, a, \alpha \models \phi_i$, for $i = 1, 2$,
$w, a, \alpha \models [n:r]\phi$	iff $w, a, \alpha. n:r \models \phi$,
$w, a, \alpha \models [n!\psi]\phi$	iff $w, \underline{n}, \alpha \models B\psi$ implies $w, a, \alpha. n!\psi \models \phi$,
$w, a, \alpha \models @_n\phi$	iff $w, \underline{n}, \alpha \models \phi$
$w, a, \alpha \models F\phi$	iff for all $b \in A$, $a \succ_w b$ implies $w, b, \alpha \models \phi$,
$w, a, \alpha \models B\phi$	iff $v, a, \beta \models \phi$ for all (v, β) such that $w \sim_a v$, $\alpha \triangleright_{(w,a)} \beta$ and $v \bowtie \beta$.

Definition 9.20 (Validities). A formula ϕ is said to be ϵ -*valid* (or simply *valid*) if $w, a, \epsilon \models \phi$ for every model (W, A, \sim, \succ, V) and pair $(w, a) \in W \times A$.

ϕ is $*$ -*valid* (or *strongly valid*) if $w, a, \alpha \models \phi$ for every pair (w, a) in every model such that $w \bowtie \alpha$.

We shall now move on to determine the logic of ϵ -validities of this asynchronous framework. Some observations about the language and semantics are in place.

Let us formalize first the above intuition that histories represent sequences of announcement and reading operators. Given a history α and a formula ϕ , we define the abbreviation $[\alpha]\phi$ recursively as follows:

$$[\epsilon]\phi = \phi, [\alpha. n:r]\phi = [\alpha][n:r]\phi, [\alpha. n!\psi]\phi = [\alpha][n!\psi]\phi.$$

(We do so dually to define $\langle \alpha \rangle \phi$.) In particular, this means that every formula in the language is of the form $[\alpha]\phi$ for some α .

The following holds:

Lemma 9.21. $w, a, \epsilon \models \langle \alpha \rangle \phi$ if and only if $w \bowtie \alpha$ and $w, a, \alpha \models \phi$. (In particular, $w \bowtie \alpha$ iff $w, *, \epsilon \models \langle \alpha \rangle \top$.)

Proof. This result follows from unpacking the semantics. Details are left to the reader. ■

9.3.3 Some considerations and examples

Before moving on to axiomatise this logic, let us make some considerations and exemplify a few situations wherein the novelty of asynchronicity might help reach places that cannot be reached in a synchronous setting:

Some beliefs are factual. We construct our belief modality departing from an initial model that presents equivalence relations. This means that in an initial situation in which no messages have been sent or received beliefs are factual. Indeed, it holds that $B\phi \rightarrow \phi$ is ϵ -valid. It is not, however, $*$ -valid in general. The relation allowing us to evaluate the ‘belief’ modality can be thought of as a relation among triples (w, a, α) given by: $(w, a, \alpha)\mathbf{R}(v, b, \beta)$ iff $w \sim_a v$ and $\alpha \triangleright_{(w,a)} \beta$; this is a combination of an equivalence relation (\sim_a) and a serial, transitive and Euclidean relation ($\triangleright_{(w,a)}$) and it is thus serial, transitive and Euclidean.

There are however some formulas which can only be believed by an agent if they are true. The reader may check that $B\phi \rightarrow \phi$ is $*$ -valid if ϕ is, for instance, $@_n p$, or $@_m B\neg p$. In Section 6.2 it is shown that this is true for all formulas in the so-called *positive fragment*, $\phi ::= p | \neg p | \phi \wedge \phi | B_a \phi$. We conjecture that an analogous result might be true for this setting.

Agents may remain ignorant of messages already sent to them. An agent may remain ignorant of an information which has already been made public to her, and even announce her ignorance before receiving the message.

Example 9.22. All agents in this and the following examples are assumed to be friends at all worlds. Nora tweets that Maruxa has received a grant ($@_m p$, where p is the indexical proposition ‘_ got a grant’), after which Maruxa, who has not received this message, tweets that she does not believe she has received the grant. She then reads Nora’s tweet and happily announces that she has indeed received it. The sequence $n!@_m p. m!\neg B p. m:r. m!p$, which encodes this exchange of information, is perfectly possible in an ASEL model.

(Note that $@_m B @_m p$ is equivalent to $@_m B p$, so $m!p$ and $m!@_m p$ encode the same announcement.)

Messages can become false before reception. Messages in ASEL, much like in most announcement frameworks like PAL, need to be true when sent. However, due to the asynchronicity of ASEL, an agent might receive a message which was true at the time and has since become false. For instance:

Example 9.23. Nora tweets that Maruxa does not know she has received a grant ($@_m \neg B p$), followed by a clarification tweet that she has indeed received it. Maruxa then reads both tweets, after which Kilian does. This is represented by the sequence $n!@_m \neg B p. n!@_m p. m:r. k:r$; by the time Kilian receives the message that Maruxa does not believe $@_m p$ holds, it has already become false. Kilian, unaware Maruxa has read the tweet, might still believe that Maruxa

does not know the information, although he necessarily considers it possible that she has.

Moore sentences are not immediately made false. In this asynchronous setting we could have a situation wherein a Moore sentence pertaining one of the receivers of a message is sent, and remains true for a while, even after being received by some agents in the network.

Example 9.24. In the above situation, Nora tweets at once that Maruxa has received the grant and does not know it yet. Immediately afterwards, Kilian receives the tweet. Maruxa does not receive any messages. After this situation, modelled by the sequence $n!(@_m p \wedge \neg @_m Bp)$. $k:r$, it still holds that $@_m p \wedge \neg @_m Bp$.

Let us now move on to axiomatising the logic.

9.3.4 Reduction axioms

We now proceed to determining the logic of asynchronous social epistemic models, ASE_L. The key to this is that all the dynamic modalities can be reduced, i.e., every formula in the language $\mathcal{L}_{\text{ASEL}}$ is provably equivalent to a formula in \mathcal{L}_{SEL} .

To start things off:

Proposition 9.25. *The following equivalences are ϵ -valid:*

- ($R_{\perp 1}$) $[\alpha. n:r]\perp \leftrightarrow [\alpha]\perp$;
- ($R_{\perp 2}$) $[\alpha. n!\phi]\perp \leftrightarrow ([\alpha]\perp \vee \neg @_n[\alpha]B\phi)$;
- (R_p) $[\alpha]p \leftrightarrow ([\alpha]\perp \vee p)$;
- (R_n) $[\alpha]n \leftrightarrow ([\alpha]\perp \vee n)$;
- (R_{\neg}) $[\alpha]\neg\phi \leftrightarrow ([\alpha]\perp \vee \neg[\alpha]\phi)$;
- (R_{\vee}) $[\alpha](\phi \wedge \psi) \leftrightarrow ([\alpha]\phi \wedge [\alpha]\psi)$;
- ($R_{@}$) $[\alpha]@_n\phi \leftrightarrow @_n[\alpha]\phi$;
- (R_F) $[\alpha]F\phi \leftrightarrow ([\alpha]\perp \vee F[\alpha]\phi)$.

Proof. Let us focus on the last two lines, leaving the rest to the reader. We prove soundness of their dual forms $\langle\alpha\rangle @_n\phi \leftrightarrow @_n\langle\alpha\rangle\phi$ and $\langle\alpha\rangle \hat{F}\phi \leftrightarrow \langle\alpha\rangle \top \wedge \hat{F}\langle\alpha\rangle\phi$.

For $R_{@}$, we note that $w, a, \epsilon \models @_n\langle\alpha\rangle\phi$ iff $w, \underline{n}, \epsilon \models \langle\alpha\rangle\phi$ iff (by Lemma 9.21) $w \bowtie \alpha$ and $w, \underline{n}, \alpha \models \phi$. The second of these items is equivalent to $w, a, \alpha \models @_n\phi$. This, along with the fact that $w \bowtie \alpha$, is equivalent by Lemma 9.21 to $w, a, \epsilon \models \langle\alpha\rangle @_n\phi$.

For R_F , we have: $w, a, \epsilon \models \langle \alpha \rangle \hat{F}\phi$ if and only if (by Lemma 9.21) $w \bowtie \alpha$ and $w, a, \alpha \models \hat{F}\phi$ iff (by Lem. 9.21 and the semantics of F) $w, a, \epsilon \models \langle \alpha \rangle \top$ and $w, b, \alpha \models \phi$ for some b such that $a \succ_w b$. This last item is equivalent to $w \bowtie \alpha$ and $w, b, \epsilon \models \langle \alpha \rangle \phi$ for some b with $a \succ_w b$, and thus $w, a, \epsilon \models \langle \alpha \rangle \top \wedge \hat{F}\langle \alpha \rangle \phi$. ■

Another observation pertains to the doxastic B modality, and will give us a hint of the reduction axiom we will find for it.

Proposition 9.26. *Let α be a history. The following holds:*

$$w, a, \epsilon \models \langle \alpha \rangle \hat{B}\phi \text{ if and only if } w, a, \epsilon \models \langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_{(w,a)} \beta} \hat{B}\langle \beta \rangle \phi.$$

Proof. As given by Lemma 9.21, if $w, a, \epsilon \models \langle \alpha \rangle \hat{B}\phi$, then necessarily $w \bowtie \alpha$ and thus $w, a, \epsilon \models \langle \alpha \rangle \top$. On the other hand, since $w, a, \alpha \models \hat{B}\phi$, there must exist a pair (v, β) such that $w \sim_a v$, and $\alpha \triangleright_{(v,a)} \beta$, and $v \bowtie \beta$, and $v, a, \beta \models \phi$. But this means that, for some β such that $\alpha \triangleright_{(w,a)} \beta$ (recall that $\triangleright_{(w,a)} = \triangleright_{(v,a)}$ in k-y-f models), and for some successor v of w , we have that $v, a, \epsilon \models \langle \beta \rangle \phi$. Therefore, $w, a, \epsilon \models \hat{B}\langle \beta \rangle \phi$, again, for some β such that $\alpha \triangleright_{(w,a)} \beta$, whence

$$w, a, \epsilon \models \bigvee_{\alpha \triangleright_{(w,a)} \beta} \hat{B}\langle \beta \rangle \phi.$$

The converse is analogous. ■

As it stands, the formula

$$\langle \alpha \rangle \hat{B}\phi \leftrightarrow \langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_{(w,a)} \beta} \hat{B}\langle \beta \rangle \phi$$

only looks like a reduction axiom. The disjunction in this formula quantifies over all possible histories which are in the $\triangleright_{(w,a)}$ relation with α ; obviously, this means that said formula need not be the same for each pair (w, a) at each model.² But it does give us a starting point. In order to find a reduction which holds everywhere, we need to quantify over all possible $\text{nom}(\alpha)$ -pseudomodels.

Given a partition A over some finite set of nominals $N \subset \text{Nom}$, and given $a \in A \cup \{\emptyset\}$ and $w \subseteq A^2$, let us define the following abbreviation:

²Compare this with the framework of AA in Chapter 6, wherein the set of agents is encoded in the syntax and there are no worries about announcements only being sent to certain agents. In that framework, the \triangleright_a relation has thus a purely syntactic definition and whether $\alpha \triangleright_a \beta$ is independent of the model. In AAL, $\langle \alpha \rangle \hat{B}_a \phi \leftrightarrow \langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \phi$ is a reduction axiom

$$\begin{aligned} \text{is}_{(w,a)}^A = & \bigwedge_{n \in a} n \wedge \bigwedge_{n \in N \setminus a} \neg n \wedge \bigwedge_{\substack{n,m \in N \\ n \equiv_A m}} @_n m \wedge \bigwedge_{\substack{n,m \in N \\ n \not\equiv_A m}} \neg @_n m \wedge \\ & \bigwedge_{\substack{n,m \in N \\ ([n],[m]) \in w}} @_n \hat{F}m \wedge \bigwedge_{\substack{n,m \in N \\ ([n],[m]) \notin w}} \neg @_n \hat{F}m, \end{aligned}$$

where \equiv_A and $[\cdot]$ denote respectively the equivalence relation and equivalence classes induced by A . The formula $\text{is}_{(w,a)}^A$ completely determines which names are equivalent in N , which equivalence classes are related by w , and which set of names corresponds to a . The following fact is thus rather obvious:

Lemma 9.27. *Given a model $M = (W, A, \sim, \succ, V)$, a pair $w \in W$, $a \in A$, and an N -pseudomodel (A', w', a') , we have: $w, a, \epsilon \models \text{is}_{(w',a')}^{A'}$ if and only if (A', w', a') is the N -pseudomodel induced by (M, w, a) . ■*

In other words, for every model (W, A, \sim, \succ, V) , every pair (w, a) and every finite set of nominals N , there exist exactly one partition A' of N , one element $a' \in A'$ and one relation $w' \subseteq A'^2$ such that $\text{is}_{(w',a')}^{A'}$ holds at (w, a, ϵ) .

In particular, for any finite set of nominals N , the finite disjunction

$$\bigvee_{\substack{(A,w,a) \\ N\text{-pseudomodel}}} \text{is}_{(w,a)}^A$$

holds everywhere. (More particularly, one exact disjunct thereof holds at every pair (w, a) .)

This observation, plus Proposition 9.26 and Lemma 9.27 give us the desired reduction:

Corollary 9.28. *The following equivalence is sound:*

$$\langle \alpha \rangle \hat{B}\phi \leftrightarrow \langle \alpha \rangle \top \wedge \bigvee_{\substack{(A,w,a) \\ \text{nom}(\alpha)\text{-pseudomodel}}} \left(\text{is}_{(w,a)}^A \wedge \bigvee_{\alpha \triangleright_{(w,a)}^A \beta} \hat{B}\langle \beta \rangle \phi \right).$$

■

The dual form of the above equivalence is:

$$[\alpha] B\phi \leftrightarrow [\alpha] \perp \vee \bigwedge_{\substack{(A,w,a) \\ \text{nom}(\alpha)\text{-pseudomodel}}} \left(\text{is}_{(w,a)}^A \rightarrow \bigwedge_{\alpha \triangleright_{(w,a)}^A \beta} B[\beta]\phi \right) \quad (R_B)$$

We are now in a position to define and show completeness of Asynchronous Social Epistemic Logic.

Definition 9.29 (The logic ASEL). The logic ASEL consists of the axioms and rules of SEL plus axioms $(R_{\perp 1})$, $(R_{\perp 2})$, (R_p) , (R_n) , (R_{\neg}) , (R_{\vee}) , $(R_{@})$, (R_F) in Prop. 9.25, and axiom (R_B) above.

Lemma 9.30. *Let $\psi \in \mathcal{L}_{\text{SEL}}$ be an announcement and reading-free formula, let $M = (W, A, \sim, \succ, V)$ be a SEL model, $w \in W$ and $a \in A$. Then $w, a, \epsilon \models^{\text{ASEL}} \psi$ iff $w, a \models_{\text{SEL}} \psi$.*

Theorem 9.31. *ASEL is sound and complete with respect to SEL models.*

Proof outline. This proof is similar to that of Theorem 6.53. Using the reduction axioms, one can show that for every formula $[\alpha]\phi \in \mathcal{L}_{\text{ASEL}}$ there is a provably equivalent formula ψ in the announcement-free language \mathcal{L}_{SEL} (in the sense that $\models_{\epsilon} [\alpha]\phi \leftrightarrow \psi$).

With this, if a formula ϕ is not a theorem of ASEL, then the corresponding announcement-free formula ψ cannot be a theorem of SEL (for it it were true everywhere, by the previous Lemma so would be ϕ), and thus there is a model refuting it; this same model refutes ϕ , again by the previous Lemma.

This is proven, using the fact that every formula of ASEL is of the form $[\alpha]\phi$ for some α , by \ll -recursion on (α, ϕ) .

For the cases $(\alpha, n!\phi)$ and $(\alpha, n:r)$, one uses the (R_{\perp}) reduction axioms and the fact that $(\alpha, \perp) \ll (\alpha, n!\phi, \perp)$ and $(\alpha, B\phi) \ll (\alpha, n!\phi, \perp)$.

For the cases (α, p) and (α, n) , one uses the corresponding reduction axioms and the fact that $(\epsilon, x) \ll (\alpha, x)$, with $x \in \text{Nom} \cup \text{Prop}$.

For the cases $(\alpha, \neg\phi)$, $(\alpha, \phi \wedge \psi)$, $(\alpha, @_n\phi)$ and $(\alpha, F\phi)$, one uses the corresponding reduction axioms and the fact that (α, ϕ) is \ll -smaller than $(\alpha, \neg\phi)$, $(\alpha, \phi \wedge \psi)$, $(\alpha, @_n\phi)$ and $(\alpha, F\phi)$.

For the case $(\alpha, B\phi)$, one uses the axiom (R_B) and the fact that, if $\alpha \triangleright_{(w,a)}^A \beta$, then $(\beta, \phi) \ll (\alpha, B\phi)$, and $(\alpha, \perp) \ll (\alpha, B\phi)$, in order to apply the induction hypothesis. ■

We note that ASEL is a finitary axiomatisation and that, since SEL has the Finite Model Property (as shown in [105, 10]), so does ASEL. Therefore:

Proposition 9.32. *ASEL is decidable.*

Discussion. We have proposed several dynamic extensions (along both dimensions) to the framework introduced in *Logic in the Community* [92], and provided their sound and complete logics.

The proposed notion of ‘social updates’ using arrow update logics is, in its current state, rather simplistic. Some tentative ideas to extend it include:

- Using arrow updates to account for the formation of ‘new friendships’.
- Establishing a notion of ‘deal breakers’. These are properties that a user finds so abhorrent that she will stop following anyone once she learns they have that property. Something like this could be modelled by adding a map to the models $db : A \rightarrow 2^{\mathcal{L}_{SEL}}$, mapping each agent to her deal breakers.

Asynchronous social epistemic logic is an attempt to somewhat realistically portray the diffusion of information in a message-centered social network such as Twitter. Many notions are not accounted for, and would be interesting directions for future investigation. These include:

- The fact that an agent might not read the entirety of the queued messages every time she logs onto Twitter. One could be tempted to make the $\triangleright_{(w,a)}$ relation to quantify over all subsequences of $\alpha \upharpoonright_{(w,a)}$, although that would give rise to a rather trivial notion of belief: agents would only believe things they already believed before checking Twitter.
- Retweets (or, more generally, spreading someone else’s message).
- Awarely or unawarely spreading misinformation; agents not believing the tweets they read; agents not trusting certain users.

For the sake of brevity, this chapter has not been concerned with the extension of SEL mentioned in Section 9.1, namely SEL_{\downarrow} , which adds ‘state variables’ x, y, \dots and an indexical operator $\downarrow x$. Most of the results in this paper can be easily extended to that framework. In particular, the reduction formulas $[n!\phi]\downarrow x.\psi \leftrightarrow \downarrow x.[n!\phi]\psi$, $[U]\downarrow x.\phi \leftrightarrow \downarrow x.[U]\phi$, and $[\alpha]\downarrow x.\phi \leftrightarrow \downarrow x.[\alpha]\phi$ (renaming bound variables whenever necessary) are valid in the different frameworks proposed here.

Conclusions

These were the verses that could be read; in the others, the writing was worm-eaten, and they were given to an academician to be deciphered. Our best information is that he has done so, after many long nights of laborious study, and intends to publish them, hoping for a third sally by Don Quixote.

*Forsi altro canterà con miglior plectio.*¹

*Ending of the first part of Don Quixote,
by Miguel de Cervantes Saavedra (1605).*

WE SHALL PUT AN END TO THIS THESIS by summing up both the most salient results that were discussed in it and, perhaps most importantly, by reminding the reader of some of the results which were not obtained; most of these have already been talked about in the last paragraph of each chapter.

We shall as well dedicate a section of this introduction to discuss possible directions these investigations could take. Some of the ideas for future research presented below were set to become, at some point, part of this dissertation; many of those questions are still worth answering, and many of the insights acquired during those abandoned explorations constitute starting points for future endeavours to solve them.

10.1 What we did and did not do

The first part of this thesis (at least this is my hope) served the double purpose of putting forward an argument for an independent study of indexed and orthogonal frames, and of taking some steps in the direction of this study. The latter came in the form of the results discussed in Chapter 3. The former was done by highlighting both the ubiquity of indexed and orthogonal frames

¹Two points of note about this quote: firstly, the last line in Italian is a bad transcription of a verse by Ariosto, *Forse altri canterà con miglior plectro*. Most English translations of *Don Quixote* seem to maintain this misquote; I do not expect the reader to show this indulgence to my errors. Secondly, there is no apparent consensus as to whether Cervantes intended (or expected) to write a sequel upon finishing the first part of *Don Quixote*. The strongest argument for the fact that he did (see e.g. [45]) rests on the last few lines of the first part, reproduced above. An often overlooked aspect of this interpretation resides in the fact that, were this indeed the case, Cervantes is referring to his future self as *someone else with a better plectrum*. I think that is quite wholesome.

in the literature and the fact that these structures are isomorphic to each other.

One of the arguments, as was claimed in Chapters 4 and 5, resides in the expectation that the results of the first part will help apply conventional techniques to unconventional bidimensional structures. This possibility was illustrated in two ways: first, Chapter 5 offered a rather simplified proof of the (already proven) completeness of SEL by using the novel technique of *indexed canonical models*; on the other hand, the end of Chapter 4 was dedicated to showing, with our isomorphisms on hand, that different species of models for STIT logics that were believed to merely ‘satisfy the same formulas’ are in fact isomorphic to each other.

I do believe this is a thread worth pulling – where can a study of indexed frames as orthogonal structures take us? Here are a few (still unanswered) questions:

- Can we simplify the conditions given in Prop. 4.4 for a frame to be isomorphic to a subset space? How much do these conditions change if we demand that the subset space be closed under intersections, or that it be a topology?

Suppose we abstractly redefine the concept of a subset space via the frames described in Prop. 4.4. Could this serve as a departure point, for instance, for a framework of multi-agent dynamic subset space logic? If so, how would it differ from the multi-agent SSL of [102] or the dynamic SSL of, e.g., [26]? We shall expand on this question in the next section.

- Can we utilise the orthogonal structure semantics described in Section 5.5 to simplify some of the proofs offered in [105], or more generally to advance the study of Social Epistemic Logic? I conjecture (perhaps baselessly) that the logic of non-rigid SEL for orthogonal structures admits a more-or-less standard, canonical-model based, completeness proof.
- How can these investigations further develop the study of multi-dimensional structures? One would have to define a suitable generalisation of the concept of orthogonality (this is briefly discussed at the end of Chapter 3).
- It is well-known that topological spaces can be seen as a generalisation of preordered sets [23]. What is the topological generalisation of preordered indexed frames? What is its ‘orthogonal’ isomorphic counterpart? We discuss some tentative ideas to this respect in the next section.

The second part of this thesis introduced several frameworks for asynchronous announcements, all of them building up from the framework AA which was the subject of Chapter 6. Variants of this framework which consider group reception of messages (Chapter 7) or the possibility of reasoning over potential future announcements (Chapter 8) were developed in those pages, as well as multiple proposals to extend the framework of SEL dynamically, the most notable being the asynchronous version of SEL presented in Chapter 9.

Some open questions about these frameworks include:

- The $[a^*]$ reading operator defined for asynchronous SEL contrasts with the $[a]$ operator used in AA; the latter represents a reading of the last message of the queue, whereas the former represents a reading of all previously unread messages. Can we consider a variant of AA with the $[a^*]$ operator instead of (or alongside) $[a]$? How does it differ from the given framework?
- How can we implement in the models discussed here concepts such as lost messages, private messages only sent to one or multiple agents, non-chronological reception of messages, etc?
- The framework for asynchronous SEL at the end of Chapter 9 admits some improvements: for instance, the requirement that an agent who sends a message must be aware of the fact that she sent a message is something that could (and should) be implemented.
- Perhaps the most interesting question: how can we define a suitable notion of knowledge in any of these frameworks, without falling into immediate circularity problems? Some ideas towards an answer are explored in the next section.

10.2 Lost causes and future perspectives

Some of the questions formulated above simply exist as questions. For most of these, time constraints did not allow for a proper exploration of their possible answers. Some others, however, have received varying degrees of attention during the course of this investigation.

Occasionally the path these questions were heading towards veered a bit far from the scope of this thesis; not any less frequently, this path seemed to drive into a brick wall. For one reason or another, the lines of research presented below were abandoned. Let us rescue some aspects of them here, along with some hopes for their resurrection in future research endeavours.

10.2.1 Orthogonal topological spaces

A *topological space* is a tuple (X, τ) where X is a nonempty set and τ is a collection of subsets of X with the following properties:

- $X, \emptyset \in \tau$;
- if $\sigma \subseteq \tau$, then $\bigcup \sigma \in \tau$;
- if $U, V \in \tau$, then $U \cap V \in \tau$.

A topological space is called *Alexandroff* if it also satisfies:

- if $\sigma \subseteq \tau$, then $\bigcap \sigma \in \tau$.

It is a well-known fact [23] that the category of Alexandroff spaces is isomorphic to the category of preordered (i.e. reflexive and transitive) frames: indeed, given a preordered frame (X, \leq) , we may define a topology τ_{\leq} on X by setting

$$U \in \tau_{\leq} \text{ iff for all } x \in U: x \leq y \Rightarrow y \in U.$$

Conversely, given a topological space (X, τ) , we may define a preorder \leq_{τ} on X by setting

$$x \leq_{\tau} y \text{ iff for all } U \in \tau: x \in U \Rightarrow y \in U.$$

It is the case that $\leq_{\tau_{\leq}} = \leq$ and, if τ is an Alexandroff topology, that $\tau_{\leq_{\tau}} = \tau$.

Now, the concept of a preorder within the context of these investigations makes sense in a very straightforward manner: we can easily talk about *pre-ordered* indexed frames, or their counterparts *preordered* orthogonal frames, by considering either indexed frames of the form

$$(W_1 \times W_2, \leq_1, \leq_2),$$

where both relations are reflexive and transitive, or (fully) orthogonal frames (X, \leq_1, \leq_2) where both relations satisfy the properties of orthogonality on top of being preorders.

Can we generalise this notion of preorder to a topological setting? That is, given a bi-topological space (X, τ_1, τ_2) , where both topologies are defined on X , is there a reasonable way to say that these topologies are *orthogonal* to each other?

I argue that the most suitable definition for this idea of orthogonality is the following:

Definition 10.1. A bitopological space (X, τ_1, τ_2) is *orthogonal* iff there exist bases² \mathcal{B}_1 and \mathcal{B}_2 for τ_1 and τ_2 respectively such that $|U_1 \cap U_2| \leq 1$ for all $U_1 \in \mathcal{B}_1$ and $U_2 \in \mathcal{B}_2$.

Let us outline a few reasons that make this into a suitable definition:

Firstly, it generalises the notion of a preordered orthogonal frame, in the sense that, given such a frame (X, \leq_1, \leq_2) , if we construct the topologies τ_{\leq_1} and τ_{\leq_2} using the method above, we obtain an orthogonal topological space wherein both topologies are Alexandroff; conversely, starting from an orthogonal topological space (X, τ_1, τ_2) and constructing the induced preorders \leq_{τ_1} and \leq_{τ_2} as above, we obtain a preordered orthogonal frame.

Secondly, not unlike how the notion of orthogonal frames generalises that of products, the notion of orthogonal topologies defined above generalises that of *topological products*.

Definition 10.2 (see e.g. [24]). The *product* of two topological spaces³ (X_1, τ_1) and (X_2, τ_2) is the space

$$(X_1 \times X_2, \tau_H, \tau_V),$$

where τ_H and τ_V are topologies which have as bases, respectively,

$$\mathcal{B}_H = \{U_1 \times \{x_2\} : U_1 \in \tau_1, x_2 \in X_2\} \text{ and } \mathcal{B}_V = \{\{x_1\} \times U_2 : x_1 \in X_1, U_2 \in \tau_2\}.$$

With this definition, it is rather obvious that there exist bases for both topologies with the property that their ‘basic’ open sets intersect at most at singletons: namely, the two bases above.

Using orthogonal topological spaces could potentially be useful to model certain bidimensional scenarios. For instance, using a topological semantics to model knowledge is something that has been seen in recent (and not so recent) literature [77, 24, 80, among many others]; according to this interpretation, an agent ‘knows’ something when she has a piece of evidence for it, and the collection of all the pieces of available evidence constitutes a topology. With respect to the other dimension, one could for instance conceive a variant of SEL wherein the ‘social connections’ are not grounded on relations but rather on something like ‘spheres of influence’, or wherein there are different levels of ‘friendship’ or ‘trust’ between the agents; both these scenarios could admit topological modellings.

²Recall that a *basis* of a topology τ is a set $\mathcal{B} \subseteq \tau$ which is closed under finite intersection and such that every set $U \in \tau$ can be expressed as a union of elements of \mathcal{B} .

³Note that this definition of *topological product* is unique to the field of modal logic. In elementary topology texts (see e.g. [79]), the *product* $(X_1, \tau_1) \times (X_2, \tau_2)$ is instead defined as the topological space $(X_1 \times X_2, \tau_{\times})$, where τ_{\times} has the set $\{U_1 \times U_2 : U_i \in \tau_i\}$ as a basis.

10.2.2 Multi-agent dynamic Subset Space Logic.

One of the main attractions of subset spaces is their ability to provide an evidence-based approach to knowledge. There exist multiple dynamic approaches to subset space logics in the literature [103, 27, 8, 38], as well as multi-agent frameworks for subset spaces [102, 61]. To my knowledge, the only framework which currently features a variant of SSL which is both multi-agent and dynamic is [39]; this work, however, does away with the ‘epistemic effort’ modality (that is, with \Box).

A framework for multi-agent subset space models with some form of dynamic operator and which maintains the effort modality (preferably an agent-specific effort modality \Box_a , as opposed to a unique, general, effort modality \Box as is the case in [80]) seems to be a highly desirable goal.

Proposition 4.4 allows to abstract from the notion of subset spaces and see them as orthogonal frames (X, R_K, R_\square) which satisfy certain properties: namely (SS1) – (SS5), as described in the aforementioned Proposition. With this in hand, a multi-agent variant is (if only seemingly) straightforward: let us simply consider a tuple

$$(X, R_K^1, R_\square^1, \dots, R_K^n, R_\square^n)$$

such that each pair of orthogonal relations (R_K^i, R_\square^i) satisfy (SS1) – (SS5).⁴ Also seemingly simple is to define a PAL-like semantics on this relational structures to account for public announcements.

I am, however, cautious about this simplicity.

Does this approach stop us from ‘jumping out of the epistemic range’, as discussed in [102]? (I conjecture it does not.) Do the properties (SS1) – (SS5) survive a public announcement operation? (I conjecture they do, although perhaps by imposing certain conditions on the admissible announcements.) If we unravel the isomorphism and ‘recover’ the multi-agent subset spaces corresponding to the tuples (X, R_K^i, R_\square^i) described above, to which extent do these look like the models in [102]? Is there some kind of link between their partition-based semantics and our equivalence relations?

10.2.3 Knowledge in an asynchronous setting

Defining a suitable notion of *knowledge* in the Asynchronous Announcement setting requires tweaking the \triangleright_a relation so that it becomes an equivalence

⁴There are of course some caveats: recall that the equivalence relations of \equiv_\square in Prop. 4.4 are supposed to represent the elements of the set X . In this multi-agent setting we should thus demand that all the R_\square^i ’s be contained in *the same* relation \equiv_\square .

relation. As discussed in Chapter 6, the obvious idea of setting

$$\alpha \equiv_a \beta \text{ iff } \alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a}$$

is largely problematic, and lends itself easily to circularity issues.⁵

We present in this section some possible (albeit underdeveloped) solutions to this problem.

Option 1. We limit the size of β : in this interpretation, a knows what she has read and she also knows how many unread announcements there are. This alternative fits well with the distributed computing tradition and it does seem to have some conceptual standing: one could think of multiple real-life situations in which this sort of knowledge would occur.

$$\alpha \equiv_a \beta \text{ iff } (\alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a} \text{ and } |\alpha|_! = |\beta|_!).$$

It is not clear to me whether we maintain circularity here, but we certainly have the problem of well-foundedness. It could be that $\alpha = p.a.\top$ and $\beta = p.a.K_1 \dots K_{1000}p$; thus, we have $\alpha \equiv_a \beta$ but $\deg \alpha = 0$ and $\deg \beta = 1000$. It does not seem clear that one could obtain a well-founded order \ll such that $(\beta, \phi) \ll (\alpha, K_a \phi)$, which one would need in order to define the semantics and to carry out most proofs. Under this definition, we also have an infinite number of β 's to worry about.

Option 2. This is a bit more creative and by far the simplest option mathematically, although admittedly conceptually inelegant. We demand β 's with not only the same announcements read by a , but rather *the same announcements*:

$$\alpha \equiv_a \beta \text{ iff } (\alpha \upharpoonright_! = \beta \upharpoonright_! \text{ and } |\alpha|_a = |\beta|_a).$$

This is indeed very easy to deal with mathematically: we have a finite number of β 's and thus the reduction axiom

$$\langle \alpha \rangle \hat{K}_a \phi \leftrightarrow \langle \alpha \rangle \top \wedge \bigvee_{\beta \equiv_a \alpha} \hat{K}_a \langle \beta \rangle \phi$$

is expressed in terms of a finite disjunction. Moreover, we can define \ll exactly as in Chapter 6, it will be well-founded, and it will be the case that $(\beta, \phi) \ll (\alpha, K_a \phi)$. With this, the proofs write themselves.

⁵Given any α , we have $\alpha \equiv_a \alpha.K_a \phi$. Suppose $w \bowtie \alpha$. Then $w, \alpha \models K_a \phi$ iff for every (t, β) with $R_a w t$ and $\alpha \equiv_a \beta$ and $t \bowtie \beta$ it is the case that $t, \beta \models \phi$. In particular, we have to check whether $t \bowtie \beta$ for $(t, \beta) = (w, \alpha.K_a \phi)$. But checking whether $w \bowtie \alpha.K_a \phi$ is the same as checking whether $w, \alpha \models K_a \phi$, and hence the circularity.

However, there is an obvious conceptual problem in assuming that only β 's with the same announcements as α are admitted. The relation \equiv_a is not to be interpreted anymore as epistemic possibility for agent a , i.e., \equiv_a is not interpreted from a 's perspective, but rather in a more holistic way, from an all-knowing entity who is aware of all sent messages. I do not think this can be satisfyingly justified.

Other solutions Let us get even more creative. We can solve the problem of circularity and well-foundedness by taking some weird turns:

Weird turn number 1: Histories do not contain epistemic modalities. Let us say we only allow histories to contain formulas in the language \mathcal{L}_{-K} , defined as:

$$\phi ::= p \mid \perp \mid \phi \wedge \phi \mid \neg\phi \mid [\phi]\phi \mid [a]\phi.$$

At first glance this seems highly undesirable because we want to be able to announce things about the epistemic state of an agent. We want to announce things such as ‘ b knows ϕ ’.

However, albeit in a rather sneaky way, such things can still be announced. Consider the history $\alpha = p.b.a.[b]\perp.a$. When a reads the second announcement, $[b]\perp$, she then knows that b has read p and therefore $w, \alpha \models K_a K_b p$ for any compatible w ! In a way the announcement $[b]\perp$ gives all agents the information that b has read every announcement up to that point. This is not as powerful expressively as allowing epistemic announcements, and I am not sure whether we can similarly express that b does not know something.

In any case, the fact that positive knowledge can still be communicated via announcements (even if we remove it from the announcement language) is a rather interesting observation in and of itself, and the reduced language would certainly help dealing with circularity issues.

Weird turn number 2. Belief as subjective certainty. We use our usual AA semantics. “Knowledge” is simply “belief” in situations in which the agent has received all announcements. Conversely, “believing” is “believing you know”:

$$w, \alpha \models K_a \phi \text{ iff } w, \alpha \models B_a \phi \text{ and } |\alpha|_! = |\alpha|_a.$$

Note that \triangleright_a restricted to the set $\{\alpha : |\alpha|_a = |\alpha|_!\}$ is an equivalence relation, and therefore the S5 axioms are $*$ -valid with this interpretation. The following is also valid:

$$B_a \phi \leftrightarrow B_a K_a \phi,$$

which is one of the theorems in Stalnaker’s logic of belief [97]⁶.

One could have some understandable conceptual gripes with the idea that something only constitutes knowledge if all announcements have been read. I myself see it as potentially problematic: for instance, an agent can never be certain about another agent’s knowledge (because an agent never knows whether other agent has read every announcement). In particular, for every $\alpha \neq \epsilon$ and $a \neq b$, we have the ϵ -validity: $\models [\alpha] \neg K_a K_b \phi$. Moreover, announcing knowledge (but not reading the announcement) makes the knowledge lost. We have the $*$ -validities

$$\models [K_a p . a] K_a p \text{ and } \models [K_a p] \neg K_a p,$$

which seem rather strange.

Or even stranger: $w, p . a \models K_a p$ but $w, p . a \models \langle p \rangle \neg K_a p$: if p is announced and a reads it, she knows it, but if it is announced again, she only believes it. We would certainly want to tweak this idea so that knowledge of positive formulas remains.

A potential solution for this is to make sure that (i) the agent somehow knows whether there are unread announcements, and (ii) *the value of propositional variables is allowed to change via an announcement*. This fixes to some extent the above conceptual problem: if an agent has received the announcement that p , but there is a new announcement, it could be the case that this new announcement is making $\neg p$ true, therefore the agent cannot be certain that p holds. This takes us to:

Weird turn number 3: assignments in Asynchronous Announcement Logic. Let us try to import some of the ideas of the public assignment framework [36] to the setting of Asynchronous Announcement Logic. An *assignment* comes in the form of a dynamic modality $[p \leftarrow \phi] \psi$, which is read “after the truth value of p is switched to the truth value of ϕ , ψ holds.

Therefore the language is:

$$\phi ::= p \mid \perp \mid \phi \wedge \phi \mid \neg \phi \mid B_a \phi \mid K_a \phi \mid [\phi] \phi \mid [a] \phi \mid [p \leftarrow \phi] \phi,$$

and the semantics for the non-standard connectives is as follows:

⁶I genuinely cannot remember whether the fact that this follows the philosophical intuitions of [97] has happened by design or by accident.

$$w \vDash \alpha.p \leftarrow \phi \quad \text{iff} \quad w \vDash \alpha;$$

$$w, \epsilon \models p \quad \text{iff} \quad w \in V(p);$$

$$w, \alpha.\phi \models p \quad \text{iff} \quad w, \alpha \models p;$$

$$w, \alpha.a \models p \quad \text{iff} \quad w, \alpha \models p;$$

$$w, \alpha.q \leftarrow \phi \models p \quad \text{iff} \quad w, \alpha \models p;$$

$$w, \alpha.p \leftarrow \phi \models p \quad \text{iff} \quad w, \alpha \models \phi;$$

$$w, \alpha \models [p \leftarrow \phi]\psi \quad \text{iff} \quad w, \alpha.p \leftarrow \phi \models \psi$$

$$w, \alpha \models K_a \phi \quad \text{iff} \quad w, \alpha \models B_a \phi \text{ and } |\alpha|_a = |\alpha|_1.$$

(The rest being of course as in AA).

All the reduction axioms one might expect hold. In particular, we have that

$$\langle \alpha.p \leftarrow \phi \rangle p \leftrightarrow \langle \alpha \rangle \phi \text{ and } \langle \alpha.p \leftarrow \phi \rangle q \leftrightarrow \langle \alpha \rangle q$$

are both valid.

Free from some of the conceptual gripes discussed above, this seems like an idea worth pursuing.

This scenic tour through the sprouts of unrealised ideas concludes the present dissertation. My hope for this chapter is that the interested reader might feel inclined to water these soils. Many of the ideas and open questions mentioned here constitute, in my opinion, rather interesting areas of future research. I will (most likely) tackle some of these investigations in the coming years. If I do not, hopefully somebody else will. In any case,

Forse altri canterà con miglior plettro.

Bibliography

- [1] T. Ågotnes, P. Balbiani, H. v. Ditmarsch, and P. Seban. Group announcement logic. *Journal of Applied Logic*, 8(1):62–81, 2010.
- [2] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of symbolic logic*, pages 510–530, 1985.
- [3] K. Apt, D. Grossi, and W. van der Hoek. Epistemic protocols for distributed gossiping. In *Proceedings of 15th TARK*, 2015.
- [4] C. Areces and B. ten Cate. Hybrid logics. *Handbook of Modal Logic*, pages 821–868, 2006.
- [5] M. Attamah, H. van Ditmarsch, D. Grossi, and W. van der Hoek. Knowledge and gossip. In *Proc. of 21st ECAI*, pages 21–26. IOS Press, 2014.
- [6] P. Balbiani. Putting right the wording and the proof of the Truth Lemma for APAL. *Journal of Applied Non-Classical Logics*, vol. 25:pp. 2–19, Apr. 2015.
- [7] P. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. D. Lima. ‘Knowable’ as ‘known after an announcement’. *Review of Symbolic Logic*, 1(3):305–334, 2008.
- [8] P. Balbiani, H. v. Ditmarsch, and A. Kudinov. Subset space logic with arbitrary announcements. In *Indian Conference on Logic and Its Applications*, pages 233–244. Springer, 2013.
- [9] P. Balbiani and S. Fernández González. Indexed frames and hybrid logics. In *Advances in Modal Logic*, 2020.
- [10] P. Balbiani and S. Fernández González. Indexed frames and hybrid logics. In *Advances in Modal Logic*, 2020.
- [11] P. Balbiani and S. Fernández González. Orthogonal frames and indexed relations. In A. Silva, R. Wassermann, and R. de Queiroz, editors, *Logic, Language, Information, and Computation*, pages 219–234, Cham, 2021. Springer International Publishing.
- [12] P. Balbiani and H. van Ditmarsch. A simple proof of the completeness of APAL. *Studies in Logic*, 8(1):65–78, 2015.

- [13] P. Balbiani, H. van Ditmarsch, and S. Fernández González. From public announcements to asynchronous announcements. In *ECAI 2020: 24th European Conference on Artificial Intelligence*, 2020.
- [14] P. Balbiani, H. van Ditmarsch, and S. Fernández González. Quantifying over asynchronous information change. In *Advances in Modal Logic*, 2020.
- [15] P. Balbiani, H. van Ditmarsch, and S. F. González. Asynchronous announcements. *ACM Transactions in Computational Logic*, 23(2), 2022.
- [16] P. Balbiani, H. van Ditmarsch, A. Herzig, and T. de Lima. Some truths are best left unsaid. In *Proc. of 9th Advances in Modal Logic*, pages 36–54. College Publications, 2012.
- [17] A. Baltag, L. 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.
- [18] A. Baltag, L. S. Moss, and S. Solecki. The logic of common knowledge, public announcements, and private suspicions. In *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, pages 43–56, 1998.
- [19] A. Baltag, A. Özgün, and A. Vargas Sandoval. Apal with memory is better. In L. S. Moss, R. de Queiroz, and M. Martinez, editors, *Logic, Language, Information, and Computation*, pages 106–129. Springer, 2018.
- [20] N. D. Belnap, M. Perloff, M. Xu, et al. *Facing the future: agents and choices in our indeterminist world*. Oxford University Press, 2001.
- [21] 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.
- [22] J. van Benthem. One is a lonely number: on the logic of communication. In *Logic colloquium 2002. Lecture Notes in Logic, Vol. 27*, pages 96–129. A.K. Peters, 2006.
- [23] J. van Benthem and G. Bezhanishvili. Modal logics of space. In *Handbook of spatial logics*, pages 217–298. Springer, 2007.

- [24] J. van Benthem, G. Bezhanishvili, B. ten Cate, and D. Sarenac. Multimodal logics of products of topologies. *Studia Logica*, 84(3):369–392, 2006.
- [25] van Benthem, J., J. Gerbrandy, T. Hoshi, and E. Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, 38:491–526, 2009.
- [26] A. Bjorndahl. Topological subset space models for public announcements. In H. van Ditmarsch and G. Sandu, editors, *Jaakko Hintikka on knowledge and game-theoretical semantics*, Outstanding contributions to logic 12, pages 165–186. Springer, 2018.
- [27] A. Bjorndahl. Topological subset space models for public announcements. In H. van Ditmarsch and G. Sandu, editors, *Jaakko Hintikka on Knowledge and Game Theoretical Semantics*, pages 165–186. Springer, 2018.
- [28] P. Blackburn, M. d. Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
- [29] P. Blackburn, M. d. Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
- [30] P. Blackburn and B. ten Cate. Pure extensions, proof rules, and hybrid axiomatics. *Studia Logica*, 84(2):277–322, 2006.
- [31] J. P. Burgess. The unreal future. *Theoria*, 44(3):157–179, 1978.
- [32] K. Chandy and J. Misra. How processes learn. In *Proc. of the 4th PODC*, pages 204–214, 1985.
- [33] R. Ciuni and E. Lorini. Comparing semantics for temporal stit logic. *Logique et Analyse*, 61(243):299–339, 2017.
- [34] A. Dabrowski, L. Moss, and R. Parikh. Topological reasoning and the logic of knowledge. *Ann. Pure Appl. Logic*, 78(1-3):73–110, 1996.
- [35] L. Dissing and T. Bolander. Implementing theory of mind on a robot using dynamic epistemic logic. In C. Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 1615–1621. International Joint Conferences on Artificial Intelligence Organization, 7 2020. Main track.

- [36] H. van Ditmarsch, A. Herzig, and T. de Lima. Public announcements, public assignments and the complexity of their logic. *Journal of Applied Non-Classical Logics*, 22(3):249–273, 2012.
- [37] H. van Ditmarsch, W. v. d. Hoek, B. Kooi, and L. Kuijer. Arbitrary arrow update logic. *Artificial Intelligence*, 242:80–106, 2017.
- [38] H. van Ditmarsch, S. Knight, and A. Özgün. Announcement as effort on topological spaces. *Electronic Proceedings in Theoretical Computer Science*, 215:283–297, 06 2016.
- [39] H. van Ditmarsch, S. Knight, and A. Özgün. Private announcements on topological spaces. *Studia Logica*, 106(3):481–513, 2018.
- [40] H. van Ditmarsch and B. Kooi. The secret of my success. *Synthese*, 151:201–232, 2006.
- [41] H. van Ditmarsch and B. Kooi. *One Hundred Prisoners and a Light Bulb*. Copernicus, 2015.
- [42] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2008.
- [43] H. van Ditmarsch, W. van der Hoek, B. Kooi, and L. Kuijer. Arbitrary arrow update logic. *Artificial Intelligence*, 242:80–106, 2017.
- [44] H. van Ditmarsch, J. van Eijck, and W. Wu. One hundred prisoners and a lightbulb — logic and computation. In *Proc. of KR 2010 Toronto*, pages 90–100, 2010.
- [45] J. Étienvre. La elusión del apócrifo en la segunda parte del quirote: final del juego. *Criticón*, (127):93–103, 2016.
- [46] S. Fernández González. Change in social networks: Some dynamic extensions of Social Epistemic Logic. *Journal of Logic and Computation*, 2022.
- [47] T. French and H. van Ditmarsch. Undecidability for arbitrary public announcement logic. In *Advances in Modal Logic 7*, pages 23–42. College Publications, 2008.
- [48] D. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyashev. *Many-dimensional modal logics: Theory and applications*. Elsevier, 2003.

- [49] G. Gargov and V. Goranko. Modal logic with names. *Journal of Philosophical Logic*, 22(6):607–636, 1993.
- [50] B. Genest, D. Peled, and S. Schewe. Knowledge = Observation + Memory + Computation. In *Proc. of 18th FoSSaCS*, LNCS 9034, pages 215–229, 2015.
- [51] K. Georgatos. Knowledge theoretic properties of topological spaces. In *International Conference on Logic at Work*, pages 147–159. Springer, 1992.
- [52] J. Gerbrandy and W. Groeneveld. Reasoning about information change. *Journal of Logic, Language, and Information*, 6:147–169, 1997.
- [53] E. Gettier. Is justified true belief knowledge? *Analysis*, 23:121–123, 1963.
- [54] 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*, pages 73–87, 2018.
- [55] J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100, 1991.
- [56] J. Hales. Arbitrary action model logic and action model synthesis. In *Proc. of 28th LICS*, pages 253–262. IEEE, 2013.
- [57] J. Hales. *Quantifying over epistemic updates*. PhD thesis, School of Computer Science & Software Engineering, University of Western Australia, 2016. <https://research-repository.uwa.edu.au/en/publications/quantifying-over-epistemic-updates>.
- [58] J. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.
- [59] J. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.
- [60] J. Halpern and L. 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.

- [61] B. Heinemann. Logics for multi-subset spaces. *Journal of Applied Non-Classical Logics*, 20(3):219–240, 2010.
- [62] M. Herlihy, D. Kozlov, and S. Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, San Francisco, CA, USA, 2013.
- [63] J. Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- [64] D. Hunt. Two problems with knowing the future. In *The Importance of Time*, pages 207–223. Springer, 1997.
- [65] A.-M. Kermarrec and M. van Steen. Gossiping in distributed systems. *SIGOPS Oper. Syst. Rev.*, 41(5):2–7, 2007.
- [66] S. Knight, B. Maubert, and F. Schwarzenruber. Asynchronous announcements in a public channel. In *Proc. of 12th ICTAC*, LNCS 9399, pages 272–289, 2015.
- [67] S. Knight, B. Maubert, and F. Schwarzenruber. Reasoning about knowledge and messages in asynchronous multi-agent systems. *Mathematical Structures in Computer Science*, 29(1):127–168, 2019.
- [68] B. Kooi and B. Renne. Arrow update logic. *Review of Symbolic Logic*, 4(4), 2011.
- [69] M. Kracht and F. Wolter. Properties of independently axiomatizable bimodal logics. *The Journal of Symbolic Logic*, 56(4):1469–1485, 1991.
- [70] A. Kshemkalyani and M. Singhal. *Distributed Computing: Principles, Algorithms, and Systems*. Cambridge University Press, New York, NY, USA, 2008.
- [71] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.
- [72] Z. Liang and J. Seligman. The dynamics of peer pressure. In *International Workshop on Logic, Rationality and Interaction*, pages 390–391, 2011.
- [73] F. Liu, J. Seligman, and P. Girard. Logical dynamics of belief change in the community. *Synthese*, 191(11), 2014.

- [74] A. Lomuscio and M. Ryan. An algorithmic approach to knowledge evolution. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 13(2):119–132, 1998.
- [75] D. Makinson. Some embedding theorems for modal logic. *Notre Dame Journal of Formal Logic*, 12(2):252–254, 1971.
- [76] Y. Moses, D. Dolev, and J. Halpern. Cheating husbands and other stories: a case study in knowledge, action, and communication. *Distributed Computing*, 1(3):167–176, 1986.
- [77] L. S. Moss and R. Parikh. Topological reasoning and the logic of knowledge. In *TARK*, volume 92, pages 95–105, 1992.
- [78] M. Mukund and M. Sohoni. Keeping track of the latest gossip in a distributed system. *Distributed Computing*, 10(3):137–148, 1997.
- [79] J. Munkres. *Topology*. Prentice Hall, 2000.
- [80] A. Özgün. *Evidence in Epistemic Logic: A Topological Perspective*. PhD thesis, University of Lorraine & University of Amsterdam, 2017.
- [81] P. Panangaden and K. Taylor. Concurrent common knowledge: Defining agreement for asynchronous systems. *Distributed Computing*, 6:73–93, 1992.
- [82] R. Parikh and R. Ramanujam. A knowledge based semantics of messages. *Journal of Logic, Language and Information*, 12:453–467, 2003.
- [83] S. Passy and T. Tinchev. An essay in combinatory dynamic logic. *Information and Computation*, 93(2):263–332, 1991.
- [84] T. Pedersen and M. Slavkovik. Formal models of conflicting social influence. In *PRIMA 2017: Principles and Practice of Multi-Agent Systems*, pages 349–365. Springer International Publishing, 2017.
- [85] J. Picado and A. Pultr. *Frames and Locales: topology without points*. Springer Science & Business Media, 2011.
- [86] J. Plaza. Logics of public communications. In *Proc. of the 4th ISMIS*, pages 201–216. Oak Ridge National Laboratory, 1989.
- [87] R. Ramanujam. Local knowledge assertions in a changing world. In Y. Shoham, editor, *Proc. of 6th TARK*, pages 1–14. Morgan Kaufmann, 1996.

- [88] R. Ramanujam. View-based explicit knowledge. *Ann. Pure Appl. Logic*, 96(1-3), 1999.
- [89] J. Ruan and M. Thielscher. A logic for knowledge flow in social networks. In *Australasian Joint Conference on Artificial Intelligence*, pages 511–520. Springer, 2011.
- [90] K. Sano. Axiomatizing epistemic logic of friendship via tree sequent calculus. In *International Workshop on Logic, Rationality and Interaction*, pages 224–239, 2017.
- [91] F. Schwarzentruher. Epistemic reasoning in artificial intelligence. Habilitation Thesis, University of Rennes, <http://people.irisa.fr/Francois.Schwarzentruher/hdr/>, 2019.
- [92] J. Seligman, F. Liu, and P. Girard. Logic in the community. In *Indian Conference on Logic and Its Applications*, pages 178–188. Springer, 2011.
- [93] J. Seligman, F. Liu, and P. Girard. Logic in the community. In *Indian Conference on Logic and Its Applications*, pages 178–188. Springer, 2011.
- [94] J. Seligman, F. Liu, and P. Girard. Facebook and the epistemic logic of friendship. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2013)*, pages 230–238, 2013.
- [95] S. Smets and F. R. Velázquez-Quesada. How to make friends: A logical approach to social group creation. In *International Workshop on Logic, Rationality and Interaction*, pages 377–390. Springer, 2017.
- [96] A. Solaki, Z. Terzopoulou, and B. Zhao. Logic of closeness revision: Challenging relations in social networks. In *Proceedings of the ESSLLI Student Sessions*, pages 123–134, 2016.
- [97] R. Stalnaker. On logics of knowledge and belief. *Philosophical Studies*, 128(1):169–199, 2005.
- [98] C. Steinsvold. *Topological Models of Belief Logics*. PhD thesis, CUNY Graduate Center, 2007.
- [99] 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.

- [100] A. Tarski. Theorems on the existence of successors of cardinals, and the axiom of choice. In *Indagationes Mathematicae (Proceedings)*, volume 57, pages 26–32. Elsevier, 1954.
- [101] R. H. Thomason. Combinations of tense and modality. In *Handbook of philosophical logic*, pages 135–165. Springer, 1984.
- [102] Y. Wáng and T. Ågotnes. Multi-agent subset space logic. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.
- [103] Y. Wáng and T. Ågotnes. Subset space public announcement logic. In *Indian Conference on Logic and Its Applications*, pages 245–257. Springer, 2013.
- [104] Z. Xiong, T. Ågotnes, J. Seligman, and R. Zhu. Towards a logic of tweeting. In *International Workshop on Logic, Rationality and Interaction*, pages 49–64, 2017.
- [105] L. Zhen. *Towards Axiomatisation of Social Epistemic Logic*. PhD thesis, University of Auckland, 2020.

Index

- N -pseudomodel, 180
- Action Model Logic, 20, 154
 - arbitrary asynchronous action models, 157
 - asynchronous action models, 154
- Arbitrary Asynchronous Announcement Logic (AAA), 139
- Arbitrary Public Announcement Logic (APAL), 20
- Arrow Update Logic, 22, 167
 - asynchronous arrow updates, 168
- Asynchronous Announcement Logic (AA), 83
 - Equivalence-based AA, 112
- Asynchronous Broadcast Logic, 111
 - DU-ABL, 116
- Asynchronous Social Epistemic Logic (ASEL), 174
- basic modal language, 14
- bisimulation, 142
- common belief, 135
- decidability, 16, 20, 23, 103, 109, 117, 140, 181
- expressivity, 100, 145
- Facebook logic, 169
- frame
 - generated subframe, 12
 - Kripke frame, 11
- fusion of two logics, 31
- history, 84, 122, 141
- hybrid logics, 61
- indexed frames, 25
- matrix enumeration, 37
- modal logics
 - KD45, 16
 - K, 15
 - S5, 16
 - S4, 16
 - normal modal logics, 15
- model, 11
- models
 - models for AA, 91
- morphisms
 - bounded morphism, 12, 36
 - isomorphism, 12
- named formula, 65
 - BCN, 65
- orthogonal frames
 - definition, 30
 - generalized, 35
- orthogonal structure, 35, 37
 - semistructure, 35, 37
 - standard orthogonal structure, 35
- Partially Synchronised Announcement Logic (PSAL), 119
- positive formulas, 83, 98, 99, 147
- products of frames, 27, 42
- Public Announcement Logic (PAL), 18
- reduction axioms, 20, 103, 118, 129, 156, 178, 192

- semantics
 - semantics of AAA, 141
 - semantics of AA, 92
 - semantics of ASEL, 175
 - semantics of PSAL, 127
 - semantics of Eq-AA, 113
- Social Epistemic Logic (SEL), 28, 59, 160
 - SEL_{\downarrow} , 74
 - SEL^+ , 63
 - SEL_{\downarrow} , 164
- STIT logic, 29, 45
 - choice B-trees, 51
 - Kamp frames, 29, 48
 - T-STIT frames, 47
- subset spaces, 27, 43
- updates
 - epistemic updates, 163
 - social updates, 167
- view relation, 84, 123, 125, 126, 141, 155