Coalition Announcements

Rustam Galimullin

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

November 2018

Abstract

Coalition announcement logic is one of the family of logics of quantified announcements. It extends public announcement logic with formulas $\langle\!\![G]\!\!\rangle\varphi$ that are read as 'there is a truthful public announcement by agents from G such that whatever agents from $A \setminus G$ announce at the same time, φ holds after the joint announcement.' The logic has enjoyed comparatively less attention than its siblings — arbitrary and group announcement logics. The reason for such a situation can be partially attributed to the inherent alternation of quantification in coalition announcements. To deal with the problem, we consider relativised group announcements that separate the coalition's announcement from the anti-coalition's response. We present coalition and relativised group announcement logic and show its completeness. Apart from that, we prove that the complexity of the model-checking problem for coalition announcement logic is PSPACE-complete in the general case, and in P in a special case of positive target formulas. We also study relative expressivity of logics of quantified announcements. In particular we show that arbitrary and coalition announcement logics are not at least as expressive as group announcement logic. Finally, we present a counter-example to the proposed definition of coalition announcements in terms of group announcements, and consider some other interesting properties.

Acknowledgements

Being a traditional part of almost every thesis, acknowledgements usually comprise of seemingly unexciting and at times annoying collections of platitudes, ready-made thanks and dull inside jokes. This recurring phenomenon, however, unveils a simultaneously mundane yet moving truth that most authors have in common: the text following the acknowledgements is massively influenced in a myriad of ways by direct, indirect, and occasionally accidental actions and words of people the author feels connected to[1]. As you, the reader, has probably already guessed, the same is applicable to my work. So I gladly honour the established tradition, and, however clumsily, make my gratitude known.

I consider the time I spent in Nottingham the most intellectually stimulating and rewarding of my life so far. For this, I am indebted to my supervisor Natasha Alechina, whose help and support made the presented work possible. From day one, I was fortunate enough to have such an outstanding mentor: not only did Natasha treat me as an equal researcher, rather than a student, she also guided me on the profound journey through a logician's day-to-day life.

This final version of the thesis has been largely improved by taking into account the suggestions and corrections of my viva examiners Thomas Ågotnes and Venanzio Capretta. Their reading of my work and discussing it with me in the examination, made me realise how unique and enjoyable the viva process is.

During these four years I was privileged to meet and work with other researchers passionate about logic. Of them, I would like to especially thank Hans van Ditmarsch and Tim French with whom I had the pleasure of sharing authorship of two papers.

I would like to thank Barteld Kooi and Allard Tamminga of Groningen University for introducing me to dynamic epistemic logic. In fact, it was during my visit to Groningen when I first heard an expression 'modal logic'. Now it is clear that this moment was pivotal in my life. When I think of either Barteld or Allard (or both) and their role in my path to the present moment, it strikes me that *everything* could have been totally different. And I am happy that it is not.

My life in the UK has been adorned by a constellation of friends. Here is a by no means an exhaustive list: Fiona Boulton, Simon Castle-Green, Jérémie Clos, Daniel Langley, Nathan Langman, Stephen Lucas, Kennedy Megran, Łukasz Targosz, Benjamin Taylor, Jonathan Thaler, Charlotte Tupper, and Jacob Wright.

Through the best of times and through the worst of times, I felt the un-

[1] Note that this connection is not necessarily reflexive, symmetric, transitive or euclidean.

remitting support from my Siberian friends. Anton Cherepov, Andrey Churikov, Galina Dorofeeva, Anton Franchuk, Alexey Gribko, Grigoriy Guryev, Evgeniy Krivobok, Roman Kulikov, Andrey Lyulyakin, Aleksandr Nikitin, Sasha Popov, Margarita Popova, and Anatoliy Shestakov — I cherish your friendship in all possible worlds.

The irony of the situation is that for the most important people in my life — my parents Rozaliya and Fanis, and my brother Tagir — everything that is written here, including this sentence, is an accumulation of lexemes of a foreign language. However, it is they who through their care and guidance made this moment possible.

Contents

Chapter 1

Introduction

Communication is a social behaviour, and as such often includes several interlocutors. It can be usually characterised by goal-orientedness, and participants in the communication are involved in a kind of game to successfully achieve their goals. These goals may differ, but learning is invariably among them. Sharing ideas, asking questions, letting others know your intentions — all these involve increasing knowledge. Contrary to the famous maxim, words sometimes are on a par with actions, and for the goals of acquiring and sharing knowledge words are the most effective actions.

A formal framework for reasoning about knowledge is epistemic logic (EL) [Fagin et al., 2004; van Ditmarsch et al., 2008], where it is possible to express propositions like 'Ann knows that Beth knows whether 16th of June is Bloomsday, and Ann does not know whether Carol knows that fact.' Already in this example our interest in the social dimension of knowledge is manifested. Indeed, in the context of communication interlocutors may possess knowledge not only of the facts they would like to share or conceal but also some higher-order knowledge about what they already know, and what others know as well.

Epistemic logic describes the *static* distribution of knowledge in some situations, and hence captures only snapshots of the learning process. In order to reason about how agents' knowledge changes as a result of communication, we should consider a *dynamic* take on knowledge. Dynamic epistemic logic [van Ditmarsch et al., 2008] is an umbrella term for various logics of knowledge and belief change. These formalisms model multiple aspects of epistemic change: suspicion [Baltag et al., 1998], lying [van Ditmarsch et al., 2012], ontic changes [van Benthem et al., 2006; van Ditmarsch and Kooi, 2008], time [Hoshi, 2010; Renne et al., 2016], and so on.

In our work we are interested in arguably the most fundamental type of epistemic actions — truthful public announcements [Plaza, 2007], [van Ditmarsch et al., 2008, Chapter 4] — that describe the situation when all agents acquire the same piece of true information, and all of them are commonly aware of everyone acquiring it. Although public announcements are quite idealised epistemic actions, they capture a plethora of communication scenarios where interlocutors trust each other and strive to decrease their ignorance. However, a formalism that allows us to reason about such epistemic events, public announcement logic (PAL), does not specify where the announcements come from. Hence, in the setting of multi-agent communication, we are particularly interested in announcements made by interlocutors, or agents. Moreover, goal-orientedness of communication suggests that we should also be able to reason about the *existence* of a sequence of announcements that achieve a given epistemic goal.

The latter requirement led to the exploration of the new exciting frontier in the area of dynamic epistemic logic. This new research is concerned with *quantification* over various epistemic actions [van Ditmarsch, 2012]. The most notable examples of such formalisms are arbitrary [Balbiani et al., 2008], group [Ågotnes et al., 2010], and coalition [Ågotnes and van Ditmarsch, 2008] announcement logics (APAL, GAL, and CAL respectively), the latter of which is the least studied one. As its name suggests, coalition announcement logic deals with coalitions and their opponents, or anti-coalitions, in epistemic communication scenarios. Such scenarios usually include situations when an epistemic goal is not only to increase knowledge of some agents but also to leave some other agents ignorant of a certain fact at the same time. For example, bidders in an auction submit their bids simultaneously, making other bidders aware of their current bid but not of the total sum at their disposal. We may be interested in whether there is a bid by a single participant which wins the auction, no matter what other bidders offer. Or, if the bidder does not have a required sum, whether there is a suitable coalition with which she can outbid everyone outside of the coalition. Other examples include communicating over an insecure channel with an eavesdropper who may add some bits of information to the messages sent between the sender and the receiver, or making everyone aware of some property of a coalition, like its number of coffee-drinkers, so that properties of individual agents in the coalition cannot be deduced no matter what agents in the anti-coalition announce.

Presence of coalitions and anti-coalitions in coalition announcement logic hints at a game-theoretic setting. Indeed, coalition announcements were partially inspired by game theory, in particular by the forcing operator [van Benthem, 2014, Chapter 11] and [van Benthem, 2001], and coalition logic [Pauly, 2002]. Thus, coalition announcement logic is one of the meeting points between games and dynamic epistemic logic.

In our work we study coalition announcements and their relation to other types of quantified announcements.

1.1 Overview and Contributions

Our thesis comprises eight chapters. We start off with the introduction (Chapter 2) of the base underlying logics, EL and PAL, along with presenting standard definitions, properties, and proof strategies of modal logic. The chapter closes with a high-level exposition of coalition logic, which can be considered as a spiritual predecessor of coalition announcement logic.

Three main logics of quantified public announcements — APAL, GAL, and CAL — are presented in Chapter 3. Not only do we introduce the formalisms but also mention all the relevant logical properties, such as complexity and expressivity. This chapter sets the stage for the rest of the thesis. We also consider a worked example.

In Chapter 4 we deal with the model-checking problem for CAL. Since the direct implementation of the truth definition for coalition announcements is impossible, we use an equivalent semantics with quantification over a finite set of strategies instead of an infinite set of possible announcements. This allows us to show that the complexity of the model-checking problem is PSPACE-complete. In addition we show that if a formula within the scope of a coalition announcement is positive, then it is enough to check the maximal informative announcement by agents in a coalition. Complexity in this case is in P. Results discussed in the chapter are presented in [Galimullin et al., 2018].

Sometimes, knowing whether a given formula of a logic is valid or not may result in stronger intuitions regarding logic's axiomatisation, expressivity, etc. Chapter 5 offers a selection of various properties of CAL and GAL. Such properties are expressed as logical formulas with operators $[\langle G\rangle]$ and $\langle[G]\rangle$ for CAL, and $[G]$ and $\langle G\rangle$ for GAL. To the best of our knowledge, none of the results presented in the chapter were proved in the literature, and some of them were mentioned as open questions. In particular, we show that $\langle[G]\rangle\varphi \leftrightarrow \langle G\rangle[A \setminus G]\varphi$ is not valid by presenting a counterexample to the right-to-left direction (and proving validity of the left-to-right direction). This formula was considered as a possible definition of coalition operators in GAL. Another intriguing result deals with commutativity of box and diamond versions of coalition and group announcement operators: the Church-Rosser principle $\Diamond\Box\varphi \rightarrow \Box\Diamond\varphi$ (where $\Diamond$ and $\Box$ are substituted with the corresponding operator) is not valid in both GAL and CAL.

Existing techniques for proving the completeness of logics of quantified public announcements [Balbiani et al., 2008; Balbiani and van Ditmarsch, 2015] seem inadequate for dealing with coalition announcements, which, being complex operators with an internal alternation of quantifiers, do not work well with the Lindenbaum lemma in the aforementioned papers. This is why we employ relativised group announcements $[G, \chi]\varphi$ that parametrise group announcements by some given formula χ that is to be announced in conjunction with G's announcement. This operator 'splits' coalition announcements and anti-coalition responses, and hence allows us to treat them separately. We present a logic with coalition and relativised group announcements in Chapter 6, and show its completeness. The chapter is based on [Galimullin and Alechina, 2017] (and the corrected version [Galimullin and Alechina, 2018]).

Relative expressivity of logics of quantified announcement has been a longstanding open question. In Chapter 7 we make some progress towards its solution. In particular, we show that CAL and APAL are not at least as expressive as GAL. To achieve this, we present two classes of models and a GAL formula that is true in one class and false in the other, and show that no CAL or APAL formula can distinguish these two classes. For evaluation of the formulas we use formula games for CAL and GAL with relativised group announcements. The latter separate the moves in a game corresponding to coalition and anti-coalition announcements.

We also mention that CAL is not at least as expressive as APAL. The chapter is based on collaborative work with Natasha Alechina, Hans van Ditmarsch, and Tim French. Formula games (Definition 7.2) were originally proposed by Tim, and contributions to the proof of Theorem 7.6 are shared equally.

In the conclusion (Chapter 8), we recapitulate the thesis, and point to some promising directions for further research.

Chapter 2

Knowledge, Announcements, and Coalitions

This chapter serves as an exposition of the basic underlying formalisms of the thesis — epistemic logic (Section 2.1) and public announcement logic (Section 2.2). The former models an agent's knowledge and ignorance of basic facts as well as her knowledge and ignorance of knowledge of other agents. Next, we move on from the static scenarios of epistemic logic to the dynamic ones of public announcement logic. The latter allows us to reason about events when all agents simultaneously receive the same piece of information, and it is common knowledge that such an event has taken place. In Section 2.3 we consider a somewhat tangential topic of coalition logic. This logic, however, is a precursor of coalition announcement logic, which is the main subject of our work, and while presenting it, we hope to indicate some the basic intuitions of the latter.

2.1 Epistemic Logic

In this section we introduce *epistemic logic* (EL) that serves as a foundation for all other logics considered in the thesis. Broadly speaking, EL studies the notion of knowledge in terms of possible states, or worlds. Such an approach to knowledge was conceived by Jaakko Hintikka in [Hintikka, 1962]. In the exposition of the subject, we limit ourselves to facts and notions that will be used in the thesis. A more comprehensive treatment of EL and modal logic in general can be found, for example, in [Blackburn et al., 2001, 2006; van Benthem, 2010; Fagin et al., 2004].

2.1.1 Introductory Example

Two agents, a and b, want to acquire the same item, and whoever offers the greatest sum, gets it. Agents may have 5, 10, or 15 pounds, and they do not know which sum the opponent has. Let agent a have 15 pounds, and agent b have 5 pounds. This situation is presented in Figure 2.1.

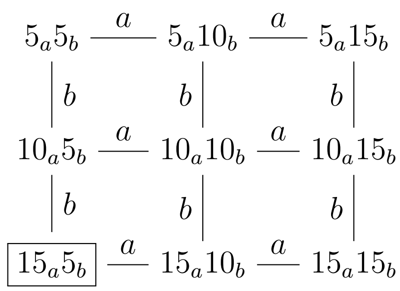

Figure 2.1: Initial model $(M, 15_a5_b)$

In the model (let us call it M), state names denote money distribution. Thus, in our designated, or actual, state 15_a5_b (boxed), agent a has 15 pounds, and agent b has 5 pounds. Formally, $(M, 15_a5_b) \models 15_a$ and $(M, 15_a5_b) \models 5_b$, or, equivalently, $(M, 15_a5_b) \models 15_a \land 5_b$. Note that in this particular example we conflate names of states and basic facts that are true in them. Labelled edges connect the states that a corresponding agent cannot distinguish. These edges represent equivalence relations, and the full version of model M is presented if Figure 2.2. Throughout the thesis we will omit reflexive and transitive arrows.

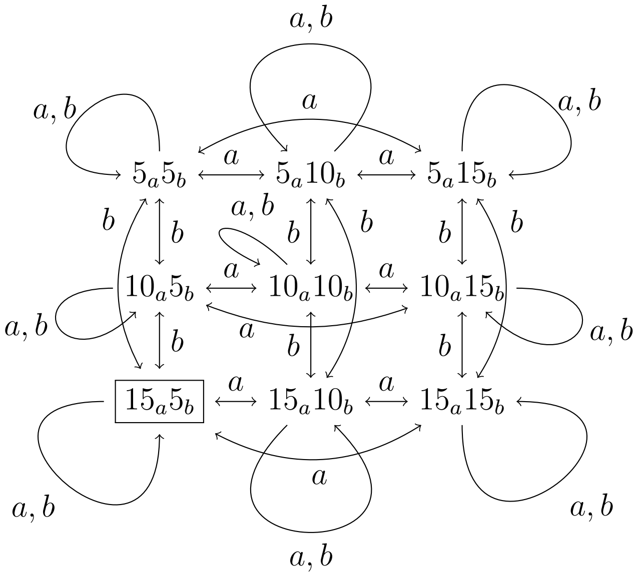

Figure 2.2: Full version of M

For example, in the actual state agent a *knows* that she has 15 pounds, but she does not know how much money agent b has; in other words, every a-arrow from state 15_a5_b leads to a state where 15_a holds, and there are a-arrows to states with 5_b, 10_b, and 15_b. Formally, $(M, 15_a5_b) \models K_a15_a \land \neg(K_a5_b \lor K_a10_b \lor K_a15_b)$.

2.1.2 Syntax and Semantics of EL

Let P denote a countable set of propositional variables, and A be a finite set of agents.

Definition 2.1 (Language of EL)**.** The *language of epistemic logic* $\mathcal{L}_{EL}$ is as follows:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi,$$

where $p \in P$, $a \in A$, and all the usual abbreviations of propositional logic (such as $\vee$, $\rightarrow$, $\leftrightarrow$) and conventions for deleting parentheses hold. Diamond version of the operator K is defined as $\widehat{K}_a\varphi := \neg K_a\neg\varphi$. The *language of propositional logic* $\mathcal{L}_{PL}$ is the one without $K_a\varphi$. Formula $K_a\varphi$ is read as 'agent a knows φ,' and the dual $\widehat{K}_a\varphi$ is read as 'agent a considers φ possible.'

Formulas of EL are interpreted on epistemic models. Figure 2.1 is an example of such a model.

Definition 2.2 (Epistemic model)**.** An *epistemic model* is a triple $M = (W, \sim, V)$, where

- W is a non-empty set of states,
- $\sim: A \rightarrow \mathcal{P}(W \times W)$ is an equivalence relation for each agent $a \in A$,
- $V : P \rightarrow \mathcal{P}(W)$ is a valuation of propositional variables $p \in P$.

A pair $(W, \sim)$ is called an *epistemic frame*, and a pair (M, w) with $w \in W$ is called a *pointed model.* M is called *finite* if W is finite. Also, we write $M_1 \subseteq M_2$ if $W_1 \subseteq W_2$, $\sim_1$ and V_1 are results of restricting $\sim_2$ and V_2 to W_1, and call M_1 a *submodel* of M_2.

Next definition specifies how truth or falsity of an epistemic formula is determined in a pointed model.

Definition 2.3 (Semantics of EL)**.** Let a pointed model (M, w) with $M = (W, \sim, V)$, $a \in A$, and $\varphi, \psi \in \mathcal{L}_{EL}$ be given. The *semantics of epistemic logic* is presented below.

$$\begin{array}{lll}
(M,w) \models p & \text{iff} & w \in V(p) \\
(M,w) \models \neg\varphi & \text{iff} & (M,w) \not\models \varphi \\
(M,w) \models \varphi \wedge \psi & \text{iff} & (M,w) \models \varphi \text{ and } (M,w) \models \psi \\
(M,w) \models K_a\varphi & \text{iff} & \text{for all } v \in W : w \sim_a v \text{ implies } (M,v) \models \varphi
\end{array}$$

The semantics for the dual of $K_a\varphi$ follows easily from the definition of $\widehat{K}_a\varphi$.

$$(M,w) \models \widehat{K}_a\varphi \quad \text{iff} \quad \text{there is } v \in W : w \sim_a v \text{ and } (M,v) \models \varphi$$

Definition 2.4 (Validity and satisfiability)**.** We call formula φ *valid* and write $\models \varphi$ if and only if for any pointed model (M, w) it holds that $(M, w) \models \varphi$. And φ is called *satisfiable* if and only if there is some (M, w) such that $(M, w) \models \varphi$.

In the thesis epistemic models with equivalence relation defined in Definition 2.2 is the only class of models we are dealing with. Hence, when we write $\models \varphi$, we mean that φ is valid on the class of epistemic models.

Definition 2.5 (Equivalence of Formulas)**.** Two formulas φ and ψ are called *equivalent*, if for all (M, w) it holds that $(M, w) \models \varphi$ if and only if $(M, w) \models \psi$.

2.1.3 Bisimulation

The basic notion of similarity in modal logic is bisimulation.

Definition 2.6 (Bisimulation)**.** Let two models $M = (W, \sim, V)$ and $M' = (W', \sim', V')$ be given. A non-empty binary relation $Z \subseteq W \times W'$ is called a *bisimulation* if and only if for all $w \in W$ and $w' \in W'$ with $(w, w') \in Z$:

- w and w' satisfy the same propositional variables;
- for all $a \in A$ and all $v \in W$: if $w \sim_a v$, then there is a v' such that $w' \sim_a v'$ and $(v, v') \in Z$;
- for all $a \in A$ and all $v' \in W'$: if $w' \sim_a v'$, then there is a v such that $w \sim_a v$ and $(v, v') \in Z$.

If there is a bisimulation between models M and M' linking states w and w', we say that (M, w) and (M', w') are bisimilar.

Note that any union of bisimulations between two models is a bisimulation, and the union of all bisimulations is a maximal bisimulation.

Definition 2.7 (Quotient model)**.** Let model M be given. The *quotient model* of $M = (W, \sim, V)$ with respect to some relation R on W is $M^R = (W^R, \sim^R, V^R)$, where $W^R = \{[w] \mid w \in W\}$ and $[w] = \{v \in W \mid wRv\}$, $[w] \sim_a^R [v]$ iff $\exists w' \in [w]$, $\exists v' \in [v]$ such that $w' \sim_a v'$ in M, and $[w] \in V^R(p)$ iff $\exists w' \in [w]$ such that $w' \in V(p)$.

Definition 2.8 (Bisimulation contraction)**.** Let model M be given. *Bisimulation contraction* of M (written $\|M\|$) is a model that is isomorphic to the quotient model of M with respect to the maximal bisimulation of M with itself. Such a maximal bisimulation is an equivalence relation.

Informally, bisimulation contraction is the minimal representation of M.
The following theorem is a well-known result in modal logic.

Theorem 2.1. Suppose (M, w) and (M', w') are bisimilar. Then for all $\varphi \in \mathcal{L}_{EL}$, $(M, w) \models \varphi$ iff $(M', w') \models \varphi$.

This fact means that no modal formula can distinguish two bisimilar states.

Corollary 2.2. $(\|M\|, [w]) \models \varphi$ iff $(M, w) \models \varphi$ for all $\varphi \in \mathcal{L}_{EL}$.

Definition 2.9 (n-bisimulation)**.** Let two models $M = (W, \sim, V)$ and $M' = (W', \sim', V')$ be given. We call pointed models (M, w) and (M', w') *n-bisimilar* if and only if there is a sequence of binary relations $Z_n \subseteq \ldots \subseteq Z_0$ with the following properties (for $i + 1 \leq n$):

- $(w, w') \in Z_n$;
- if $(v, v') \in Z_0$, then v and v' satisfy the same propositional variables;
- for all $a \in A$ and all $u \in W$: if $v \sim_a u$ and $(v, v') \in Z_{i+1}$, then there is a u' such that $v' \sim_a u'$ and $(u, u') \in Z_i$;
- for all $a \in A$ and all $u' \in W'$: if $v' \sim_a u'$ and $(v, v') \in Z_{i+1}$, then there is a u such that $v \sim_a u$ and $(u, u') \in Z_i$.

Note that bisimulation implies n-bisimulation, and not vice versa.

Intuitively, n-bisimulation means that the models behave similarly up to a certain depth n.

Let us consider an example. In Figure 2.3, states u and u', and t and t' are bisimilar, and each EL formula that satisfies one of the states in a pair, satisfies the other.

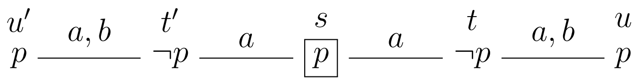

Figure 2.3: Model (M, s)

Bisimulation contraction of (M, s) that basically results in 'collapsing' sets of mutually bisimilar states into 'representative' states is shown in Figure 2.4. Note that (M, s) and $(\|M\|, s)$ satisfy the same formulas, and we cannot 'reduce' $(\|M\|, s)$ any further without compromising its bisimilarity with (M, s).

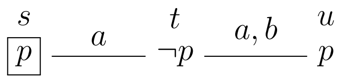

Figure 2.4: Model $(\|M\|, s)$

2.1.4 Satisfiability and model checking for EL

Two basic reasoning problems in logic are satisfiability and model checking. The common procedure for both of the problems is checking an input for some property, and returning 'yes' or 'no'.

Definition 2.10 (Satisfiability)**.** Let some formula φ be given. The *satisfiability problem* is the problem to determine whether there is a model (M, w) that satisfies φ.

Definition 2.11 (Model checking)**.** Let some finite model (M, w) and some formula φ be given. The *model checking problem* is the problem to determine whether φ is satisfied in (M, w) (or, equivalently, whether $(M, w) \models \varphi$ holds).

Theorem 2.3 ([Halpern and Moses, 1992])**.** The satisfiability problem for EL is NP-complete in the single-agent case, and PSPACE-complete in the multi-agent case.

Theorem 2.4 ([Halpern and Moses, 1992])**.** The model checking problem for EL is in P.

2.1.5 Axiomatisation and completeness of EL

Semantics tells us what formulas of $\mathcal{L}_{EL}$ are true in a given model. If we want to have a minimal set of formulas and rules of inference that allows us to derive all the validities of EL, we need a sound and complete axiomatisation.

Definition 2.12 (Axioms and rules)**.** The *axiom system for* EL is presented below.

all instantiations of propositional tautologies	
$K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$	distribution
$K_a\varphi \to \varphi$	truth
$K_a\varphi \to K_aK_a\varphi$	positive introspection
$\neg K_a\varphi \to K_a\neg K_a\varphi$	negative introspection
If $\vdash \varphi$ and $\vdash \varphi \to \psi$, then $\vdash \psi$	modus ponens
If $\vdash \varphi$, then $\vdash K_a\varphi$	necessitation

Given some axiomatisation X, if there is a finite sequence of formulas $\varphi_1, \dots, \varphi_{n-1}, \varphi$ such that every formula is either an instance of an axiom schema, or a result of application of rules of inference, then we say that φ is *derivable* in X (or, equivalently, φ is a *theorem* of X), written $\vdash_X \varphi$. However, if X is clear from the context, we will omit X and write $\vdash \varphi$.

Axioms *truth*, *positive introspection*, and *negative introspection* correspond to restrictions on the accessibility relation $\sim$. This correspondence is summarised in the table below. Recall that in our case $\sim$ is an equivalence relation.

Reflexivity	$K_a\varphi \to \varphi$	$\forall x : R(x, x)$
Transitivity	$K_a\varphi \to K_aK_a\varphi$	$\forall x, y, z : R(x, y) \wedge R(y, z) \to R(x, z)$
Euclidity	$\neg K_a\varphi \to K_a\neg K_a\varphi$	$\forall x, y, z : R(x, y) \wedge R(x, z) \to R(y, z)$

Different combinations of these axioms result in different modal logics. For example, the truth axiom marks the difference between epistemic logic and doxastic logic (logic of belief). Indeed, if an agent *knows* something, then it must be true. On the other hand, if an agent *believes* something, then it may turn out to be false (i.e. agents can have false beliefs). In the thesis we consider only epistemic logic, and more information on other similar logics and axioms may be found, for example, in [Fagin et al., 2004; van Benthem, 2010].

Consider the distribution axiom and the rule of necessitation. They mean that agents know all theorems of the logic, and their knowledge is closed under inference. Dealing with such idealised agents gives rise to the problem of *logical omniscience* (see [Fagin et al., 2004, Chapter 9]). This problem, however, is a whole separate avenue of research, and we do not consider it in the thesis.

Definition 2.13 (Soundness)**.** Let an axiom system X be given. Then, X is called *sound* if for every formula φ it holds that $\vdash_X \varphi$ implies $\models \varphi$.

Definition 2.14 (Completeness)**.** Let an axiom system X be given. Then, X is called *complete* if for every formula φ it holds that $\models \varphi$ implies $\vdash_X \varphi$.

It is usually straightforward to prove soundness: we show that all axioms are indeed valid, and that rules of inference preserve validity. Proving completeness, however, is trickier. There are several approaches to show completeness, and here we sketch the standard one.

The basic idea is to reason about the contraposition: if a formula is not derivable, then there is a model in which it is false. A special model that demonstrates this for all formulas is called the *canonical model*.

In order to construct the canonical model, we use maximal consistent sets.

Definition 2.15 (Maximal consistent set)**.** A set of formulas $\Gamma \subseteq \mathcal{L}_{EL}$ is a *maximal consistent set* (MCS) if $\Gamma \nvdash \bot$, and there is no $\Delta \subseteq \mathcal{L}_{EL}$ such that $\Gamma \subset \Delta$ and $\Delta \nvdash \bot$.

The next result states that every formula is an element of an MCS.

Lemma 2.5 (Lindenbaum)**.** If Γ is a consistent set, then there exists a maximal consistent set Δ such that $\Gamma \subseteq \Delta$.

We can now define the canonical model.

Definition 2.16 (Canonical model)**.** The *canonical model* $M^c = (W^c, \sim^c, V^c)$ is defined as follows:

- $W^c = \{\Gamma \mid \Gamma \text{ is an MCS}\}$,
- $\Gamma \sim_a^c \Delta$ iff $\{K_a\varphi \mid K_a\varphi \in \Gamma\} = \{K_a\varphi \mid K_a\varphi \in \Delta\}$,
- $V^c(p) = \{\Gamma \in W^c \mid p \in \Gamma\}$.

Next lemma ties together proof theory and semantics: a formula is true in a state of the canonical model (which is an MCS by Definition 2.16) if and only if it belongs to that MCS.

Lemma 2.6 (Truth)**.** For every $\varphi \in \mathcal{L}_{EL}$ and every MCS Γ, $(M^c, \Gamma) \models \varphi$ iff $\varphi \in \Gamma$.

Finally, the completeness of EL.

Theorem 2.7. For all $\varphi \in \mathcal{L}_{EL}$, $\models \varphi$ implies $\vdash \varphi$.

Proof. Suppose that $\nvdash \varphi$. Then $\{\neg\varphi\}$ is a consistent set, and it is a subset of some MCS Γ by Lindenbaum Lemma. By Truth Lemma we have that $(M^c, \Gamma) \models \neg\varphi$, and hence $\not\models \varphi$. □

2.1.6 Common and Distributed Knowledge

The EL setting allows us to reason not only about knowledge of individual agents but about knowledge of groups of agents as well. Such knowledge is called *group knowledge* and is written as $E_G\varphi$, meaning that 'everyone in G knows φ.'

Definition 2.17 (Group Knowledge)**.** Given $G \subseteq A$ and $\varphi \in \mathcal{L}_{EL}$, *group knowledge* is defined as

$$E_G\varphi := \bigwedge_{a\in G} K_a\varphi.$$

Note that $E_G\varphi$ is not necessarily an equivalence relation.

In terms of semantics, accessibility relations for group knowledge of G is a union of relations of agents from G, i.e.

$$\sim_{E_G} = \bigcup_{a\in G} \sim_a .$$

The ultimate version of the higher-order knowledge is *common knowledge*. It is common knowledge that φ if everybody knows that φ, everybody knows that everybody knows that φ, and so on. Informally, common knowledge is 'what every fool knows.'

Definition 2.18 (Common Knowledge)**.** Given $G \subseteq A$, the accessibility relation for *common knowledge* $\sim_{C_G}$ is defined as the reflexive transitive closure of $\sim_{E_G}$. Intuitively, common knowledge can be thought as an infinite conjunction

$$C_G\varphi := \bigwedge_{n=0}^{\infty} E_G^n\varphi,$$

where $E_G^n\varphi$ is a shorthand for $\overbrace{E_G \ldots E_G}^{n \text{ times}}\varphi$. However, this representation is not a formula of EL since it is infinite.

Another important notion of group knowledge is distributed knowledge. *Distributed knowledge* among a set of agents G that φ means that if the agents could combine their knowledge, they would be able to deduce φ. Informally, distributed knowledge is 'what a wise person knows.'

Definition 2.19 (Distribute Knowledge)**.** Given $G \subseteq A$, the accessibility relation for *distributed knowledge* $\sim_{D_G}$ is defined as

$$\sim_{D_G} = \bigcap_{a \in G} \sim_a .$$

Note that common and distributed knowledge cannot be expressed in standard EL, and in order to reason about them, we should add $C_G\varphi$ and $D_G\varphi$ in the definition of the language. In that case, the semantics looks as follows.

Definition 2.20. Let (M, w), $G \subseteq A$, and $\varphi \in \mathcal{L}_{EL}$ be given.

$$\begin{array}{ll} (M, w) \models E_G\varphi & \text{for all } v \in W : w \sim_{E_G} v \text{ implies } (M, v) \models \varphi \\ (M, w) \models C_G\varphi & \text{for all } v \in W : w \sim^*_{E_G} v \text{ implies } (M, v) \models \varphi \\ (M, w) \models D_G\varphi & \text{for all } v \in W : w \sim_{D_G} v \text{ implies } (M, v) \models \varphi \end{array}$$

where $\sim^*_{E_G}$ is the reflexive transitive closure of $\sim_{E_G}$.

Intuitively, if everybody knows that φ, then every a-transition from w leads to a state where φ holds. It is common knowledge that φ, if any G-path of any length ends in a φ-state, where a G-path is a sequence of a-transitions such that $a \in G$. Finally, it is distributed knowledge that φ, if a φ-state is reachable by *every* agent in G.

2.2 Public Announcement Logic

Whereas epistemic logic deals with a static distribution of knowledge of agents, *dynamic epistemic logic* [van Ditmarsch et al., 2008; Moss, 2015] is used to describe how this knowledge distribution changes as a result of various epistemic actions. The most studied dynamic epistemic operator is *public announcement.* A (truthful) public announcement can be described as an action of simultaneously informing all the agents of some true formula φ, and it is common knowledge among the agents that all of them accept φ. A logic for reasoning about public announcements, *public announcement logic* (PAL), was introduced in [Plaza, 2007]. PAL allows us to reason about how agents' knowledge changes after acquiring new information. The presentation of the section follows closely [van Ditmarsch et al., 2008, Chapter 4].

Let us return to the example in Figure 2.1. Suppose that agents bid in order to buy the item. Once one of the agents, let us say a, announces her bid, she also wants the other agent to remain ignorant of the total sum at her disposal. Formally, we can express this goal as formula $\varphi := K_b(10_a \vee 15_a) \wedge \neg(K_b 10_a \vee K_b 15_a)$ (for bid 10 by agent a). If a commits to pay 10 pounds, agent b knows that a has 10 or more pounds, but she does not know the exact amount. This condition can be achieved by a public announcement of $K_a 10_a \vee K_a 15_a$. In other words, agent a commits to pay 10 pounds, which denotes that she has at least that sum at her disposal. Formally, $(M, 15_a 5_b) \models [K_a 10_a \vee K_a 15_a]\varphi$, and the result of such announcement is shown in Figure 2.5.

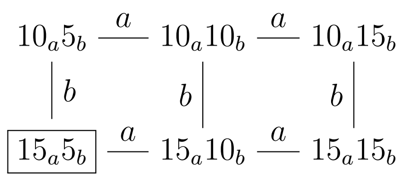

Figure 2.5: Updated model $(M, 15_a5_b)^{K_a10_a \vee K_a15_a}$

In the figure, all states that do not satisfy formula $K_a10_a \vee K_a15_a$ and all corresponding relations are removed. Note that in the original model $(M, 15_a5_b)$ formula φ was false, whereas in the updated model $(M, 15_a5_b)^{K_a10_a \vee K_a15_a}$ formula φ is true.

2.2.1 Syntax and Semantics of PAL

Definition 2.21 (Language of PAL)**.** The *language of public announcement logic* $\mathcal{L}_{PAL}$ is as follows:

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

where $p \in P$, $a \in A$, $[\varphi]$ is a public announcement, and all the usual abbreviations of propositional logic and conventions for deleting parentheses hold. The operator $\langle\varphi\rangle\psi$ is by definition $\neg[\varphi]\neg\psi$. Formulas $[\psi]\varphi$ and $\langle\psi\rangle\varphi$ are read as 'after public announcement of ψ, φ holds.'

Next, we define model updates that are results of public announcements.

Definition 2.22 (Updated model)**.** For a pointed model (M, w) and φ, an *updated model* $(M, w)^\varphi$ is a restriction of the original model to the states where φ holds and to corresponding relations. Let $[\![\varphi]\!]_M = \{w \mid (M, w) \models \varphi\}$. Then $W^\varphi = [\![\varphi]\!]_M$, $\sim_a^\varphi = \sim_a \cap ([\![\varphi]\!]_M \times [\![\varphi]\!]_M)$ for all $a \in A$, and $V^\varphi(p) = V(p) \cap [\![\varphi]\!]_M$. A model which results in subsequent updates of (M, w) with formulas $\varphi_1, \ldots, \varphi_n$ is denoted $(M, w)^{\varphi_1,\ldots,\varphi_n}$. We will also sometimes write $(M, w)^X = (W^X, \sim^X, V^X)$, where $W^X = X$, $w \in X$, $\sim_a^X = \sim_a \cap (X \times X)$ for all $a \in A$, and $V^X(p) = V(p) \cap X$.

Definition 2.23 (Semantics of PAL)**.** Let a pointed model (M, w), and $\varphi, \psi \in \mathcal{L}_{PAL}$ be given. The *semantics of* PAL is as in Definition 2.3 plus the following:

$$(M, w) \models [\psi]\varphi \quad \text{iff} \quad (M, w) \models \psi \text{ implies } (M, w)^\psi \models \varphi$$

The semantics for the dual operator is as follows:

$$(M, w) \models \langle\psi\rangle\varphi \quad \text{iff} \quad (M, w) \models \psi \text{ and } (M, w)^\psi \models \varphi$$

Note that $[\psi]\varphi$ is vacuously true if ψ is false, i.e. every φ is true after a false announcement. Also, it is easy to see that the diamond version of the public announcement operator implies the box one: $\models \langle\psi\rangle\varphi \to [\psi]\varphi$. This fact means that for a truthful public announcements there is just one deterministic outcome.

2.2.2 Axiomatisation and Completeness of PAL

Axiomatisation of PAL extends the one for EL.

Definition 2.24 (Reduction axioms for PAL)**.** Let $a \in A$, and $\varphi, \psi, \chi \in \mathcal{L}_{PAL}$. The *axiom system for* PAL includes all the axiom schemata and rules of inference from Definition 2.12 plus the following:

$[\psi]p \leftrightarrow (\psi \to p)$	atomic permanence
$[\psi]\neg\varphi \leftrightarrow (\psi \to \neg[\psi]\varphi)$	announcement and negation
$[\psi](\varphi \wedge \chi) \leftrightarrow ([\psi]\varphi \wedge [\psi]\chi)$	announcement and conjunction
$[\psi]K_a\varphi \leftrightarrow (\psi \to K_a[\psi]\varphi)$	announcement and knowledge
$[\psi][\chi]\varphi \leftrightarrow [\psi \wedge [\psi]\chi]\varphi$	announcement composition
If $\vdash \varphi$, then $\vdash [\psi]\varphi$	announcement necessitation

These axioms are *reduction axioms*, since they 'push through' the public announcement operator, and allow us to get rid of it altogether.

To prove the completeness of PAL, we define a translation function that reduces the complexity of a PAL formula within the scope of a public announcement operator. Applying this translation successively yields an EL formula that is equivalent to the original PAL one.

Definition 2.25 (Translation)**.** The *translation* $t : \mathcal{L}_{PAL} \to \mathcal{L}_{EL}$ is as follows:

$$\begin{array}{lcl}
t(p) & = & p \\
t(\neg\varphi) & = & \neg t(\varphi) \\
t(\varphi \wedge \psi) & = & t(\varphi) \wedge t(\psi) \\
t(K_a\varphi) & = & K_a t(\varphi) \\
t([\psi]p) & = & t(\psi \to p) \\
t([\psi]\neg\varphi) & = & t(\psi \to \neg[\psi]\varphi) \\
t([\psi](\varphi \wedge \chi)) & = & t([\psi]\varphi \wedge [\psi]\chi) \\
t([\psi]K_a\varphi) & = & t(\psi \to K_a[\psi]\varphi) \\
t([\psi][\chi]\varphi) & = & t([\psi \wedge [\psi]\chi]\varphi)
\end{array}$$

Lemma 2.8. For all $\varphi \in \mathcal{L}_{PAL}$, $\vdash \varphi \leftrightarrow t(\varphi)$.

This allows us to prove completeness of PAL via reduction to EL.

Theorem 2.9. PAL is sound and complete.

An alternative axiomatisation of PAL and completeness proof without using reduction axioms are presented in [Wang and Cao, 2013].

2.2.3 Expressivity and Complexity of PAL

Reduction axioms for PAL (and in particular Lemma 2.8) indicate that PAL and EL express the same properties of epistemic models. In other words, PAL and EL have the same expressive power.

Definition 2.26 (Expressivity)**.** Let two languages $\mathcal{L}_1$ and $\mathcal{L}_2$ that are interpreted in the same class of models be given. We say that $\mathcal{L}_2$ is *at least as expressive as* $\mathcal{L}_1$ (denoted as $\mathcal{L}_1 \preceq \mathcal{L}_2$) if and only if for every formula $\varphi_1 \in \mathcal{L}_1$ there is an equivalent formula $\varphi_2 \in \mathcal{L}_2$. If $\mathcal{L}_1 \preceq \mathcal{L}_2$ and $\mathcal{L}_2 \preceq \mathcal{L}_1$, then we write $\mathcal{L}_1 \equiv \mathcal{L}_2$ and say that $\mathcal{L}_1$ and $\mathcal{L}_2$ have the *same expressive power* (or that they are equally expressive). If $\mathcal{L}_2$ is *not at least as expressive as* $\mathcal{L}_1$, we write $\mathcal{L}_1 \not\preceq \mathcal{L}_2$. We say that $\mathcal{L}_1$ and $\mathcal{L}_2$ are *incomparable* if and only if $\mathcal{L}_1 \not\preceq \mathcal{L}_2$ and $\mathcal{L}_2 \not\preceq \mathcal{L}_1$.

We will, however, abuse the notation and write $L_1 \preceq L_2$, $L_1 \not\preceq L_2$, and $L_1 \equiv L_2$, for $\mathcal{L}_{L_1} \preceq \mathcal{L}_{L_2}$, $\mathcal{L}_{L_1} \not\preceq \mathcal{L}_{L_2}$, and $\mathcal{L}_{L_1} \equiv \mathcal{L}_{L_2}$ correspondingly.

Lemma 2.8 implies the following fact.

Theorem 2.10 ([Plaza, 2007])**.** PAL $\equiv$ EL, i.e. PAL and EL are equally expressive.

For other results on the expressivity of various DELs, see [van Ditmarsch et al., 2008, Chapter 8].

Although it might seem that reduction axioms make public announcement operator excessive, it was shown [Lutz, 2006; French et al., 2013] that PAL is exponentially more succinct than EL, i.e. using EL formulas instead of equivalent PAL ones generally requires exponentially more space.

Despite PAL being more succinct, the computational profile of PAL is the same as that of EL.

Theorem 2.11 ([Lutz, 2006])**.** The satisfiability problem for PAL is NP-complete in the single-agent case, and PSPACE-complete in the multi-agent case.

Theorem 2.12 ([van Benthem et al., 2006])**.** The model checking problem for PAL is in P.

2.2.4 Positive Fragment

An interesting property of public announcements is that formula φ, after being announced, does not necessarily remain true, that is $\varphi \to [\varphi]\varphi$ is not valid. The most (in)famous example of such a sentence φ is Moore sentence: $\varphi := p \land \neg K_a p$. If some fact p is true and agent a does not know it, then after this information is announced, a cannot help but learn p, making this φ false. Formally, $\models [p \land \neg K_a p]\neg(p \land \neg K_a p)$. For some formulas φ, however, $[\varphi]\varphi$ is a validity. Such formulas are called *successful*, and the fragment of PAL that contains only successful formulas is called *positive* PAL [van Ditmarsch and Kooi, 2006]. This fragment, however, does not contain *all* successful formulas: an interesting observation is that inconsistent formulas are successful.

Definition 2.27 (Positive Fragment)**.** The language $\mathcal{L}_{PAL^+}$ of *the positive fragment of public announcement logic* is defined by the following BNF:

$$\varphi ::= p \mid \neg p \mid (\varphi \land \varphi) \mid (\varphi \lor \varphi) \mid K_a\varphi \mid [\neg\varphi]\varphi,$$

where $p \in P$ and $a \in A$. *Positive fragment of* EL $\mathcal{L}_{EL^+}$ is the one without $[\neg\varphi]\varphi$.

Positive formulas remain true in all further restrictions of a model.

Definition 2.28 (Preservation)**.** Formula φ is *preserved under submodels* if for any models M_1 and M_2, $M_2 \subseteq M_1$ and $(M_1, w) \models \varphi$ implies $(M_2, w) \models \varphi$.

Theorem 2.13 ([van Ditmarsch and Kooi, 2006])**.** If $\varphi \in \mathcal{L}_{PAL^+}$, then φ is preserved under submodels.

Theorem 2.13 holds for EL^+ in both directions [Andréka et al., 1995]. The same can be said about PAL^+.

Theorem 2.14. If $\varphi \in \mathcal{L}_{PAL}$ is preserved under submodels, then φ is equivalent to some $\psi \in \mathcal{L}_{PAL^+}$.

Proof. Let φ be preserved under submodels. This formula has an equivalent EL formula $t(\varphi)$ due to the translation function (Definition 2.25). Since $t(\varphi)$ is preserved and $t(\varphi) \in \mathcal{L}_{EL}$, it has an equivalent $t(\varphi)^+ \in \mathcal{L}_{EL^+}$ [Andréka et al., 1995]. From $\mathcal{L}_{EL^+} \subset \mathcal{L}_{PAL^+}$ we conclude that $t(\varphi)^+ \in \mathcal{L}_{PAL^+}$. Thus a preserved φ has a PAL^+ equivalent. □

2.3 Coalition Logic

One of the basic formalisms for reasoning about coalitional powers in multi-agent systems is *Coalition Logic* (CL) [Pauly, 2002] that extends the language of propositional logic with the coalition operator $\langle\langle G\rangle\rangle\varphi$. Given a set of agents $G \subseteq A$ and a formula φ of coalition logic, $\langle\langle G\rangle\rangle\varphi$ means that 'agents from G have a joint strategy to guarantee that φ is true.' CL is not an epistemic logic per se, however, it predates coalition announcement logic (Section 3.3) and can be considered as a somewhat more general case of the latter.

2.3.1 Syntax and Semantics of CL

Definition 2.29 (Language of CL)**.** The *language of coalition logic* $\mathcal{L}_{CL}$ is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \langle\langle G\rangle\rangle\varphi,$$

where $p \in P$, $a \in A$, $G \subseteq A$ and all the usual abbreviations of propositional logic and conventions for deleting parentheses hold. The dual coalition operator $[\![G]\!]\varphi$ is defined as $\neg\langle\langle G\rangle\rangle\neg\varphi$ and is read 'whatever agents from G do, they cannot avoid φ.' Two special cases of coalitions are the *grand coalition* A and the *empty coalition* $\emptyset$.

The semantics of CL can be given in terms of neighbourhood models [Pacuit, 2017] and *concurrent game models* (CGM) [Ågotnes and van Ditmarsch, 2014; Ågotnes et al., 2015]. In our exposition we choose the latter, since CL coalition operators in GCM semantics are easier to contrast with CAL operators (Section 3.3).

Definition 2.30 (CGM)**.** A *concurrent game model* (CGM) M is a tuple (St, V, Act, act, out), where

- St is a non-empty set of states.
- $V : P \to \mathcal{P}(St)$ is a valuation of propositional variables.
- Act is a non-empty finite set of actions.
- $act : A \times St \to \mathcal{P}(Act)$ assigns to each player in each state a non-empty set of actions. The *set of joint actions for a coalition* G in s is $act(a, s) \times \ldots \times act(b, s)$ for $a, \ldots, b \in G$. We denote such a set as $ACT^G(s)$ and its elements as $act^G(s)$.
- out is a transition function that maps each state $s \in St$ and joint action $act^A(s)$ to a state $out(s, act^A(s)) \in St$.

Pair (M, s) is called a *pointed CGM.*

Definition 2.31 (Semantics of CL)**.** Let a pointed CGM (M, s), and $\varphi \in \mathcal{L}_{CL}$ be given. The *semantics of* CL is the same as for propositional logic plus the following:

$$\begin{aligned}(M, w) \models \langle\langle G \rangle\rangle\varphi \quad \text{iff} \quad & \exists act^G(s) \in ACT^G(s), \forall act^{A \setminus G}(s) \in ACT^{A \setminus G}(s) : \\ & (M, out(s, act^G(s) \cup act^{A \setminus G}(s))) \models \varphi\end{aligned}$$

where $act^G(s) \cup act^{A \setminus G}(s)$ is $(\alpha_a, \ldots, \alpha_b)$ for $a, \ldots, b \in A$.

Thus operator $\langle\langle G \rangle\rangle\varphi$ should be read as 'G has an action to achieve φ no matter what other agents $A \setminus G$ do.'

Let us consider an example of a CGM presented in Figure 2.6. Three researchers, a, b, and c, decide to spend some of their research budget on an espresso machine. Currently they have none in their office (state ~~`coffee`~~), and they are deliberating between two options (states `coffee I` and `coffee II`). In order to choose the espresso machine, researchers have to vote. Thus, the set of available actions $Act = \{v_I, v_{II}, e\}$, the set of all joint actions is $\{(v_I, v_I, v_I)$, (v_I, v_I, v_{II}), (v_I, v_{II}, v_I), (v_{II}, v_I, v_I), (v_{II}, v_{II}, v_{II}), (v_{II}, v_{II}, v_I), (v_{II}, v_I, v_{II}), (v_I, v_{II}, v_{II}), $(e, e, e)\}$, where (v_I, v_I, v_{II}) means that a and b vote for the first machine, and c votes for the second one, and (e, e, e) means that all of the agents are enjoying coffee. Note that only one action is possible in `coffee I` and `coffee II`.

Assume that some p holds in `coffee I`, some q holds in `coffee II`, and both p and q are false in ~~`coffee`~~. Hence we have that $(M, \sout{\texttt{coffee}}) \models \neg p \wedge \neg q$. The grand coalition consisting of all agents can choose either espresso machine, i.e. $(M, \sout{\texttt{coffee}}) \models \langle\langle \{a, b, c\} \rangle\rangle p \wedge \langle\langle \{a, b, c\} \rangle\rangle q$. However, there is not enough budget to have both, $(M, \sout{\texttt{coffee}}) \models [\![\{a, b, c\}]\!] \neg(p \wedge q)$. In this scenario it is also possible for two agents to team up and force some outcome no matter what the third agent does. Formally,

$$(M, \sout{\texttt{coffee}}) \models \bigwedge \begin{pmatrix} \langle\langle \{a, b\} \rangle\rangle p \wedge \langle\langle \{a, b\} \rangle\rangle q \\ \langle\langle \{a, c\} \rangle\rangle p \wedge \langle\langle \{a, c\} \rangle\rangle q \\ \langle\langle \{b, c\} \rangle\rangle p \wedge \langle\langle \{b, c\} \rangle\rangle q \end{pmatrix}.$$

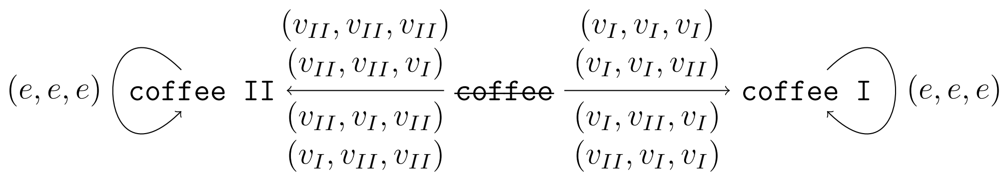

Figure 2.6: Researchers and Espresso Machines

Finally, a single agent cannot force any outcome in the best-of-three voting, i.e.

$$(M, \sout{\texttt{coffee}}) \models \bigwedge \begin{pmatrix} [\![\{a\}]\!]p \wedge [\![\{a\}]\!]q \\ [\![\{b\}]\!]p \wedge [\![\{b\}]\!]q \\ [\![\{c\}]\!]p \wedge [\![\{c\}]\!]q \end{pmatrix}.$$

2.3.2 Axiomatisation of CL

A sound and complete axiomatisation of CL was provided in [Pauly, 2002].

Definition 2.32 (Axiomatisation of CL)**.** Let $G, H \subseteq A$ and $\varphi, \psi \in \mathcal{L}_{CL}$. The *axiom system for* CL is presented below.

propositional tautologies,	
$\neg\langle\!\langle G\rangle\!\rangle\bot$	liveness,
$\langle\!\langle G\rangle\!\rangle\top$	safety,
$\neg\langle\!\langle\emptyset\rangle\!\rangle\neg\varphi \to \langle\!\langle A\rangle\!\rangle\varphi$	A-maximality,
$\langle\!\langle G\rangle\!\rangle(\varphi \wedge \psi) \to \langle\!\langle G\rangle\!\rangle\varphi$	outcome-monotonicity,
$\langle\!\langle G\rangle\!\rangle\varphi \wedge \langle\!\langle H\rangle\!\rangle\psi \to \langle\!\langle G \cup H\rangle\!\rangle(\varphi \wedge \psi)$, if $G \cap H = \emptyset$	superadditivity,
If $\vdash \varphi$ and $\varphi \to \psi$, then $\vdash \psi$	modus ponens,
If $\vdash \varphi \leftrightarrow \psi$, then $\vdash \langle\!\langle G\rangle\!\rangle\varphi \leftrightarrow \langle\!\langle G\rangle\!\rangle\psi$	equivalence.

Intuitively, liveness expresses the fact that every action of a coalition leads to some outcome, and safety assumes that every coalition has some choice. A-maximality indicates the relationship between the empty and grand coalitions. According to outcome-monotonicity, if a coalition can force some outcome, then the coalition can force any superset of that outcome. Superadditivity allows disjoint coalitions to combine their choices.

Theorem 2.15 ([Pauly, 2002])**.** CL is sound and complete.

Finally, we mention the complexity results for CL.

Theorem 2.16 ([Pauly, 2002])**.** The satisfiability problem for CL is PSPACE-complete.

Theorem 2.17 ([Alur et al., 2002])**.** The model checking problem for CL is in P.

In this section we do not consider an epistemic variant of CL. Such an extension of the logic with individual, distributed, and common knowledge operators has been studied in [Ågotnes and Alechina, 2012].

Chapter 3

Logics of Quantified Public Announcements

One of the ways to generalise PAL is to allow for quantification over public announcements. Such a generalisation was first implemented in [Balbiani et al., 2008], where the authors introduce Arbitrary Public Announcement Logic (Section 3.2) with operators for quantification over announcements of epistemic formulas. Restricting such a quantification to knowledge formulas of agents gives rise to Group Announcement Logic (Section 3.1). A formalism that allows us to consider not only announcements by a group of agents but counter-announcements by their opponents as well, is Coalition Announcement Logic (Section 3.3). This logic, as opposed to group and arbitrary announcement logics, entertains a more game-theoretic setting, where some coalition may not have a strategy to achieve their goal in the presence of the anti-coalition, even though the goal is achievable in the non-competitive setting. In Section 3.4 we present a worked example for group and coalition announcement logics.

3.1 Group Announcement Logic

Group Announcement Logic (GAL) [Ågotnes and van Ditmarsch, 2008; Ågotnes et al., 2010] is an extension of PAL with group announcement modalities $[G]\varphi$ (and its dual $\langle G\rangle\varphi$). Alternatively, GAL can be considered as a restriction of *Arbitrary Public Announcement Logic* (APAL) [Balbiani et al., 2008]: instead of quantifying over *all* epistemic formulas, we quantify over formulas *known by agents in a group* G. Although APAL predates GAL, we start our discussion with the latter since APAL is not in the primary focus of our work.

Formula $\langle G\rangle\varphi$ should be read as 'there is a truthful announcement by agents from group G such that φ holds after that announcement.' In this context a truthful announcement means that agents actually know formulas they announce. In other words, announcement of φ_a by agent $a \in G$ is interpreted as $K_a\varphi$. Similarly, $[G]\varphi$ is read 'whatever agents from group G announce, φ holds afterwards.'

In Figure 2.5 model $(M, 15_a5_b)$ is updated with the announcement $K_a10_a \vee K_a15_a$. Coincidentally, the same model is the result of updating $(M, 15_a5_b)$ with

$K_a(10_a \vee 15_a)$. Note that this formula is a knowledge formula of agent a. Hence, we can conclude that *there is* an announcement by agent a such that φ is true afterwords; from $(M, 15_a5_b) \models \langle K_a(10_a \vee 15_a)\rangle\varphi$ it follows that $(M, 15_a5_b) \models \langle\{a\}\rangle\varphi$. Announcement by groups of more than one agent are conjunctions of corresponding announcements by agents.

3.1.1 Syntax and Semantics of GAL

As usual, P is a countable set of propositional variables, and A is a finite set of agents.

Definition 3.1 (Language of GAL)**.** The *language of group announcement logic* $\mathcal{L}_{GAL}$ is inductively defined as

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [\varphi]\varphi \mid [G]\varphi,$$

where $p \in P$, $a \in A$, $G \subseteq A$ and all the usual abbreviations of propositional logic and conventions for deleting parentheses hold. The operator $\langle G\rangle\varphi$ is by definition $\neg[G]\neg\varphi$.

Let us denote by ψ_G formula $\bigwedge_{a\in G} K_a\psi_a$ such that $\psi_a \in \mathcal{L}_{EL}$. We refer to this fragment as $\mathcal{L}_{EL}^G$. So ψ_G's are shorthands for conjunctions of knowledge formulas of agents from group G.

Definition 3.2 (Semantics of GAL)**.** Let a pointed model (M, w), and $\varphi \in \mathcal{L}_{GAL}$ be given. The *semantics of* GAL is as in Definition 2.23 plus the following:

$$(M, w) \models [G]\varphi \quad \text{iff} \quad \forall\psi_G : (M, w) \models [\psi_G]\varphi$$

The semantics for the dual operator is as follows:

$$(M, w) \models \langle G\rangle\varphi \quad \text{iff} \quad \exists\psi_G : (M, w) \models \langle\psi_G\rangle\varphi$$

In the definition of the semantics, the group announcement operator $[G]$ only quantifies over epistemic formulas (and hence, by Theorem 2.10, over PAL formulas) known to agents. This allows us to avoid circularity in the definition. More on this issue and alternative semantics for quantified announcements is in [van Ditmarsch et al., 2016].

It is easy to see that the following proposition holds.

Proposition 3.1. Let some $\varphi \in \mathcal{L}_{GAL}$ be given. Then $\models [G]\varphi$ if and only if for all ψ_G it holds that $\models [\psi_G]\varphi$.

Note that generalisations of Proposition 3.1 to $[\psi][G]\varphi$, $K_a[G]\varphi$, and $\psi \to [G]\varphi$ also hold. To address this variety of possible rules of inference, we consider a more succinct way of their representation. Ultimately we require premises to be expressions of depth n of the type

$$\varphi_1 \to \Box_1(\varphi_2 \to \ldots(\varphi_n \to \Box_n\sharp)\ldots),$$

where $\Box_i$ is either K_a or $[\psi]$ for some $a \in A$ and $\psi \in \mathcal{L}_{GAL}$, atom $\sharp$ denotes a placement of a formula to which a derivation is applied, and some φ's and $\Box$'s can be omitted. This condition is captured by *necessity forms* originally introduced by Goldblatt in [Goldblatt, 1982, Chapter 2] (under the name *admissibility forms*).

Definition 3.3. (Necessity forms) Let $\varphi \in \mathcal{L}_{GAL}$, then *necessity forms* are inductively defined as follows:

$$\eta(\sharp) ::= \sharp \mid \varphi \to \eta(\sharp) \mid K_a\eta(\sharp) \mid [\varphi]\eta(\sharp).$$

The atom $\sharp$ has a unique occurrence in each necessity form. The result of the replacement of $\sharp$ with φ in some $\eta(\sharp)$ is denoted as $\eta(\varphi)$. The dual of a necessity form $\eta(\varphi)$ is a *possibility form* $\eta\{\varphi\}$ that is defined as $\eta(\varphi) ::= \neg\eta\{\neg\varphi\}$.

3.1.2 Axiomatisation of GAL

In this section we present an axiomatisation of GAL.

Definition 3.4 (Axiomatisation of GAL)**.** Let $a \in A$, $G \subseteq A$, $\psi_G \in \mathcal{L}^G_{EL}$, and $\varphi, \psi, \chi \in \mathcal{L}_{GAL}$. The *axiom system for* GAL is an extension of PAL and is presented below.

$(A0)$	propositional tautologies,
$(A1)$	$K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$,
$(A2)$	$K_a\varphi \to \varphi$,
$(A3)$	$K_a\varphi \to K_aK_a\varphi$,
$(A4)$	$\neg K_a\varphi \to K_a\neg K_a\varphi$,
$(A5)$	$[\varphi]p \leftrightarrow (\varphi \to p)$,
$(A6)$	$[\varphi]\neg\psi \leftrightarrow (\varphi \to \neg[\varphi]\psi)$,
$(A7)$	$[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$,
$(A8)$	$[\varphi]K_a\psi \leftrightarrow (\varphi \to K_a[\varphi]\psi)$,
$(A9)$	$[\varphi][\psi]\chi \leftrightarrow [\varphi \wedge [\varphi]\psi]\chi$,
$(A10)$	$[G]\varphi \to [\psi_G]\varphi$,
$(R0)$	If $\vdash \varphi$ and $\vdash \varphi \to \psi$, then $\vdash \psi$,
$(R1)$	If $\vdash \varphi$, then $\vdash K_a\varphi$,
$(R2)$	If $\vdash \varphi$, then $\vdash [\psi]\varphi$,
$(R3)$	If $\vdash \varphi$, then $\vdash [G]\varphi$,
$(R4)$	If $\forall\psi_G : \vdash \eta([\psi_G]\varphi)$, then $\vdash \eta([G]\varphi)$.

Rule $R4$ makes the axiomatisation infinitary. Whereas finitary rules require a finite number of premises, *infinitary rule* $R4$ requires an infinite number of formulas $\eta([\psi_G]\varphi)$ to be theorems of GAL in order to conclude that $\eta([G]\varphi)$ is a theorem of GAL as well. Hence, in what follows, we override Definition 2.12 so that $\vdash \varphi$ means that φ belongs to the smallest subset of $\mathcal{L}_{GAL}$ that contains all the axioms $A0 - A10$ and closed under rules of inference $R0 - R4$. Such a set is called GAL, and we refer to elements of GAL as *theorems*. Set GAL is infinite, and if $\eta([\psi_G]\varphi)$ are all theorems, then $\eta([G]\varphi)$ is also a theorem. Note that for finitary systems, both definitions of a theorem – as an element of a set, and as a final line in a derivation — are equivalent.

The completeness of the axiom system for GAL is due to [Ågotnes et al., 2010; Balbiani et al., 2008; Balbiani and van Ditmarsch, 2015; Balbiani, 2015].

Theorem 3.2. The axiom system for GAL is sound and complete.

3.1.3 Complexity and Logical Properties

Let us mention some interesting validities of GAL.

Proposition 3.3 ([Ågotnes et al., 2010])**.** All of the following are valid.

1. $\langle G \rangle p \leftrightarrow p$,
2. $\langle \emptyset \rangle \varphi \leftrightarrow [\emptyset]\varphi \leftrightarrow \varphi$,
3. $\varphi \to \langle G \rangle \varphi$,
4. $\langle G \rangle \langle H \rangle \varphi \to \langle G \cup H \rangle \varphi$,
5. $\langle G \rangle \langle G \rangle \varphi \leftrightarrow \langle G \rangle \varphi$.

The first validity states that agents' announcements cannot alter propositional variables. The second one expresses the property that the empty group is powerless. A group variant of the truth axiom is presented in 3. According to 4, if consecutive announcements by groups achieve φ, then φ can be achieved by a single joint announcement. The validity of the other direction of 4 used to be an open question. In Chapter 5 we settle it by presenting a counterexample. Property 5 is a corollary of 4. Even though 4 and 5 reduce sequences of announcements to a single announcement, in some scenarios agents may be unaware of the consequences of their joint announcements. This leads to a distinction between *being able* to execute some protocol and *knowing how* to execute it. For the discussion see [Ågotnes et al., 2010].

Proposition 3.4. Let $G, H \subseteq A$, and $\varphi, \psi \in \mathcal{L}_{GAL}$. Property $(\langle G \rangle \varphi \land \langle H \rangle \psi) \to \langle G \cup H \rangle (\varphi \land \psi)$ is not a validity of GAL.

Proof. Consider some pointed model (M, w) such that $(M, w) \models \langle \{a\} \rangle K_c p \land \langle \{b\} \rangle \neg K_c p$. Trivially, this does not imply $(M, w) \models \langle \{a, b\} \rangle (K_c p \land \neg K_c p)$. □

Invalid property in Proposition 3.4 is the superadditivity axiom of CL. This implies that GAL is not a coalition logic in a sense that there are some validities of CL that are not valid in GAL [Ågotnes and van Ditmarsch, 2014]. This is due to the fact that GAL operators do not take into account agents outside of a given group (if there are any). A logic with quantified announcements that is a coalition logic is considered in Section 3.3.

Now we present the complexity profile of GAL.

Theorem 3.5 ([Ågotnes et al., 2016])**.** The satisfiability problem for GAL is undecidable.

Theorem 3.6 ([Ågotnes et al., 2010])**.** The model checking problem for GAL is PSPACE-Complete.

Finally, we state a known expressivity result for GAL: GAL is as expressive as EL in the single-agent case, and strictly more expressive in the multi-agent case.

Theorem 3.7. GAL $\equiv$ EL in the single-agent case, and EL $\preceq$ GAL, and GAL $\not\preceq$ EL in the multi-agent case.

3.2 Arbitrary Public Announcement Logic

The expansion of the domain of quantification in GAL to the set of *all* epistemic formulas, rather than knowledge formulas of agents, results in another well-known formalism — *Arbitrary Public Announcement Logic* (APAL). Modalities of APAL, $\Box\varphi$ and $\Diamond\varphi$, mean that 'after all (some) announcements of epistemic formulas, φ holds.' From a logical perspective, APAL is quite similar to GAL, and hence our exposition of the logic is rather condensed.

3.2.1 Syntax, Semantics, and Axiomatisation of APAL

Definition 3.5 (Language of APAL)**.** The *language of group announcement logic* $\mathcal{L}_{GAL}$ is inductively defined as

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [\varphi]\varphi \mid \Box\varphi,$$

where $p \in P$, $a \in A$, and all the usual abbreviations of propositional logic and conventions for deleting parentheses hold. The operator $\Diamond\varphi$ is by definition $\neg\Box\neg\varphi$.

Definition 3.6 (Semantics of APAL)**.** Let a pointed model (M, w), and $\varphi \in \mathcal{L}_{GAL}$ be given. The *semantics of* APAL is as in Definition 2.23 plus the following:

$$(M, w) \models \Box\varphi \quad \text{iff} \quad \forall\psi \in \mathcal{L}_{EL} : (M, w) \models [\psi]\varphi$$

The semantics for the dual operator is as follows:

$$(M, w) \models \Diamond\varphi \quad \text{iff} \quad \exists\psi \in \mathcal{L}_{EL} : (M, w) \models \langle\psi\rangle\varphi$$

Similarly to the case of GAL, the quantification is restricted to epistemic formulas.

Definition 3.7 (Axiomatisation of APAL)**.** The *axiom system for* APAL is the same as the one for GAL with the following exceptions:

$$\begin{array}{ll} (A10) & \Box\varphi \rightarrow [\psi]\varphi, \text{ where } \psi \in \mathcal{L}_{EL}, \\ (R3) & \text{If } \vdash \varphi, \text{ then } \vdash \Box\varphi, \\ (R4) & \text{If } \forall\psi \in \mathcal{L}_{EL} : \vdash \eta([\psi]\varphi), \text{ then } \vdash \eta(\Box\varphi). \end{array}$$

APAL was shown to be complete in [Balbiani et al., 2008; Balbiani and van Ditmarsch, 2015; Balbiani, 2015].

Theorem 3.8. The axiom system for APAL is sound and complete.

3.2.2 Expressivity and Complexity of APAL

Regarding the relation between APAL and EL, we have exactly the same result as for GAL.

Theorem 3.9 ([Balbiani et al., 2008])**.** In the single-agent case APAL $\equiv$ EL, and APAL $\not\preceq$ EL and EL $\preceq$ APAL in the multi-agent case.

When APAL and GAL are compared to each other, it turns out that GAL is not at least as expressive as APAL, i.e. there are some properties of epistemic models expressible in APAL that cannot be expressed in GAL.

Theorem 3.10 ([Ågotnes et al., 2010])**.** Generally, APAL $\not\preceq$ GAL. In the special case when an agent with the identity relation is present, APAL $\preceq$ GAL.

Whether GAL $\not\preceq$ APAL was posed as an open question in [Ågotnes et al., 2010]. We solve this problem in Chapter 7.

The complexity profile of APAL is the same as of GAL.

Theorem 3.11 ([Ågotnes et al., 2016])**.** The satisfiability problem for APAL is undecidable.

Theorem 3.12 ([Balbiani et al., 2008])**.** The model checking problem for GAL is PSPACE-Complete.

3.3 Coalition Announcement Logic

Coalition Announcement Logic (CAL) was proposed in [Ågotnes and van Ditmarsch, 2008] as a generalisation of PAL. CAL modalities $[\langle G\rangle]$ and $\langle[G]\rangle$, as opposed to the GAL ones, are interpreted as double quantifications of the type $\forall\exists$ and $\exists\forall$ over epistemic formulas known by agents. Thus, $[\langle G\rangle]\varphi$ means 'whatever agents from G announce, there is an announcement by agents from $A \setminus G$ such that φ holds afterwards. And its dual, $\langle[G]\rangle\varphi$, is read as 'there is a truthful announcement made by the agents in G such that no matter what the agents not in G simultaneously announce, φ holds afterwards.'

CAL modalities were motivated by coalition logic [Pauly, 2002] and van Benthem's playability operator [van Benthem, 2014, Chapter 11] and [van Benthem, 2001]. Among other logics of quantified announcements — APAL and GAL — CAL is the least studied one.

Returning to the example in Figure 2.5, not only is φ true after a truthful announcement $K_a(10_a \vee 15_a)$ by a, but also it is true *no matter what* agent b announces at the same time. Formally, we write $(M, 15_a5_b) \models \langle[\{a\}]\rangle\varphi$. A more complex example is presented Section 3.4.

Definition 3.8 (Language of CAL)**.** The *language of coalition announcement logic* $\mathcal{L}_{CAL}$ is by the following BNF:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [\varphi]\varphi \mid [\langle G\rangle]\varphi,$$

where $p \in P$, $a \in A$, $G \subseteq A$ and all the usual abbreviations of propositional logic and conventions for deleting parentheses hold. The operator $\langle\!\langle G \rangle\!\rangle\varphi$ is by definition $\neg[\langle G \rangle]\neg\varphi$.

As in the case of GAL, we denote by ψ_G formula $\bigwedge_{a \in G} K_a \psi_a$ such that $\psi_a \in \mathcal{L}_{EL}$, and refer to this fragment as $\mathcal{L}_{EL}^G$.

Definition 3.9 (Semantics of CAL)**.** Let a pointed model (M, w), and $\varphi \in \mathcal{L}_{CAL}$ be given. The *semantics of* CAL is as in Definition 2.23 plus the following:

$$(M, w) \models [\langle G \rangle]\varphi \quad \text{iff} \quad \forall \psi_G, \exists \chi_{A \setminus G} : (M, w) \models \psi_G \\ \text{implies } (M, w) \models \langle \psi_G \wedge \chi_{A \setminus G} \rangle \varphi$$

The semantics for the dual operator is as follows:

$$(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi \quad \text{iff} \quad \exists \psi_G, \forall \chi_{A \setminus G} : (M, w) \models \psi_G \text{ and } (M, w) \models [\psi_G \wedge \chi_{A \setminus G}]\varphi$$

In order to avoid circularity in the definition above, operators $[\langle G \rangle]$ and $\langle\!\langle G \rangle\!\rangle$ quantify only over epistemic formulas known to agents.

Interestingly, the two special cases of the grand coalition A and the empty coalition $\emptyset$ are duals. Let us consider their semantics:

$$(M, w) \models \langle\!\langle A \rangle\!\rangle\varphi \quad \text{iff} \quad \exists \psi_A : (M, w) \models \langle \psi_A \rangle \varphi$$
and
$$(M, w) \models \langle\!\langle \emptyset \rangle\!\rangle\varphi \quad \text{iff} \quad \forall \psi_A : (M, w) \models [\psi_A] \varphi$$

This definition is used in the following proposition.

Proposition 3.13. All of the following is valid.

1. $\langle\!\langle G \rangle\!\rangle p \leftrightarrow p$,
2. $[\langle A \rangle]\varphi \leftrightarrow \langle\!\langle \emptyset \rangle\!\rangle\varphi$,
3. $\varphi \to \langle\!\langle A \rangle\!\rangle\varphi$.

Property 1 states that agents cannot change values of propositional variables. This follows from the fact that the only available action in CAL, public announcement, is a purely epistemic action. The second formula captures the semantics of grand and empty coalitions. Finally, property 3 is a CAL variant of the truth axiom. Note that 3 is not valid for $G \subsetneq A$.

CAL is a coalition logic: axioms and rules of inference of CL are valid or validity preserving in CAL [Ågotnes and van Ditmarsch, 2008] (see Section 5.3 for the proof). Moreover, CAL can express some properties that are not valid in CL, e.g. formula 1 above.

Other known results for CAL include the following theorems.

Theorem 3.14 ([Ågotnes et al., 2016])**.** The satisfiability problem for CAL is undecidable.

The complexity of the model checking for CAL was an open question. In Chapter 4 we show that it is PSPACE-complete.

Theorem 3.15. CAL $\equiv$ EL in the single-agent case, and EL $\preceq$ CAL and CAL $\not\preceq$ EL in the multi-agent case.

To the best of our knowledge, relative expressivity of CAL and other logics of quantified announcements has not been studied yet. We address this open problem in Chapter 7, and demonstrate that GAL $\not\preceq$ CAL and APAL $\not\preceq$ CAL.

3.4 Households and Burglars: An Example

In city N. local authorities decided to gather information about and present statistics on electricity consumption. This information should be submitted by each neighbourhood in the city, indicating the total number of households that have been using electricity last month. Data about neighbourhoods is public, and data about individual households is private, i.e. particular users of electricity are not revealed, but the total number of such users in the area is common knowledge. And there is a reason for such a requirement.

A group of local burglars is also interested in the public report on electricity consumption: they hope to deduce which households have not used electricity recently since it is an indication that property owners are not in their houses these days (most probably, they are on vacation). Burglars, however, want to be quite certain that the house of their choice is empty, and hence they would not risk burglary unless they know for sure that the property owners are away. Also, they are highly reluctant to lurk around a neighbourhood trying to learn who is away, such a behaviour is quite suspicious. Thus, their only way to know about 'vacant' households is through the public report.

In the city, there is a small neighbourhood of only four houses: a, b, c, and d. They are situated around a park in a circular fashion such that neighbours on the left and on the right are equidistant. The park is quite large and dwellers of the houses tend to know only their closest neighbours on the left and on the right. Thus, for example, owner of c knows owners of b and d, and about their plans, but she is unaware about plans of owners of a.

Epistemic model TES describing the given situation is presented in Figure 3.1. In the model, state-names indicate which owners are at home; for instance, 1001 means that a and d are at home, and that b and c are not. Burglars v (for 'villains') do not have any information regarding home-owners, and their epistemic relation is universal. We do not present the relation in the figure for readability.

Let the actual state be 0101, and let 0101 also abbreviate $\neg p_a \wedge p_b \wedge \neg p_c \wedge p_d$, where p_i stands for 'owner i is at home.' Note that neither burglars nor the owners possess the full information about the neighbourhood: $(TES, 0101) \models \neg(K_a 0101 \vee K_b 0101 \vee K_c 0101 \vee K_d 0101 \vee K_v 0101)$. Also note that dwellers are aware of their own states and states of their left- and right-hand-side neighbours,

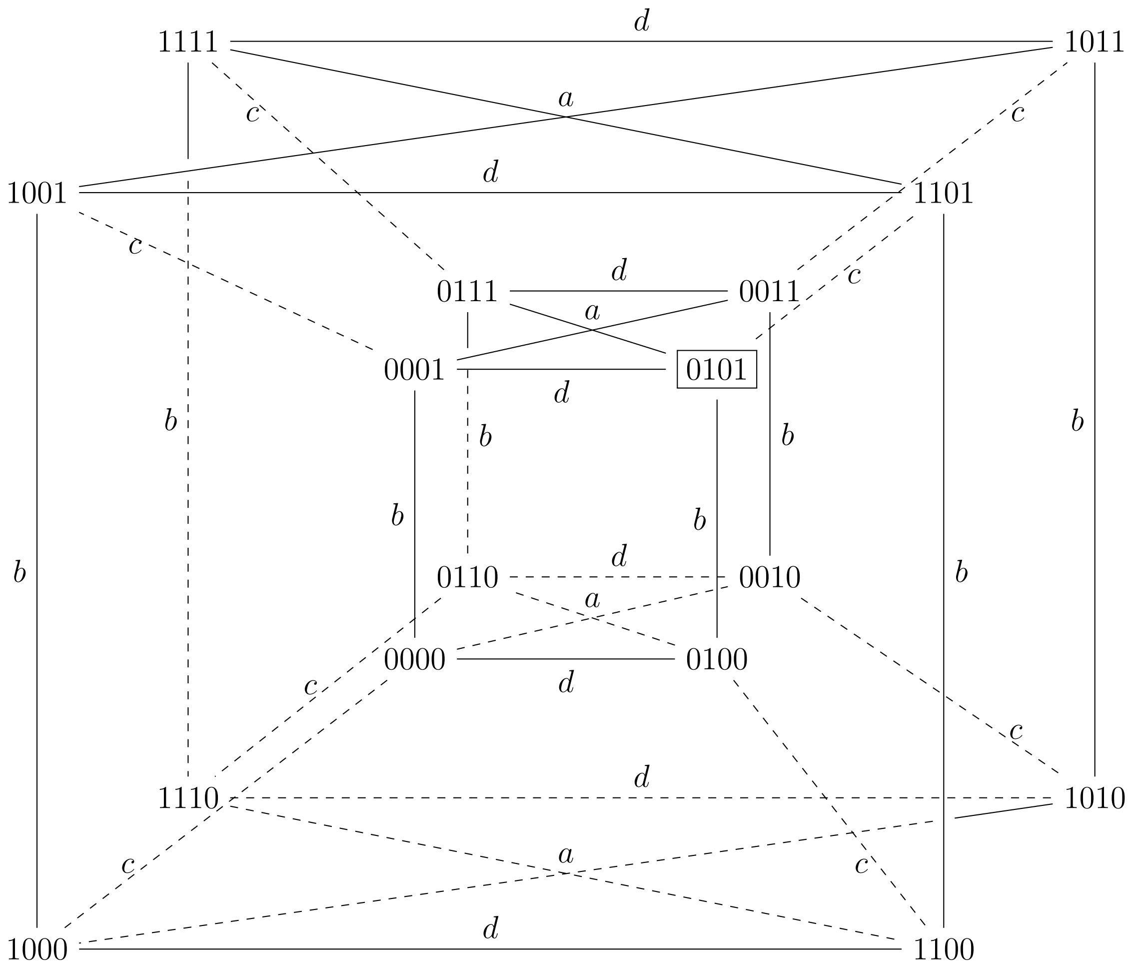

Figure 3.1: Model $(TES, 0101)$

but not about a state of the farthest owner. E.g. $(TES, 0101) \models K_c\neg p_c \wedge K_c p_b \wedge K_c p_d \wedge \neg(K_c\neg p_a \vee K_c p_a)$.

Information agents a, b, c and d want to submit is 'two household in our neighbourhood have been using electricity these days.' This sentence, however, should conform to the requirement that exact households remain unknown to the public outside the neighbourhood. We can express this goal as formula

$$sofa := K_v \bigvee \begin{pmatrix} p_a \wedge p_b \wedge \neg p_c \wedge \neg p_d \\ p_a \wedge \neg p_b \wedge p_c \wedge \neg p_d \\ p_a \wedge \neg p_b \wedge \neg p_c \wedge p_d \\ \neg p_a \wedge p_b \wedge p_c \wedge \neg p_d \\ \neg p_a \wedge p_b \wedge \neg p_c \wedge p_d \\ \neg p_a \wedge \neg p_b \wedge p_c \wedge p_d \end{pmatrix} \wedge \neg \bigvee \begin{pmatrix} K_v p_a \vee K_v \neg p_a \\ K_v p_b \vee K_v \neg p_b \\ K_v p_c \vee K_v \neg p_c \\ K_v p_d \vee K_v \neg p_d \end{pmatrix},$$

where $sofa$ stands for 'the state of affairs.' A successful group announcement by agents to achieve this goal is possible when everyone announces 'I know that if I have not been using electricity recently, then at least one of my neighbours on the left and on the right has, and if I have been using it, then one of the neighbours must be on vacation.' Formally, such an announcement can be expressed by the following formula:

$$mis := \bigwedge \begin{pmatrix} K_a((\neg p_a \to (p_d \lor p_b)) \land (p_a \to \neg(p_d \land p_b))) \\ K_b((\neg p_b \to (p_a \lor p_c)) \land (p_b \to \neg(p_a \land p_c))) \\ K_c((\neg p_c \to (p_b \lor p_d)) \land (p_c \to \neg(p_b \land p_d))) \\ K_d((\neg p_d \to (p_c \lor p_a)) \land (p_d \to \neg(p_c \land p_a))) \end{pmatrix},$$

where mis stands for 'mutual informative statement.'

Thus we have that $(TES, 0101) \models \langle mis \rangle sofa$. Since mis is an announcement of agents' knowledge, we can conclude that there is an announcement by a, b, c and d such that $sofa$ holds in the resulting model, i.e. $(TES, 0101) \models \langle \{a, b, c, d\} \rangle sofa$. Result of updating $(TES, 0101)$ with mis is presented in Figure 3.2.

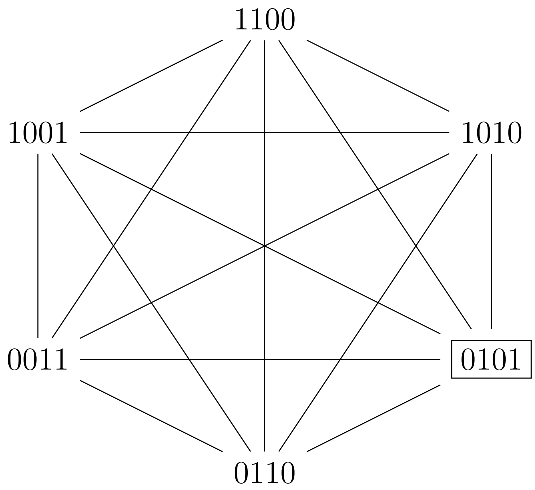

Figure 3.2: Model $(TES, 0101)^{mis}$

All the relations in the model are v equivalence relations. Hence, indeed, in $(TES, 0101)^{mis}$ exactly two households have been using electricity recently, and although the public (and burglars as well) knows that fact, they cannot name particular houses that are 'vacant.' A 'side-effect' of group announcement mis is that all residents in the neighbourhood know exactly who is on vacation, and it is common knowledge.

Note that we can state a fact stronger than $(TES, 0101) \models \langle \{a, b, c, d\} \rangle sofa$. Since v's relation is universal, they cannot preclude the group to make $sofa$ true whatever they (i.e. v) announce. In other words, $(TES, 0101) \models \langle\!\langle \{a, b, c, d\} \rangle\!\rangle sofa$.

Interestingly, in this particular example even two agents can make an announcement such that $sofa$ holds in the resulting model. Consider the following announcement by agents a and b:

$$mis_{a,b} := K_a((p_a \to \neg p_d) \land (\neg p_a \to p_d)) \land K_b((p_b \to \neg p_c) \land (\neg p_b \to p_c)).$$

The resulting updated model is shown in Figure 3.3 (all the relations are v-relations).

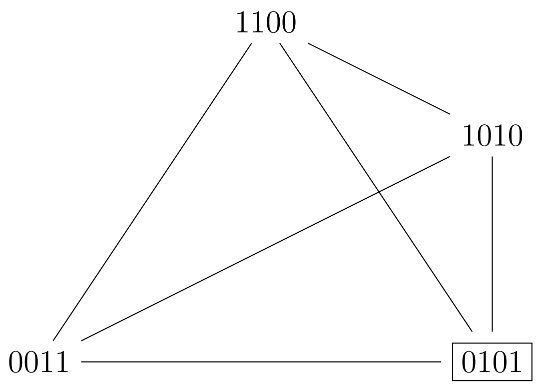

Figure 3.3: Model $(TES, 0101)^{mis_{a,b}}$

The reader can verify that $(TES, 0101)^{mis_{a,b}} \models sofa$, and hence $(TES, 0101) \models \langle\{a, b\}\rangle sofa$. Note that compared to model $(TES, 0101)^{mis}$ (Figure 3.2), model $(TES, 0101)^{mis_{a,b}}$ has fewer states. This means that owners gave a bit more information than necessary, but they still managed to inform authorities that exactly two households have been using electricity while not revealing the exact state of affairs.

Even though two owners can make a successful announcement, they must ensure that none of the other agents has not been conspiring with burglars. For assume this is the case that agent c, for example, decides to reveal to burglars which houses are empty. She can pass the following information with a's and b's submission: $K_c(\neg p_c \land p_b \land p_d)$. This announcement made in conjunction with $mis_{a,b}$ results in a singleton model with 0101 as the only state. Moreover, whatever a and b announce, c always has an announcement to make $sofa$ false in the resulting model (and, alas, to let the burglars know that she is on vacation). Formally, we have that $(TES, 0101) \models \neg\langle\!\langle\{a, b\}\rangle\!\rangle sofa$, or, equivalently, $(TES, 0101) \models [\![\{a, b\}]\!]\neg sofa$. Hence, in this particular example, property owners should always cooperate if they want to inform authorities about electricity consumption and keep the burglars away.

We have seen that an announcement by two owners is enough to make $sofa$ true. What about the single-agent case? As owners possess information about themselves and two closest neighbours, they do not know the actual state of the world, i.e. they do not have enough information about their farthest neighbour. However, it is possible for some agents to make an announcement such that it informs the public that at least two of the households have been using electricity recently, and particular users and non-users remain incognito. Formally, such a target formula is as follows:

$$sofa_1 := K_v \bigvee \begin{pmatrix} p_a \wedge p_b \wedge \neg p_c \wedge \neg p_d \\ p_a \wedge \neg p_b \wedge p_c \wedge \neg p_d \\ p_a \wedge \neg p_b \wedge \neg p_c \wedge p_d \\ \neg p_a \wedge p_b \wedge p_c \wedge \neg p_d \\ \neg p_a \wedge p_b \wedge \neg p_c \wedge p_d \\ \neg p_a \wedge \neg p_b \wedge p_c \wedge p_d \\ p_a \wedge p_b \wedge p_c \wedge \neg p_d \\ p_a \wedge p_b \wedge \neg p_c \wedge p_d \\ p_a \wedge \neg p_b \wedge p_c \wedge p_d \\ \neg p_a \wedge p_b \wedge p_c \wedge p_d \\ p_a \wedge p_b \wedge p_c \wedge p_d \end{pmatrix} \wedge \neg \bigvee \begin{pmatrix} K_v p_a \vee K_v \neg p_a \\ K_v p_b \vee K_v \neg p_b \\ K_v p_c \vee K_v \neg p_c \\ K_v p_d \vee K_v \neg p_d \end{pmatrix}.$$

Agent a, for instance, can make $sofa_1$ true in $(TES, 0101)$ by announcing

$$mis_a := K_a((\neg p_a \rightarrow (p_b \wedge p_d)) \wedge (p_a \rightarrow (p_b \vee p_d))).$$

The result of such an announcement is presented in Figure 3.4 (all relations are v-relations).

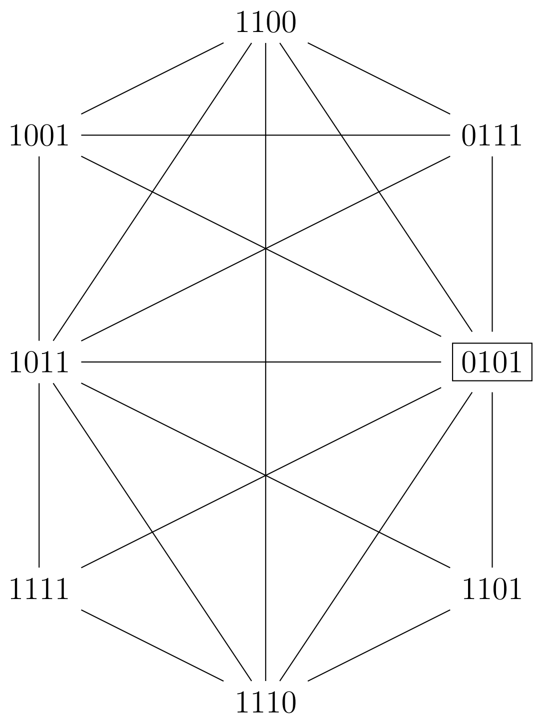

Figure 3.4: Model $(TES, 0101)^{mis_a}$

It easy to verify that $(TES, 0101)^{mis_a} \models sofa_1$. Hence, it also holds that $(TES, 0101) \models \langle\{a\}\rangle sofa_1$, and, obviously, $(TES, 0101) \models \neg\langle\!\langle\{a\}\rangle\!\rangle sofa_1$.

Chapter 4

Model Checking for CAL

The complexity of the model checking problem for GAL is known to be PSPACE-complete [Ågotnes et al., 2010]. In this chapter we show that the same holds for CAL. Since both $[\langle G\rangle]\varphi$ and $\langle[G]\rangle\varphi$ are defined using quantification over an infinite number of epistemic formulas, direct application of the semantics of CAL in the model checking algorithm is not possible. In Section 4.1 we employ distinguishing formulas to make a transition from an infinite number of possible announcements to a finite number of strategies available to a coalition of agents. This allows us to present the model checking algorithm for CAL in Section 4.2, and prove the complexity result. Moreover, we demonstrate that if a formula within the scope of a coalition announcement operator is a positive PAL or GAL formula, then the complexity of the model checking problem is in P. This chapter is based on [Galimullin et al., 2018].

4.1 Strategies of Groups of Agents on Finite Models

4.1.1 Distinguishing Formulas

In this section we introduce distinguishing formulas that are satisfied in only one (up to bisimulation) state in a finite model. The discussion is based on [van Ditmarsch et al., 2014]. Although agents know and can possibly announce an infinite number of formulas, using distinguishing formulas allows us to consider only finitely many different announcements. This is done by associating strategies of agents with corresponding distinguishing formulas, where a strategy of agent a is a union of a-equivalence classes.

Here and subsequently, all epistemic models are *finite* and *bisimulation contracted*. Also, without loss of generality, we assume that the set of propositional variables P is finite. It follows from the fact that in a finite epistemic model $M = (W, \sim, V)$ there are $2^{|W|}$ unique truth assignments for a propositional variable, and a truth assignment for any $p_{2^{|W|}+1}$ will repeat one from $p_1, \ldots, p_{2^{|W|}}$.

We continue with the formal definition of distinguishing formulas.

Definition 4.1 (Distinguishing Formula)**.** Let a finite epistemic model M be given. Formula $\delta_{S,S'}$ is called *distinguishing* for $S, S' \subseteq W$ if $S \subseteq [\![\delta_{S,S'}]\!]_M$ and $S' \cap [\![\delta_{S,S'}]\!]_M = \emptyset$. If a formula distinguishes state w from all other non-bismilar states in M, we write δ_w.

Proposition 4.1 (van Ditmarsch et al. [2014])**.** Let a finite epistemic model M be given. Every pointed model (M, w) is distinguished from all other non-bisimilar pointed models (M, v) by some distinguishing formula $\delta_w \in \mathcal{L}_{EL}$.

Given a finite model (M, w), a distinguishing formula δ_w is constructed recursively as follows:

$$\delta_w^{k+1} := \delta_w^0 \wedge \bigwedge_{a \in A} (\bigwedge_{w \sim_a v} \widehat{K}_a \delta_v^k \wedge K_a \bigvee_{w \sim_a v} \delta_v^k),$$

where $0 \leq k < |W|$, and δ_w^0 is the conjunction of all literals that are true in w, i.e. $\delta_w^0 := \bigwedge_{w \in V(p)} p \wedge \bigwedge_{w \notin V(p)} \neg p$. So, formula $\delta_w := \delta_w^{|W|}$.

Assumptions regarding some given model being finite and bisimulation contracted are of vital importance for the construction of distinguishing formulas. If the model is infinite, then we may either need an infinite amount of propositional variables to describe the given state, or there may be infinite branches of accessibility relations. If the model is not bisimulation contracted, i.e. there are bisimilar states in the model, then distinguishing formulas cease to be unique — the same formula describes all bisimilar states in the model.

Having defined distinguishing formulas for states, we can define distinguishing formulas for sets of states.

Definition 4.2. Let some finite and bisimulation contracted model (M, w), and a set S of states in M be given. A *distinguishing formula for S* is

$$\delta_S := \bigvee_{w \in S} \delta_w.$$

Let us recall the bidding example from Section 2.1.1, and construct distinguishing formula $\delta_{15_a 5_b}$.

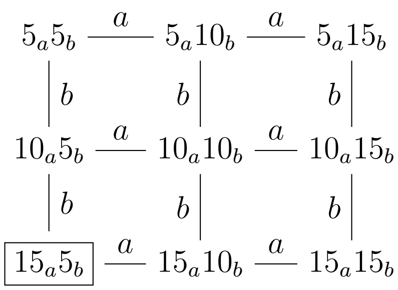

Figure 4.1: Model $(M, 15_a 5_b)$

First, we start with the propositional description of the state:

$$\delta^0_{15_a5_b} := 15_a \wedge 5_b \wedge \neg 10_a \wedge \neg 5_a \wedge \neg 10_b \wedge \neg 15_b.$$

Let us assume that we calculated δ^0's in the same fashion for all other states. Next, we proceed with the first iteration of $\delta_{15_a5_b}$:

$$\delta^1_{15_a5_b} := \delta^0_{15_a5_b} \wedge \bigwedge \left(\begin{array}{c} \widehat{K}_a\delta^0_{15_a5_b} \wedge \widehat{K}_a\delta^0_{15_a10_b} \wedge \widehat{K}_a\delta^0_{15_a15_b} \wedge K_a \left(\bigvee \begin{array}{c} \delta^0_{15_a5_b} \\ \delta^0_{15_a10_b} \\ \delta^0_{15_a15_b} \end{array} \right) \\ \widehat{K}_b\delta^0_{15_a5_b} \wedge \widehat{K}_b\delta^0_{10_a5_b} \wedge \widehat{K}_b\delta^0_{5_a5_b} \wedge K_b \left(\bigvee \begin{array}{c} \delta^0_{15_a5_b} \\ \delta^0_{10_a5_b} \\ \delta^0_{5_a5_b} \end{array} \right) \end{array} \right).$$

The process continues for $|W|$ iterations. Informally, each iteration of a distinguishing formula construction adds one step of 'reachable distance' from a given state. Hence, in our example with 9 states distinguishing formula $\delta_{15_a5_b}$ looks as follows (assuming that all previous $\delta^k_{15_a5_b}$'s have been calculated):

$$\delta^9_{15_a5_b} := \delta^0_{15_a5_b} \wedge \bigwedge \left(\begin{array}{c} \widehat{K}_a\delta^8_{15_a5_b} \wedge \widehat{K}_a\delta^8_{15_a10_b} \wedge \widehat{K}_a\delta^8_{15_a15_b} \wedge K_a \left(\bigvee \begin{array}{c} \delta^8_{15_a5_b} \\ \delta^8_{15_a10_b} \\ \delta^8_{15_a15_b} \end{array} \right) \\ \widehat{K}_b\delta^8_{15_a5_b} \wedge \widehat{K}_b\delta^8_{10_a5_b} \wedge \widehat{K}_b\delta^8_{5_a5_b} \wedge K_b \left(\bigvee \begin{array}{c} \delta^8_{15_a5_b} \\ \delta^8_{10_a5_b} \\ \delta^8_{5_a5_b} \end{array} \right) \end{array} \right).$$

Note that since models we are dealing with in this chapter are finite, distinguishing formulas always exist.

4.1.2 Strategies

In this section we introduce strategies and connect them to public announcements using distinguishing formulas. Intuitively, strategies are sets of states that agents can ensure to be in the updated model by announcing a formula that holds in those states. The formal definition of a strategy is presented below.

Definition 4.3 (Strategies)**.** Let $M/a = \{[w_1]_a, \ldots, [w_n]_a\}$ be the set of a-equivalence classes in M. A *strategy* X_a for an agent a in a finite model (M, w) is a union of equivalence classes of a including $[w]_a$. *The set of all available strategies* of a is $S(a, w) = \{[w]_a \cup X_a : X_a \subseteq M/a\}$. *Group strategy* X_G is defined as $\bigcap_{a \in G} X_a$ for all $a \in G$. The set of available strategies for a group of agents G is $S(G, w) = \{\bigcap_{a \in G} X_a : X_a \in S(a, w)\}$.

Strategies can only be implemented by agents, and generally public announcements do not correspond to strategies. Consider model $(M, 15_a5_b)$ in Figure 4.1, and an a-equivalence class $\{15_a5_b, 15_a10_b, 15_a15_b\}$. This equivalence class is also a strategy X_a. It is easy to see that public announcement of $(15_a \wedge$

$5_b) \vee (15_a \wedge 10_b)$ does not correspond to any strategy of a, that is for all X_a: $[\![(15_a \wedge 5_b) \vee (15_a \wedge 10_b)]\!] \neq X_a$. Implementing a strategy X_a results in an updated model where agent a knows some φ, and in our example it is not the case that $K_a((15_a \wedge 5_b) \vee (15_a \wedge 10_b))$.

Note, that for any (M, w) and $G \subseteq A$, $S(G, w)$ is not empty, since the trivial strategy that includes all the states of the current model is available to all agents.

Proposition 4.2. In a finite model (M, w), for any $G \subseteq A$, $S(G, w)$ is finite.

Proof. Due to the fact that in a finite model there is a finite number of equivalence classes for each agent. □

Thus, in Figure 4.1 there are three a-equivalence classes: $\{15_a 5_b, 15_a 10_b, 15_a 15_b\}$, $\{10_a 5_b, 10_a 10_b, 10_a 15_b\}$, and $\{5_a 5_b, 5_a 10_b, 5_a 15_b\}$. Let us designate them by the first element of a corresponding set, i.e. $\mathtt{15_a5_b}$, $\mathtt{10_a5_b}$, and $\mathtt{5_a5_b}$. The set of all available strategies of agent a in $(M, 15_a 5_b)$ is $\{\mathtt{15_a5_b}, \mathtt{15_a5_b} \cup \mathtt{10_a5_b}, \mathtt{15_a5_b} \cup \mathtt{5_a5_b}, \mathtt{15_a5_b} \cup \mathtt{10_a5_b} \cup \mathtt{5_a5_b}\}$. Similarly, the set of all available strategies of agent b in $(M, 15_a 5_b)$ is $\{\mathtt{15_a5_b}, \mathtt{15_a5_b} \cup \mathtt{15_a10_b}, \mathtt{15_a5_b} \cup \mathtt{15_a15_b}, \mathtt{15_a5_b} \cup \mathtt{15_a10_b} \cup \mathtt{15_a15_b}\}$. Finally, there is a group strategy for agents a and b that contains only two states – $15_a 5_b$ and $10_a 5_b$. This strategy is an intersection of a's $\mathtt{15_a5_b} \cup \mathtt{10_a5_b}$ and b's $\mathtt{15_a5_b}$, that is $\{15_a 5_b, 15_a 10_b, 15_a 15_b, 10_a 5_b, 10_a 10_b, 10_a 15_b\} \cap \{15_a 5_b, 10_a 5_b, 5_a 5_b\}$.

Now we tie together announcements and strategies. Each of infinitely many possible announcements by agents in a finite model corresponds to a set of states where it is true (a strategy). In a finite bisimulation contracted model, each strategy is definable by a distinguishing formula, hence it corresponds to an announcement. This allows us to consider finitely many strategies instead of considering infinitely many possible announcements: there are only finitely many non-equivalent (in terms of model updates) announcements for each finite model, and each of them is equivalent to a distinguishing formula of some strategy.

Given a finite and bisimulation contracted model (M, w) and strategy X_G, a distinguishing formula δ_{X_G} for X_G can be obtained from Definition 4.2 as $\bigvee_{w \in X_G} \delta_w$.

Next, we show that agents know their strategies and thus can make corresponding announcements.

Proposition 4.3. Let agent a have strategy X_a in some finite bisimulation contracted (M, w). Then $(M, w) \models K_a \delta_{X_a}$. Also, let $X_G := X_a \cap \ldots \cap X_b$ be a strategy, then $(M, w) \models K_a \delta_{X_a} \wedge \ldots \wedge K_b \delta_{X_b}$, where $a, \ldots, b \in G$.

Proof. We show just the first part of the proposition, since the second part follows easily. By the definition of a strategy, $X_a = [w_1]_a \cup \ldots \cup [w_n]_a$ for some $[w_1]_a, \ldots, [w_n]_a \in M/a$. For every equivalence class $[w_i]_a$ there is a corresponding distinguishing formula $\delta_{[w_i]_a}$. Since for all $v \in [w_i]_a$, $(M, v) \models \delta_{[w_i]_a}$ (by Proposition 4.1 and Definition 4.2), we have that $(M, v) \models K_a \delta_{[w_i]_a}$. The same holds for other equivalence classes of a including the one with w, and we have $(M, w) \models (K_a \delta_{[w_1]_a} \vee \ldots \vee K_a \delta_{[w_n]_a})$, which implies $(M, w) \models K_a(\delta_{[w_1]_a} \vee \ldots \vee \delta_{[w_n]_a})$. Note that $\delta_{[w_1]_a} \vee \ldots \vee \delta_{[w_n]_a}$ is a distinguishing formula of strategy X_a, so we can write

$(M, w) \models K_a\delta_{X_a}$. Finally, having defined $K_a\delta_{X_a}, \ldots, K_b\delta_{X_b}$ for all $a, \ldots, b \in G$, group strategy X_G in (M, w) corresponds to $(M, w) \models K_a\delta_{X_a} \wedge \ldots \wedge K_b\delta_{X_b}$. □

The following proposition (which follows from Propositions 4.1 and 4.3) states that given a strategy, corresponding public announcement yields exactly the model with states specified by the strategy.

Proposition 4.4. *Given a finite bisimulation contracted model $M = (W, \sim, V)$ and a strategy X_a, $W^{K_a\delta_{X_a}} = X_a$. More generally, $W^{K_a\delta_{X_a}\wedge\ldots\wedge K_b\delta_{X_b}} = X_G$, where $a, \ldots, b \in G$.*

We also show that true group announcements correspond to group strategies.

Proposition 4.5. *Given some model (M, w) and formula ψ_G such that $(M, w) \models \psi_G$, there is a strategy X_G such that $X_G = W^{\psi_G}$.*

Proof. Assume that $(M, w) \models \psi_G$. Formula ψ_G is an abbreviation for $K_a\psi_a \wedge \ldots \wedge K_b\psi_b$, where $a, \ldots, b \in G$ and $\psi_a \in \mathcal{L}_{EL}$. Let us consider $K_a\psi_a$. By the semantics we have that $(M, w) \models K_a\psi_a$ holds if and only if for all v reachable from w via a, $(M, v) \models \psi_a$. Note that all states reachable from the given one via a form an a-equivalence class $[w]_a$. In the same way $K_a\psi_a$ may be true in other a-equivalence classes $[u]_a, \ldots [t]_a$. Hence, formula $K_a\psi_a$ holds in the union of these equivalence classes, i.e. it holds in $W^{K_a\psi_a} = [w]_a \cup \ldots \cup [t]_a$. By Definition 4.3, $[w]_a \cup \ldots \cup [t]_a$ is a strategy X_a of agent a.

Now assume that we have defined strategies of all $a, \ldots, b \in G$ in this fashion. Note that $W^{K_a\psi_a\wedge\ldots\wedge K_b\psi_b} = W^{K_a\psi_a} \cap \ldots \cap W^{K_b\psi_b}$, and for all agents i, $W^{K_i\psi_i} = X_i$. Hence we have that $W^{K_a\psi_a\wedge\ldots\wedge K_b\psi_b} = W^{K_a\psi_a} \cap \ldots \cap W^{K_b\psi_b} = X_a \cap \ldots \cap X_b$, and the latter is a group strategy X_G. □

Now, let us reformulate semantics for the group and coalition announcement operators in terms of strategies.

Proposition 4.6. *For a finite bisimulation contracted model (M, w) we have that*

$$\begin{array}{lll} (M, w) \models \langle G\rangle\varphi & \text{iff} & \exists X_G \in S(G, w) : (M, w)^{X_G} \models \varphi, \\ (M, w) \models \langle\!\langle G\rangle\!\rangle\varphi & \text{iff} & \exists X_G \in S(G, w)\ \forall X_{A\setminus G} \in S(A \setminus G, w) : (M, w)^{X_G \cap X_{A\setminus G}} \models \varphi. \end{array}$$

Proof. Case $\langle G\rangle\varphi$: From left to right. Assume that for some pointed model we have that $(M, w) \models \langle G\rangle\varphi$. By the semantics this means that $\exists\psi_G$ such that $(M, w) \models \langle\psi_G\rangle\varphi$. Formula ψ_G is true in some set of states that is an intersection of unions of agents' equivalence classes. Note that this set is some strategy X_G, that is $W^{\psi_G} = X_G$ (by Proposition 4.5).

From right to left. Let there be some strategy X_G such that $(M, w)^{X_G} \models \varphi$, then, by Propositions 4.3 and 4.4, there is an announcement of distinguishing formulas by agents from G such that $X_G = W^{K_a\delta_{X_a}\wedge\ldots\wedge K_b\delta_{X_b}}$, and hence $(M, w)^{K_a\delta_{X_a}\wedge\ldots\wedge K_b\delta_{X_b}} \models \varphi$ and $(M, w) \models K_a\delta_{X_a} \wedge \ldots \wedge K_b\delta_{X_b}$. The latter is equivalent to $(M, w) \models \langle G\rangle\varphi$ by the semantics.

Case $\langle\!\langle G \rangle\!\rangle\varphi$: *From left to right.* Suppose that for some (M, w) it holds that $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi$. By the definition of semantics this is equivalent to $\exists\psi_G, \forall\chi_{A\setminus G}$: $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi$. As in the previous case, ψ_G correspond to some group strategy X_G such that $W^{\psi_G} = X_G$. The same holds for announcements by $A \setminus G$. The only interesting case to consider is a false announcement by that coalition, i.e. some $\chi_{A\setminus G}$ such that $(M, w) \not\models \chi_{A\setminus G}$. This announcement, however, makes $(M, w) \models [\psi_G \wedge \chi_{A\setminus G}]\varphi$ trivially true. Hence, it is enough to consider just 'meaningful' true announcements by $A \setminus G$. So, for each formula $\psi_G \wedge \chi_{A\setminus G}$ there is a strategy $X_G \cap X_{A\setminus G}$ (by Proposition 4.5) such that $W^{\psi_G \wedge \chi_{A\setminus G}} = X_G \cap X_{A\setminus G}$.

From right to left. Assume that there is some strategy X_G such that for all strategies $X_{A\setminus G}$ it holds that $(M, w)^{X_G \wedge X_{A\setminus G}} \models \varphi$. By Propositions 4.3 and 4.4 this means that for some $K_G\delta_{X_G}$ and all $K_{A\setminus G}\delta_{X_{A\setminus G}}$ it holds that $(M, w) \models K_G\delta_{X_G} \wedge K_{A\setminus G}\delta_{X_{A\setminus G}}$ and $(M, w)^{K_G\delta_{X_G} \wedge K_{A\setminus G}\delta_{X_{A\setminus G}}} \models \varphi$. Note that in a model there is always an infinite number of formulas such that $(M, w) \models K_a\psi_a \leftrightarrow K_a\chi_a$ [Ågotnes and van Ditmarsch, 2011]. Therefore, there is an infinite amount of possible equivalent announcements for agents from $A \setminus G$ of the form $\chi_{A\setminus G}$. So, we have that for some $K_G\delta_{X_G}$ and for all $\chi_{A\setminus G}$ it holds that $(M, w) \models K_G\delta_{X_G} \wedge \chi_{A\setminus G}$ and $(M, w)^{K_G\delta_{X_G} \wedge \chi_{A\setminus G}} \models \varphi$. We can relax the assumption of $\chi_{A\setminus G}$ being true, and rewrite the latter as $(M, w) \models K_G\delta_{X_G}$ and $(M, w) \models \chi_{A\setminus G}$ implies $(M, w)^{K_G\delta_{X_G} \wedge \chi_{A\setminus G}} \models \varphi$, which is equivalent to $(M, w) \models K_G\delta_{X_G} \wedge [K_G\delta_{X_G} \wedge \chi_{A\setminus G}]\varphi$ for all $\chi_{A\setminus G}$, and this is $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi$ by the semantics. □

So, we have tied together announcements and strategies via distinguishing formulas. From now on, we may abuse notation and write M^{X_G}, meaning that M^{X_G} is an update of model M by a joint announcement of agents G that corresponds to strategy X_G.

Sometimes we may be interested in situations where it is beneficial for agents to be as informative as possible (or, equivalently, leave as little uncertainty as possible). The type of announcements that fulfil those requirements is defined in Definition 4.4.

Definition 4.4 (Maximally Informative Announcement)**.** Let some finite bisimulation contracted model (M, w) and G be given. A *maximally informative announcement* is a formula ψ_G such that $w \in W^{\psi_G}$ and for all ψ'_G such that $w \in W^{\psi'_G}$ it holds that $W^{\psi_G} \subseteq W^{\psi'_G}$. For finite models such an announcement always exists [Ågotnes and van Ditmarsch, 2011]. We will call the corresponding strategy X_G *the strongest strategy* on a given model.

Intuitively, the strongest strategy is the smallest available strategy. Note that in a bisimulation contracted model (M, w), the strongest strategy of agents G is $X_G = [w]_a \cap \ldots \cap [w]_b$ for $a, \ldots, b \in G$, that is agents' strategies consist of the single equivalence classes that include the current state.

In model $(M, 15_a5_b)$ in Figure 4.1 a's strongest strategy is $\{15_a5_b, 15_a10_b, 15_a15_b\}$, and b's strongest strategy is $\{15_a5_b, 10_a5_b, 5_a5_b\}$. So, the strongest strategy of group $\{a, b\}$ is the intersection of strongest strategies of agents from the group:

$\{15_a5_b, 15_a10_b15_a15_b\} \cap \{15_a5_b, 10_a5_b, 5_a5_b\} = \{15_a5_b\}$. Corresponding announcements are, respectively, $K_a(\delta_{15_a5_b} \vee \delta_{15_a10_b} \vee \delta_{15_a15_b})$, $K_b(\delta_{15_a5_b} \vee \delta_{10_a5_b} \vee \delta_{5_a5_b})$, and $K_a(\delta_{15_a5_b} \vee \delta_{15_a10_b} \vee \delta_{15_a15_b}) \wedge K_b(\delta_{15_a5_b} \vee \delta_{10_a5_b} \vee \delta_{5_a5_b})$.

4.2 Model Checking Algorithm for CAL

Employing strategies allows for a rather simple model checking algorithm for CAL. We switch from quantification over an infinite number of epistemic formulas, to quantification over a finite set of strategies (Section 4.2.1). Moreover, we show that if the target formula is a positive PAL or GAL formula, then model checking is even more effective (Section 4.2.2).

4.2.1 General Case

Algorithm 1 takes a finite model M, a state w of the model, and some $\varphi_0 \in \mathcal{L}_{CAL}$ as an input, and returns $true$ if φ_0 is satisfiable in the model, and $false$ otherwise. Strictly speaking, for completeness' sake, the algorithm works for a language that includes *both* group and coalition announcements.

Algorithm 1: $mc(M, w, \varphi_0)$

1: **case** φ_0**:**
2: p : **if** $w \in V(p)$ **then return** $true$ **else return** $false$;
3: $\neg\varphi$: **if** $mc(M, w, \varphi)$ **then return** $false$ **else return** $true$;
4: $\varphi \wedge \psi$: **if** $mc(M, w, \varphi) \wedge mc(M, w, \psi)$ **then return** $true$ **else return** $false$;
5: $K_a\varphi$: **for all** v such that $w \sim_a v$
  if $\neg mc(M, v, \varphi)$ **then return** $false$
  return $true$;
6: $[\psi]\varphi$: compute the ψ-submodel M^ψ of M
  if $w \in W^\psi$ **then return** $mc(M^\psi, w, \varphi)$ **else return** $true$;
7: $\langle G\rangle\varphi$: compute $(\|M\|, w)$ and the set of strategies $S(G, w)$
  for all $X_G \in S(G, w)$
    if $mc(\|M\|^{X_G}, w, \varphi)$ **then return** $true$
  return $false$;
8: $\langle\!\langle G\rangle\!\rangle\varphi$: compute $(\|M\|, w)$ and sets of strategies $S(G, w)$ and $S(A \setminus G, w)$
  for all $X_G \in S(G, w)$
    $check = true$
    for all $X_{A\setminus G} \in S(A \setminus G, w)$
      if $\neg mc(\|M\|^{X_G \cap X_{A\setminus G}}, w, \varphi)$ **then** $check = false$
    if $check$ **then return** $true$
  return $false$.

Now we show the correctness of the algorithm.

Proposition 4.7. Let (M, w) and $\varphi \in \mathcal{L}_{CAL}$ be given. Algorithm $mc(M, w, \varphi)$ returns *true* iff $(M, w) \models \varphi$.

Proof. By a straightforward induction on the complexity of φ. We use Proposition 4.6 to prove the case for $\langle\!\langle G \rangle\!\rangle$:
$\Rightarrow$: Suppose $mc(M, w, \langle\!\langle G \rangle\!\rangle\varphi)$ returns *true*. By line 8 this means that for some strategy X_G and all strategies $X_{A\setminus G}$, $mc(\|M\|^{X_G \cap X_{A\setminus G}}, w, \varphi)$ returns *true*. By the Induction Hypothesis, $(\|M\|, w)^{X_G \cap X_{A\setminus G}} \models \varphi$ for some X_G and all $X_{A\setminus G}$, and $(\|M\|, w) \models \langle\!\langle G \rangle\!\rangle\varphi$ by the semantics.

$\Leftarrow$: Let $(\|M\|, w) \models \langle\!\langle G \rangle\!\rangle\varphi$, which means that there is some strategy X_G such that for all $X_{A\setminus G}$, $(\|M\|, w)^{X_G \cap X_{A\setminus G}} \models \varphi$. By the Induction Hypothesis, the latter holds iff for some X_G and for all $X_{A\setminus G}$, $mc(\|M\|^{X_G \cap X_{A\setminus G}}, w, \varphi)$ returns *true*. By line 8, we have that $mc(\|M\|, w, \langle\!\langle G \rangle\!\rangle\varphi)$ returns *true*. □

Proposition 4.8. Model checking for CAL is PSPACE-complete.

Proof. All the cases of the model checking algorithm apart from cases for $\langle G \rangle\varphi$ and $\langle\!\langle G \rangle\!\rangle$ require polynomial time (and polynomial space as a consequence). Moreover, computing the bisimulation contraction of a given model is known to be in P [Paige and Tarjan, 1987]. Cases for $\langle G \rangle$ and $\langle\!\langle G \rangle\!\rangle$ iterate over exponentially many strategies. However each iteration can be computed using only polynomial amount of space to represent $(\|M\|, w)$ (which contains at most the same number of states as the input model M) and the result of the update (which is a submodel of $(\|M\|, w)$) and make a recursive call to check whether φ holds in the update. By reusing space for each iteration, we can compute the cases for $\langle G \rangle$ and $\langle\!\langle G \rangle\!\rangle$ using only a polynomial amount of space.

Hardness can be obtained by a slight modification of the proof of PSPACE-hardness of the model-checking problem for GAL in [Ågotnes et al., 2010]. The proof encodes satisfiability of a quantified boolean formula (QBF) as a problem whether a particular GAL formula is true in a model corresponding to the QBF. We highlight just some parts of the proof from [Ågotnes et al., 2010]. Given some QBF $\Psi := Q_1 x_1 \ldots Q_n x_n \Phi(x_1, \ldots, x_n)$, the authors construct a model that depends on the number of variables in the formula (see Figure 4.2).

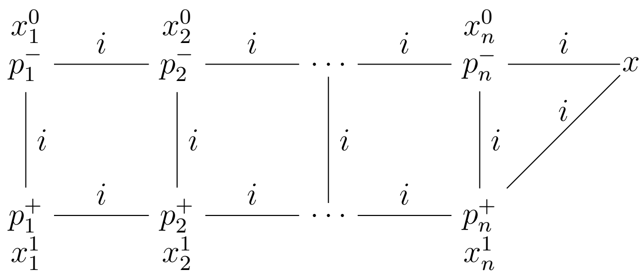

Figure 4.2: Model M that corresponds to a QBF.

Apart from agent i, whose relation is universal, there is also agent g, whose relation is an identity. Next, the authors define properties q_j 'only one x_j, upper or lower, is in the model' and r_j 'x_j is in both upper and lower rows.' These properties are used to recursively define a GAL formula $\psi(\Psi)$ that will be then evaluated in model $(M, x) \models \psi(\Psi)$. An example of a corresponding GAL formula for the given QBF $\forall x_1 \exists x_2 \forall x_3 : \Phi(x_1, x_2, x_3)$ is $K_i[g](q_1 \wedge r_2 \wedge r_3 \rightarrow \widehat{K}_i\langle g\rangle(q_1 \wedge q_2 \wedge r_3 \wedge K_i[g](q_1 \wedge q_2 \wedge q_3 \rightarrow \Phi(\widehat{K}_i p_1^+, \widehat{K}_i p_2^+, \widehat{K}_i p_3^+))))$.

For our proof, however, it is enough to notice the following. Since the encoding uses only two agents: an omniscient g and a universal i, we can replace $[g]$ and $\langle g\rangle$ with $[\langle g\rangle]$ and $\langle[g]\rangle$ (since i's only strategy is equivalent to $\top$ and no other GAL operators are used in the encoding) and obtain a CAL encoding. □

4.2.2 Positive Case

In this section we demonstrate the following result: if in a given formula of $\mathcal{L}_{CAL}$ (or the language extended with group announcements) subformulas within scopes of coalition announcement operators are positive PAL (or GAL) formulas, then the complexity of the model checking is polynomial.

Allowing coalition announcement modalities to bind only positive formulas is a natural restriction. Positive formulas have a special property: if the sum of knowledge of agents in G (their distributed knowledge) includes a positive formula φ, then φ can be made common knowledge by a group or coalition announcement by G. Formally, for a positive φ, $(M, w) \models D_G\varphi$ implies $(M, w) \models \langle[G]\rangle C_G\varphi$. See [van Ditmarsch et al., 2018; van Ditmarsch and Kooi, 2006], and also [Ågotnes and Wáng, 2017] where this is called *resolving* distributed knowledge. In other words, positive epistemic formulas can always be resolved through cooperative communication.

Negative formulas, however, do not have this property. For example, it can be distributed knowledge of agents a and b that p and $\neg K_b p$: $D_{\{a,b\}}(p \wedge \neg K_b p)$. However, it is impossible to achieve common knowledge of this formula: $C_{\{a,b\}}(p \wedge \neg K_b p)$ is inconsistent, since it implies both $K_b p$ and $\neg K_b p$. Going back to the example in Section 4.1, it is distributed knowledge of a and b that $K_a 15_a$ and $K_b 5_b$. Both formulas are positive and can be made common knowledge if a and b honestly report the amount of money they have. However, it is also distributed knowledge that $\neg K_a 5_b$ and $\neg K_b 15_a$. The conjunction

$$K_a 15_a \wedge K_b 5_b \wedge \neg K_a 5_b \wedge \neg K_b 15_a$$

is distributed knowledge, but it cannot be made common knowledge for the same reasons as above.

The positive fragment of PAL was presented in Section 2.2.4. Here we present the positive fragment of GAL.

Definition 4.5 (Positive Fragment)**.** The language $\mathcal{L}_{GAL^+}$ of *the positive fragment of group announcement logic* is defined by the following BNF:

$$\varphi ::= p \mid \neg p \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid K_a\varphi \mid [\neg\varphi]\varphi \mid [G]\varphi,$$

where $p \in P$, $a \in A$, and $G \subseteq A$.

Recall the definition of preserved formulas.

Definition 4.6 (Preservation)**.** Formula φ is *preserved under submodels* if for any models M_1 and M_2, $M_2 \subseteq M_1$ and $(M_1, w) \models \varphi$ implies $(M_2, w) \models \varphi$.

A known result that we use in this section states that formulas of $\mathcal{L}_{PAL^+}$ are preserved under submodels [van Ditmarsch and Kooi, 2006]. Let us show that this is also true about formulas of $\mathcal{L}_{GAL^+}$.

Proposition 4.9. Let M_1, M_2 with $M_2 \subseteq M_1$, $w \in W_2$ and a positive formula $\varphi \in \mathcal{L}_{GAL^+}$ be given. If $(M_1, w) \models \varphi$, then $(M_2, w) \models \varphi$.

Proof. We extend the corresponding proof from [Balbiani et al., 2008].

Case $[G]\varphi$: Suppose towards a contradiction that $(M_1, w) \models [G]\varphi$ and $(M_2, w) \not\models [G]\varphi$. Then, there is a formula ψ_G such that $(M_2, w) \models \langle \psi_G \rangle \neg\varphi$. By the semantics we have $(M_2, w)^{\psi_G} \not\models \varphi$. Since $M_2^{\psi_G} \subseteq M_2 \subseteq M_1$, and by contraposition of the Induction Hypothesis, we have that $(M_1, w) \not\models \varphi$. But at the same time, $(M_1, w) \models [G]\varphi$ implies $(M_1, w) \models [K_G\top]\varphi$, where $K_G\top := K_a\top \land \ldots \land K_b\top$ and $a, \ldots, b \in G$. The latter is equivalent to $(M_1, w) \models \varphi$. Hence, we have a contradiction. □

Proposition 4.10. $[G]\varphi \leftrightarrow \varphi$ is valid for all positive φ.

Proof. Left-to-right direction holds due to the fact that G can announce $\top_G$, and right-to-left direction holds due to Proposition 4.9. □

Since positive fragments are normally defined to be universal fragments, and both CAL modalities contain existential quantification over formulas, it would appear that CAL modalities cannot occur in its own positive fragment.

Next, we show the following interesting fact.

Proposition 4.11. $\langle\!\langle G \rangle\!\rangle\varphi \leftrightarrow [\![\langle A \backslash G \rangle]\!]\varphi$ is valid for positive φ on finite bisimulation contracted models.

Proof. The left-to-right direction is generally valid and we omit the proof (see Proposition 5.7). Suppose that $(M, w) \models [\![\langle A \setminus G \rangle]\!]\varphi$. By Proposition 4.6 we have that for all $X_{A\backslash G}$, there is some X_G such that $(M, w)^{X_{A\backslash G} \cap X_G} \models \varphi$. This implies that $(M, w)^{\top_{A\backslash G} \cap X_G} \models \varphi$ for the trivial strategy $\top_{A\backslash G}$ and some X_G. The latter is equivalent to $(M, w)^{X_G} \models \varphi$. Since φ is positive (and hence preserved under submodels), $(M, w)^{X'_G} \models \varphi$, where X'_G is the strongest strategy of G. The latter implies (again, due to the fact that φ is positive) that for all updates of the form $X'_G \cap X_{A\backslash G}$ (since they generate a submodel of $(M, w)^{X'_G}$), we also have $(M, w)^{X'_G \cap X_{A\backslash G}} \models \varphi$. And this is $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi$ by Proposition 4.6. □

Now we are ready to deal with the model checking for the positive case. Note that we do need to check case $[G]\varphi$ due to Proposition 4.10.

Proposition 4.12. Let $\varphi \in \mathcal{L}_{CAL} \cup \mathcal{L}_{GAL}$ be a formula such that all its subformulas ψ that are within scopes of $\langle G \rangle$ or $\langle\!\langle G \rangle\!\rangle$ belong to fragment $\mathcal{L}_{PAL^+}$ (or $\mathcal{L}_{GAL^+}$). Then the model checking problem for CAL and GAL is in P.

$\langle G\rangle\varphi, \langle\!\langle G\rangle\!\rangle\varphi$: compute $(||M||, w)$ and $(||M||^{X_G}, w)$, where X_G corresponds to the strongest strategy of G,
if $mc(||M||^{X_G}, w, \varphi)$ **then return** *true* **else return** *false*.

Proof. For this particular case we modify Algorithm 1 by inserting the following instead of cases on lines 7 and 8:

For all subformulas of φ_0, the algorithm calls are in P. Consider the modified call for $\langle\!\langle G\rangle\!\rangle\varphi$. It requires constructing a single update model given a specified strategy, which is a simple case of restricting the input model to the set of states in the strategy. This can be done in polynomial time. Then we call the algorithm on the updated model for φ, which by assumption requires polynomial time. □

Now, let us show that the algorithm is correct.

Proposition 4.13. Let (M, w) and $\varphi \in \mathcal{L}_{PAL^+}$ (or in $\varphi \in \mathcal{L}_{GAL^+}$) be given. The modified algorithm $mc(M, w, \varphi)$ returns *true* iff $(M, w) \models \varphi$.

Proof. By induction on φ. We show the case for $\langle\!\langle G\rangle\!\rangle\varphi$:

$\Rightarrow$: Suppose that $mc(M, w, \langle\!\langle G\rangle\!\rangle\varphi)$ returns *true*. This means that $mc(\|M\|^{X_G}, w, \varphi)$ returns *true*, where X_G is the strongest strategy of G. By the induction hypothesis, we have that $(\|M\|, w)^{X_G} \models \varphi$. Since φ is positive, for all stronger updates $X_G \cap X_{A\setminus G}$ it holds that $(\|M\|, w)^{X_G \cap X_{A\setminus G}} \models \varphi$, which is $(\|M\|, w) \models \langle\!\langle G\rangle\!\rangle\varphi$ by Proposition 4.6. Finally, the latter model is bisimilar to (M, w) and hence $(M, w) \models \langle\!\langle G\rangle\!\rangle\varphi$.

$\Leftarrow$: Let $(M, w) \models \langle\!\langle G\rangle\!\rangle\varphi$. By Proposition 4.6 this means that there is some X_G such that for all $X_{A\setminus G}$: $(M, w)^{X_G \cap X_{A\setminus G}} \models \varphi$. Set of all $X_{A\setminus G}$'s also includes the trivial strategy $\top_{A\setminus G}$, and we have $(M, w)^{X_G \cap \top_{A\setminus G}} \models \varphi$, which is equivalent to $(M, w)^{X_G} \models \varphi$. Since φ is positive and hence preserved under submodels, $(M, w)^{X'_G} \models \varphi$, where X'_G is the strongest strategy of G. By the induction hypothesis, we have that $mc(\|M\|^{X'_G}, w, \varphi)$ returns *true*. And by line 8 of the modified algorithm, we conclude that $mc(\|M\|, w, \langle\!\langle G\rangle\!\rangle\varphi)$ returns *true*. □

Note that in this particular case we cannot use the duality of coalition announcements. Indeed, if we try to rewrite $[\langle G\rangle]\varphi$ into $\neg\langle\!\langle G\rangle\!\rangle\neg\varphi$, we negate positive formula φ and it ceases to be positive. However, we can resolve the case of $[\langle G\rangle]\varphi$ by translating the formula into equivalent $\langle\!\langle A \setminus G\rangle\!\rangle\varphi$, which is allowed by Proposition 4.11.

Chapter 5

Logical Properties of GAL and CAL

Validity and non-validity of certain logical formulas may shed light on some internal properties of the logic as well as build (or disproof) intuitions about how this logic may be (dis)similar to some other one. In Section 5.1 we study how uniting and decoupling groups and coalitions of agents affects their powers to achieve some configurations of a given model. Moreover, we investigate some relations between box and diamond versions of group and coalition announcement operators. After that, in Section 5.2, we consider properties that capture some aspects of the interaction between CAL and GAL. Particularly, we demonstrate that a proposed definition of the CAL modality in terms of GAL modalities $\langle\!\langle G \rangle\!\rangle \varphi \leftrightarrow \langle G \rangle [A \setminus G] \varphi$ is not valid. Finally, it is shown in Section 5.3 that axioms of coalition logic (Section 2.3) remain valid if they are expressed in CAL. Throughout the chapter we assume that axioms of PAL remain valid for the richer languages of CAL and GAL. This is due to the fact that the validity of PAL schemata does not depend on the structure of formulas involved (and hence the proofs are not inductive) [van Ditmarsch et al., 2008, Chapter 4].

5.1 Cooperation and Ordering

5.1.1 Virtues of Cooperation

Intuition suggests that various groups and coalitions of agents, when united, can do no worse than if they were acting on their own. In this section we show that this intuition is indeed true.

We start with a somewhat obvious statement: if some configuration of a model can be achieved by a coalition, then the configuration can be achieved by a superset of the coalition.

Proposition 5.1. $\langle\!\langle G \rangle\!\rangle \varphi \to \langle\!\langle G \cup H \rangle\!\rangle \varphi$, where $G, H \subseteq A$, is valid.

Proof. Let $(M, w) \models \langle\!\langle G \rangle\!\rangle \varphi$ for some arbitrary (M, w). By the semantics of CAL this is equivalent to

$$\exists\psi_G, \forall\chi_{A\setminus G} : (M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi.$$

Let us consider formula $\chi_{A\setminus G}$: since $A \setminus G = A \setminus (G \cup H) \cup H \setminus G$, we can 'unpack' the formula into $\chi_{A\setminus(G\cup H)}$ and $\chi_{H\setminus G}$. Hence we have

$$\exists\psi_G, \forall\chi_{H\setminus G}, \forall\chi_{A\setminus(G\cup H)} : (M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{H\setminus G} \wedge \chi_{A\setminus(G\cup H)}]\varphi.$$

The latter implies

$$\exists\psi_G, \exists\top_{H\setminus G}, \forall\chi_{A\setminus(G\cup H)} : (M, w) \models \psi_G \wedge \top_{H\setminus G} \wedge [\psi_G \wedge \top_{H\setminus G} \wedge \chi_{A\setminus(G\cup H)}]\varphi,$$

where $\top_{H\setminus G} := \bigwedge_{a\in H\setminus G} K_a\top$. Combining ψ_G and $\top_{H\setminus G}$ into a single announcement $\psi_{G\cup H}$ by the united coalition $G \cup H$, we conclude that

$$\exists\psi_{G\cup H}, \forall\chi_{A\setminus(G\cup H)} : (M, w) \models \psi_{G\cup H} \wedge [\psi_{G\cup H} \wedge \chi_{A\setminus(G\cup H)}]\varphi.$$

This is equivalent to $(M, w) \models \langle\!\langle G \cup H \rangle\!\rangle\varphi$ by the semantics. □

It was shown in [Ågotnes et al., 2010] that $\langle G\rangle\varphi \leftrightarrow \langle G\rangle\langle G\rangle\varphi$. This property demonstrates that within the framework of GAL a multiple-step strategy of a group can be executed in a single step. Whether this is true for CAL is an open question. We show, however, that if the truth of some φ can be achieved by two consecutive coalition announcements by G, then whatever agents from $A \setminus G$ announce, they cannot preclude G from making φ true.

Proposition 5.2. $\langle\!\langle G\rangle\!\rangle\langle\!\langle G\rangle\!\rangle\varphi \to [\![A \setminus G]\!]\varphi$ is valid.

Proof. Suppose that for some (M, w) it holds that $(M, w) \models \langle\!\langle G\rangle\!\rangle\langle\!\langle G\rangle\!\rangle\varphi$. This is equivalent to

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G, \forall\chi'_{A\setminus G} : (M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}](\psi'_G \wedge [\psi'_G \wedge \chi'_{A\setminus G}]\varphi).$$

Since $\chi'_{A\setminus G}$ quantifies over all epistemic formulas known to $A\setminus G$, it also quantifies over $\top_{A\setminus G} := \bigwedge_{a\in A\setminus G} K_a\top$. Hence it is implied that

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G : (M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}](\psi'_G \wedge [\psi'_G \wedge \top_{A\setminus G}]\varphi),$$

which is equivalent to

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G : (M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\psi'_G \wedge [\psi_G \wedge \chi_{A\setminus G}][\psi'_G]\varphi.$$

Using PAL validity $[\psi]\chi \wedge [\psi][\chi]\varphi \leftrightarrow [\psi]\chi \wedge [\psi]\langle\chi\rangle\varphi$, we get

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G : (M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\psi'_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\langle\psi'_G\rangle\varphi.$$

Next, we use PAL validity $[\psi]\varphi \leftrightarrow (\psi \to \langle\psi\rangle\varphi)$:

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G : (M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\psi'_G \wedge (\psi_G \wedge \chi_{A\setminus G} \to \langle\psi_G \wedge \chi_{A\setminus G}\rangle\langle\psi'_G\rangle\varphi).$$

By propositional reasoning the latter implies

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G : (M,w) \models \psi_G \land (\psi_G \to (\chi_{A\setminus G} \to \langle\psi_G \land \chi_{A\setminus G}\rangle\langle\psi'_G\rangle\varphi),$$

and this implies

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G : (M,w) \models \chi_{A\setminus G} \to \langle\psi_G \land \chi_{A\setminus G}\rangle\langle\psi'_G\rangle\varphi.$$

Finally, by PAL axiom $\langle\psi\rangle\langle\chi\rangle\varphi \leftrightarrow \langle\psi \land [\psi]\chi\rangle\varphi$, we have that

$$\exists\psi_G, \forall\chi_{A\setminus G}, \exists\psi'_G : (M,w) \models \chi_{A\setminus G} \to \langle\psi_G \land \chi_{A\setminus G} \land [\psi_G \land \chi_{A\setminus G}]\psi'_G\rangle\varphi.$$

We can move $\exists\psi_G$ within the scope of $\forall\chi_{A\setminus G}$, and morph ψ_G and $[\psi_G \land \chi_{A\setminus G}]\psi'_G$ into a single announcement by G.

The latter is $(M,w) \models [\langle A \setminus G\rangle]\varphi$ by the semantics of CAL. □

Whether $\langle\!\langle G\rangle\!\rangle\langle\!\langle G\rangle\!\rangle\varphi \to \langle\!\langle G\rangle\!\rangle\varphi$ is valid is an open question. We conjecture that the property is not valid. Consider $\langle\!\langle G\rangle\!\rangle\langle\!\langle G\rangle\!\rangle\varphi$: after initial announcement, coalition G has a consecutive announcement to make φ true. This announcement, however, depends on the choice of $A \setminus G$ in the first operator. In other words, a consecutive announcement by G may vary depending on the initial announcement by $A \setminus G$. Hence, it seems highly counterintuitive that G has a single announcement that can incorporate all possible simultaneous announcements by $A \setminus G$ in a general (infinite) case.

Formula $\langle G\rangle\langle H\rangle\varphi \to \langle G \cup H\rangle\varphi$ is a validity of GAL [Ågotnes et al., 2010]. Again, it is unknown whether the same property holds for coalition operators, and, for the same reasons as for Proposition 5.2, we conjecture that the corresponding formula is not valid in CAL.

We show, however, a generalisation of Proposition 5.2.

Proposition 5.3. $\langle\!\langle G\rangle\!\rangle\langle\!\langle H\rangle\!\rangle\varphi \to [\langle A \setminus (G \cup H)\rangle]\varphi$ is valid.

Proof. Let $(M,w) \models \langle\!\langle G\rangle\!\rangle\langle\!\langle H\rangle\!\rangle\varphi$. By Proposition 5.1 applied twice, we have $(M,w) \models \langle\!\langle G \cup H\rangle\!\rangle\langle\!\langle G \cup H\rangle\!\rangle\varphi$, and by Proposition 5.2, the latter implies $(M,w) \models [\langle A \setminus (G \cup H)\rangle]\varphi$. □

Next, we show that splitting an announcement by a unified coalition into consecutive announcements by sub-coalitions may decrease their power to force certain outcomes. Whether $\langle\!\langle G \cup H\rangle\!\rangle\varphi \to \langle\!\langle G\rangle\!\rangle\langle\!\langle H\rangle\!\rangle\varphi$ is valid was mentioned as an open question in [Ågotnes et al., 2016]. We settle this problem by presenting a counterexample.

Proposition 5.4. $\langle\!\langle G \cup H\rangle\!\rangle\varphi \to \langle\!\langle G\rangle\!\rangle\langle\!\langle H\rangle\!\rangle\varphi$ is not valid.

Proof. Let $G = \{a\}, H = \{b\}$, and $\varphi := K_b(p \land q \land r) \land \neg K_a(p \land q \land r) \land \neg K_c(p \land q \land r)$. Formula φ says that agent b knows that the given propositional variables are true, and agents a and c do not. Consider model (M, pqr) in Figure 5.1 (reflexive and transitive arrows are omitted for convenience). Names of the states in the model show values of propositional variables; for example, $(M, p\overline{q}r) \models p \land \neg q \land r$.

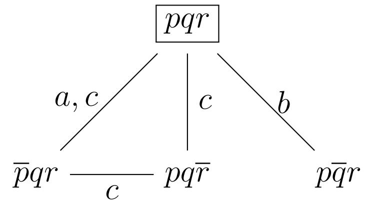

Figure 5.1: Counterexample

By the semantics $(M, pqr) \models \langle\!\langle \{a, b\} \rangle\!\rangle \varphi$ if and only if $\exists \psi_a, \exists \psi_b, \forall \chi_c : (M, pqr) \models \psi_a \wedge \psi_b \wedge [\psi_a \wedge \psi_b \wedge \chi_c]\varphi$. Let ψ_a be $K_a q$, and ψ_b be $K_b \top$. Observe that $(M, pqr) \models K_a q \wedge K_b \top$. This announcement leads to b learning that q. Moreover, c does not know any formula that she can announce to avoid φ. An informal argument is as follows. By announcing $K_a q$ agent a 'chooses' a union of a-equivalence classes $\{pqr, \overline{p}qr, pq\overline{r}\}$ (and b 'chooses' the whole model). Any simultaneous 'choice' of c includes $\{pqr, \overline{p}qr, pq\overline{r}\}$ as a subset. Thus, intersection of $\{pqr, \overline{p}qr, pq\overline{r}\}$ and any of unions of c-equivalence classes is $\{pqr, \overline{p}qr, pq\overline{r}\}$, and φ is true in such a restriction of the model.

Let us show that $(M, pqr) \not\models \langle\!\langle \{a\} \rangle\!\rangle \langle\!\langle \{b\} \rangle\!\rangle \varphi$, or, equivalently, $(M, pqr) \models [\langle \{a\} \rangle][\langle \{b\} \rangle]\neg\varphi$. According to the semantics, $\forall \psi_a, \exists \chi_b, \exists \chi_c$: $(M, pqr) \models \psi_a \rightarrow \langle \psi_a \wedge \chi_b \wedge \chi_c \rangle [\langle \{b\} \rangle]\neg\varphi$. Assume that for an arbitrary ψ_a, announcements by b and c are $K_b p$ and $K_c \top$ correspondingly. Then $(M, pqr) \models \psi_a \wedge [\psi_a \wedge K_b p \wedge K_c \top][\langle \{b\} \rangle]\neg\varphi$. Note that no matter what a announces, $K_b p$ 'forces' her to learn that $p \wedge q \wedge r$, and whatever is announced in the updated model $(M, pqr)^{\psi_a \wedge K_b p \wedge K_c \top}$, a's knowledge of $p \wedge q \wedge r$ and, hence, falsity of φ remains. Thus we reached a contradiction. □

The same counterexample can be used to demonstrate that $[\langle A \setminus (G \cup H) \rangle]\varphi \rightarrow \langle\!\langle G \rangle\!\rangle \langle\!\langle H \rangle\!\rangle \varphi$ is not valid, where $A \setminus (G \cup H) = \{c\}$. In the proof of Proposition 5.4 we show that $(M, pqr) \models \langle\!\langle \{a, b\} \rangle\!\rangle \varphi$. Using Proposition 5.7 we obtain $(M, pqr) \models [\langle c \rangle]\varphi$. The rest of the proof remains the same.

Corollary 5.5. $[\langle A \setminus (G \cup H) \rangle]\varphi \rightarrow \langle\!\langle G \rangle\!\rangle \langle\!\langle H \rangle\!\rangle \varphi$ is not valid.

To the best of our knowledge, the group announcement version of Proposition 5.4 has not been considered. We show that the property does not hold for GAL operators as well. To derive a contradiction, we use the intuition that separated groups, while being able to force a certain configuration of a model when united, may lack discerning power on their own. Contrast this to the proof of Proposition 5.4, where the contradiction was derived on the basis that former partners may spoil each other's strategies when pitched against one another.

Proposition 5.6. $\langle G \cup H \rangle \varphi \rightarrow \langle G \rangle \langle H \rangle \varphi$ is not valid[1].

Proof. Consider the model in Figure 5.2.

[1]The original idea of an infinite-grid counterexample is by Tim French. Here we present its finite and reworked version.

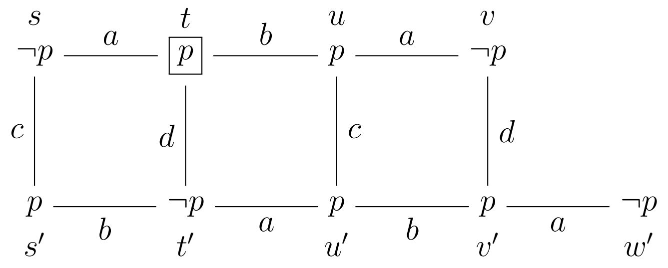

Figure 5.2: Model M_1

Note that M_1 is bisimulation contracted, and (M_1, t) can be distinguished from other states in this proof by some distinguishing formula φ_1. Also let $G = \{a, d\}$ and $H = \{b, c\}$. Next consider model M_2 in Figure 5.3.

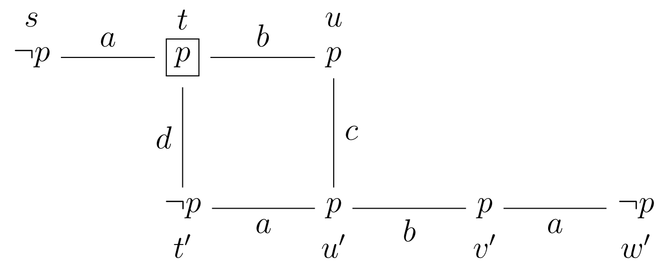

Figure 5.3: Model M_2

Again, M_2 is bisimulation contracted, and let some φ_2 be a distinguishing formula for (M_2, t). The union of all agents in model M_1 can make φ_2 true, i.e. $(M_1, t) \models \langle\{a, d\} \cup \{b, c\}\rangle\varphi_2$. Indeed, a possible mutual choice for the agents is as follows: $X_a = \{s, t, u, v, t', u', v', w'\}$, $X_b = \{s, t, u, s', t', u', v', w'\}$, $X_c = \{s, t, u, s', t', u', v', w'\}$, and $X_d = \{s, t, u, v, t', u', v', w'\}$. Hence the corresponding group announcement reduce M_1 to $X_a \cap X_b \cap X_c \cap X_d = \{s, t, u, t', u', v', w'\}$ which is exactly model M_2.

Now we show that $(M_1, t) \models [\{a, d\}][\{b, c\}]\neg\varphi_2$, or, informally, any successive announcements by the corresponding groups do not result in M_2. Since we are interested only in group announcements that can lead to M_2, and due to the fact that M_2 is bisimulation contracted, we do not consider announcements that result in a model with fewer states than M_2.

There are only two such strategies for $\{a, d\}$. First strategy is the trivial one — a and d announce $K_a\top$ and $K_d\top$. Such an announcement leaves M_1 intact. It is easy to see that whatever $\{b, c\}$ announce afterwards, they cannot both retain only states of M_2. The closest they can get to M_2 is M_3, which is presented in Figure 5.4. Clearly, M_3 is not bisimilar to M_2.

The second meaningful restriction of M_1 by $\{a, d\}$ is shown in Figure 5.5.

It might seem that the only difference between M_2 and M_4 is state v. Observe, however, that v is bisimilar to t', and any announcement by $\{b, c\}$ that 'deletes' v will also 'delete' t' (see Figure 5.6).

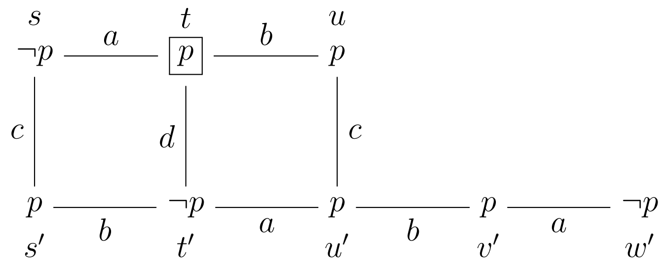

Figure 5.4: Model M_3

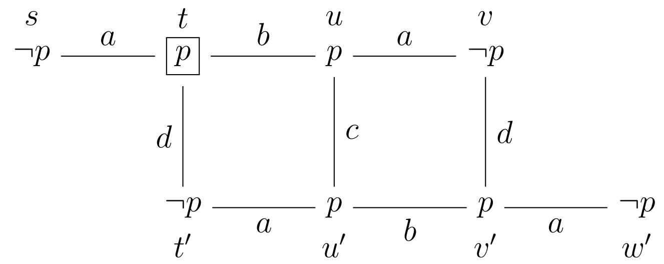

Figure 5.5: Model M_4

Thus we showed that $(M_1, t) \models [\{a, d\}][\{b, c\}]\neg\varphi_2$, which is equivalent to $(M_1, t) \not\models \langle\{a, d\}\rangle\langle\{b, c\}\rangle\varphi_2$. □

5.1.2 Boxes, Diamonds, and Church-Rosser

In this section we consider two rather straightforward results for coalition announcement operators, and demonstrate that the Church-Rosser property does holds in neither GAL nor CAL (although it holds in APAL [Balbiani et al., 2008]).

We start with the fact that if coalition G has an announcement such that they can achieve φ no matter what agents outside of the coalition announce at the same time, then for every possible announcement by $A \setminus G$ there is a corresponding 'counter-announcement' such that φ holds afterwards.

Proposition 5.7. $\langle\!\lbrack G\rbrack\!\rangle\varphi \to [\!\langle A \setminus G\rangle\!]\varphi$ is valid.

Proof. Assume that for some arbitrary (M, w) we have that $(M, w) \models \langle\!\lbrack G\rbrack\!\rangle\varphi$. By the semantics this is equivalent to $\exists\psi_G, \forall\chi_{A\setminus G} : (M, w) \models \psi_G \land [\psi_G \land \chi_{A\setminus G}]\varphi$, and the latter implies $\forall\chi_{A\setminus G}, \exists\psi_G : (M, w) \models \psi_G \land [\psi_G \land \chi_{A\setminus G}]\varphi$. Using the validity of PAL $\models [\psi]\varphi \leftrightarrow (\psi \to \langle\psi\rangle\varphi)$, we have that $\forall\chi_{A\setminus G}, \exists\psi_G : (M, w) \models \psi_G \land (\psi_G \land \chi_{A\setminus G} \to \langle\psi_G \land \chi_{A\setminus G}\rangle\varphi)$, which implies, by propositional reasoning, $\forall\chi_{A\setminus G}, \exists\psi_G : (M, w) \models \chi_{A\setminus G} \to \langle\psi_G \land \chi_{A\setminus G}\rangle\varphi$. The latter is $(M, w) \models [\!\langle A \setminus G\rangle\!]\varphi$ by the semantics of CAL. □

The other direction of Proposition 5.7 is not valid. An intuitive explanation is that even though $A \setminus G$ may have a 'counter-announcement' to every announcement by G, they may, at the same time, lack the single 'universal' announcement for all possible G's announcements.

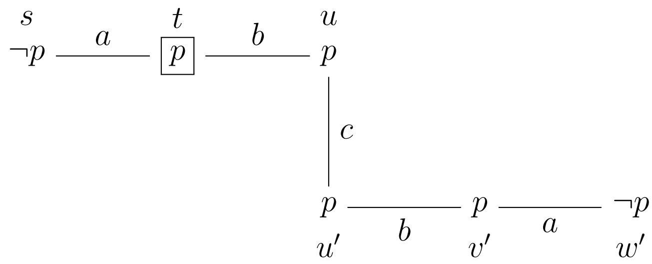

Figure 5.6: Model M_5

Proposition 5.8. $[\langle G\rangle]\varphi \to \langle[A \setminus G]\rangle\varphi$ is not valid.

Proof. We present a counterexample. Consider the model in Figure 5.7.

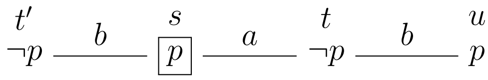

Figure 5.7: Model M_1

Pointed model (M_1, s) can be described by formula $\varphi_1 := p \wedge \widehat{K}_b(\neg p \wedge K_a\neg p) \wedge \widehat{K}_a(\neg p \wedge \widehat{K}_b K_a p)$.

Let us also consider some submodels of M_1 presented in Figure 5.8.

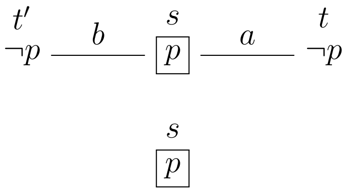

Figure 5.8: Models M_2 (top) and M_3 (bottom)

Corresponding distinguishing formulas for (M_2, s) and (M_3, s) are $\varphi_2 := p \wedge \widehat{K}_b(\neg p \wedge K_a\neg p) \wedge \widehat{K}_a(\neg p \wedge K_b\neg p)$ and $\varphi_3 := p \wedge K_a p \wedge K_b p$.

Let $G = \{a\}$, $A \setminus G = \{b\}$, and $\varphi := \varphi_1 \vee \varphi_2 \vee \varphi_3$.

First we show that $(M_1, s) \models [\langle a\rangle]\varphi$. By the semantics of CAL this is equivalent to $\forall\psi_a, \exists\chi_b$: $(M_1, s) \models \psi_a \to \langle\psi_a \wedge \chi_b\rangle\varphi$. In the model, agent a has four possible 'choices': $\{s, t\}, \{t', s, t\}, \{s, t, u\}$, and $\{t', s, t, u\}$. Options $\{t', s, t\}$ and $\{t', s, t, u\}$ clearly satisfy φ as agent b announces $K_b\top$ at the same time. If agent a 'chooses' to announce $\{s, t\}$ or $\{s, t, u\}$, agent b can announce $\{t', s\}$ at the same time, and such a joint announcement results in $\{s\}$ that satisfies φ_3.

Now, let us show that $(M_1, s) \not\models \langle[b]\rangle\varphi$, which is equivalent to $(M_1, s) \models [\langle b\rangle]\neg\varphi$. Agent b has only two options — $\{t', s, t, u\}$ and $\{t', s\}$. In the first case, agent a simultaneously 'chooses' $\{s, t\}$ that leads to φ being false. In the second case, agent a announces $K_a\top$, and φ is false in the resulting model. □

The Church-Rosser principle, $\Diamond\Box\varphi \to \Box\Diamond\varphi$, where $\Diamond$ and $\Box$ are some modal operators, corresponds to the confluence frame property $\forall x, y, z(xRy \wedge xRz \to \exists w(yRw \wedge zRw))$ (see [Blackburn et al., 2001, Chapter 3]). We are interested in how group boxes and diamonds commute together. In Proposition 5.9 we show that the Church-Rosser property does not hold for group announcements. An intuitive explanation of this fact may be that knowledge of agents changes as a model is updated. Hence, they may lose their original strategies and discerning power as a result of an announcement by some other group. In other words, the order of announcements matters.

Proposition 5.9. $\langle G\rangle[H]\varphi \to [H]\langle G\rangle\varphi$ is not valid.

Proof. The counterexample model is the same as in Proposition 5.14 and presented in Figure 5.9.

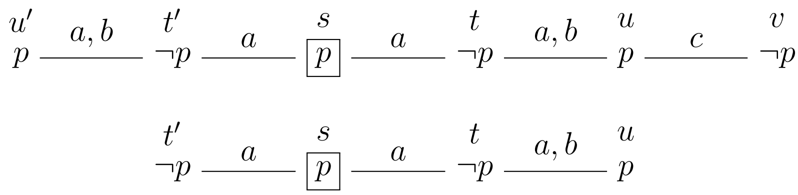

Figure 5.9: Models M_1 (top) and M_2 (bottom)

Formula φ is $\widehat{K}_a K_b \neg p \wedge \widehat{K}_a(\widehat{K}_b p \wedge \widehat{K}_b \neg p)$, and $(M_2, s) \models \varphi$ and $(M_1, s) \not\models \varphi$.

First we show that $(M_1, s) \models \langle\{a\}\rangle[\{b, c\}]\neg\varphi$. Let a's announcement be $K_a(\neg p \to K_c \neg p)$. Update of (M_1, s) with this announcement $(M_1, s)^{K_a(\neg p \to K_c \neg p)}$ is shown in Figure 5.10.

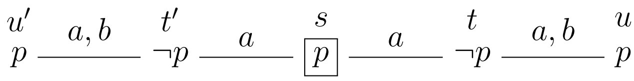

Figure 5.10: Model $(M_1, s)^{K_a(\neg p \to K_c \neg p)}$

Note that in this model states t and t', and u and u' became bisimilar. Hence, no matter what agents b and c announce, they cannot get rid of u' without 'deleting' u as well. In other words, agents b and c cannot make φ true. This establishes $(M_1, s) \models \langle\{a\}\rangle[\{b, c\}]\neg\varphi$.

The remaining part of the proof is to show that $(M_1, s) \not\models [\{b, c\}]\langle\{a\}\rangle\neg\varphi$, or, equivalently, $(M_1, s) \models \langle\{b, c\}\rangle[\{a\}]\varphi$. Let b and c's announcement be $K_c(p \to (K_b p \vee \widehat{K}_c \neg p))$ and $K_b(\neg p \to \widehat{K}_b p)$. Such a mutual announcement results in model M_2. Observe that in (M_2, s), since the whole model is an a-equivalence class, agent a has no announcement to modify it. Moreover, $(M_2, s) \models \varphi$, and hence $(M_1, s) \models \langle\{b, c\}\rangle[\{a\}]\varphi$. $\Box$

Remark. In [Ågotnes et al., 2010] (as well as in [van Ditmarsch, 2012]) it was claimed that the Church-Rosser property holds for GAL. However, we presented a counterexample to the property in Proposition 5.9. One of the problems with the proof in [Ågotnes et al., 2010] is the transition from $\langle\psi\rangle[\chi]\varphi$ to $\langle\psi \wedge \chi\rangle\varphi$, where ψ and χ are announcements of values of propositional variables by agents (and thus are positive formulas). We present a small counterexample to $\langle\psi \wedge \chi\rangle\varphi \rightarrow \langle\psi\rangle\langle\chi\rangle\varphi$.

Proposition 5.10. Let ψ and χ be positive formulas. The following are not valid: $\langle\psi \wedge \chi\rangle\varphi \rightarrow \langle\psi\rangle\langle\chi\rangle\varphi$ and $\langle\psi\rangle\langle\chi\rangle\varphi \rightarrow \langle\psi \wedge \chi\rangle\varphi$.

Proof. Consider model M in Figure 5.11. State names represent values of propositional variables in them, and agent a's relation is the identity.

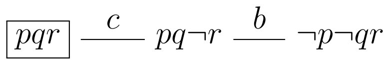

Figure 5.11: A counterexample

Let $\psi := K_a p$, $\chi := K_b q$, and $\varphi := K_c r$. It is easy to see that $(M, pqr) \models \langle K_a p \wedge K_b q\rangle K_c r$, since the updated model $(M, pqr)^{K_a p \wedge K_b q}$ consists only of one state — pqr.

Successive announcements, on the other hand, do not make $K_c r$ true. Indeed, after announcement of $K_a p$, the resulting model consists of pqr and $pq\neg r$. In both states $K_b q$ holds, hence this announcement does not modify the updated model, and c considers r and $\neg r$ possible. Formally, $(M, pqr) \models [K_a p][K_b q]\neg K_c r$.

The same argument holds in the other direction. $\square$

The Church-Rosser principle is not valid in CAL as well.

Proposition 5.11. $\langle\!\langle G\rangle\!\rangle[\![H]\!]\varphi \rightarrow [\![H]\!]\langle\!\langle G\rangle\!\rangle\varphi$ is not valid.

Proof. Consider models in Figures 5.7 and 5.8. Also let $G = \{a\}$, $H = \{b\}$, and $\varphi_2 := p \wedge \widehat{K}_b \neg p \wedge \widehat{K}_a \neg p$, $\varphi_3 := p \wedge K_a p \wedge K_b p$, and $\varphi := \varphi_2 \vee \varphi_3$. Note that $(M_2, s) \models \varphi_2$ and $(M_3, s) \models \varphi_3$.

First we show that $(M_1, s) \models \langle\!\langle\{a\}\rangle\!\rangle[\![\{b\}]\!]\varphi$, which means that agent a has a strategy X_a such that whichever strategy X_b agent b simultaneously implements, $[\![\{b\}]\!]\varphi$ holds in the resulting model. Consider a's strategy $X_a = \{t', s, t\}$. Agent b has only two options in (M_1, s): $X_b^1 = \{t', s, t, u\}$ and $X_b^2 = \{t', s\}$. Two possible resulting models are presented in Figure 5.12

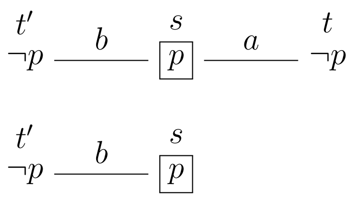

Figure 5.12: Resulting Models $(M_1, s)^{X_a \cap X_b^1}$ (top) and $(M_1, s)^{X_a \cap X_b^2}$ (bottom)

Next let us examine further model updates by coalition announcement $[\langle\{b\}\rangle]$. Again, there are only two options for agent b in $(M_1, s)^{X_a \cap X_b^1}$: $Y_b^1 = \{t', s, t\}$ and $Y_b^2 = \{t', s\}$. On Y_b^1 agent a responds with the same strategy, and on Y_b^2 she responds with $\{s, t\}$ that results in the model with single state s. In both cases φ holds. In $(M_1, s)^{X_a \cap X_b^2}$ agent b has only trivial strategy, and a responds with $\{s\}$ yielding the single-state model and making φ true. Hence, $(M_1, s) \models \langle\!\langle\{a\}\rangle\!\rangle[\langle\{b\}\rangle]\varphi$.

Now we show that $(M_1, s) \not\models [\langle\{b\}\rangle]\langle\!\langle\{a\}\rangle\!\rangle\varphi$, or, equivalently, that $(M_1, s) \models \langle\!\langle\{b\}\rangle\!\rangle[\langle\{a\}\rangle]\neg\varphi$. Let b's strategy be the trivial one, i.e. $X_b = \{t', s, t, u\}$. Results of updates of (M_1, s) with various a's strategies are presented in Figure 5.13.

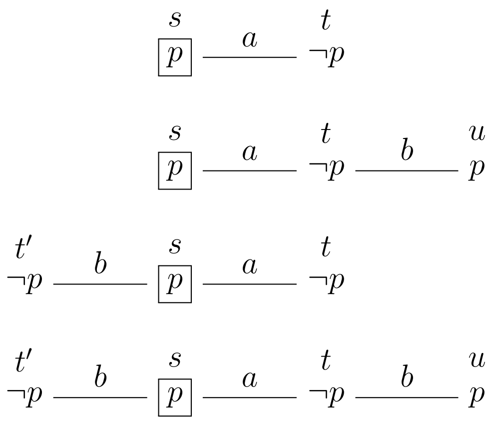

Figure 5.13: Models $(M_1, s)^{X_b \wedge X_a^1}$, $(M_1, s)^{X_b \wedge X_a^2}$, $(M_1, s)^{X_b \wedge X_a^3}$, and $(M_1, s)^{X_b \wedge X_a^4}$

Finally, we consider further updates of the models in Figure 5.13 by $[\langle\{a\}\rangle]$. It is easy to see that any further announcements by a in models $(M_1, s)^{X_b \wedge X_a^1}$ and $(M_1, s)^{X_b \wedge X_a^2}$ can be countered by the trivial strategy of b so that $\neg\varphi$ is true in resulting models. In model $(M_1, s)^{X_b \wedge X_a^3}$ agent b responds with $\{t', s, t\}$ on a's strategy $\{s, t\}$, and with $\{t', s\}$ on a's $\{t', s, t\}$; in both restrictions φ is false. Cases for $(M_1, s)^{X_b \wedge X_a^4}$ are the same as for other updates of (M_1, s). Thus, $(M_1, s) \models \langle\!\langle\{b\}\rangle\!\rangle[\langle\{a\}\rangle]\neg\varphi$. □

5.2 Interaction Between Coalition and Group Announcements

We start this section with somewhat basic results concerning interaction between GAL and CAL operators.

In Proposition 5.12 formula 1 states that if a coalition can force some outcome, then they can achieve the outcome by a group announcement. Property 2 shows that coalition and group announcements are equivalent for the grand coalition A. That an anti-coalition cannot undo the result of a coalition announcement is presented in 3. Finally, property 4 states that if a coalition can force some outcome, then they can achieve the outcome by making one additional group announcement. The converse, however, is not valid (formula 5).

Proposition 5.12. 1–4 are valid, and 5 is not valid.

1. $\langle\!\langle G \rangle\!\rangle\varphi \rightarrow \langle G \rangle\varphi$,
2. $\langle\!\langle A \rangle\!\rangle\varphi \leftrightarrow \langle A \rangle\varphi$,
3. $\langle\!\langle G \rangle\!\rangle\varphi \leftrightarrow \langle\!\langle G \rangle\!\rangle[A \setminus G]\varphi$,
4. $\langle\!\langle G \rangle\!\rangle\varphi \rightarrow \langle\!\langle G \rangle\!\rangle\langle G \rangle\varphi$,
5. $\langle\!\langle G \rangle\!\rangle\langle G \rangle\varphi \rightarrow \langle\!\langle G \rangle\!\rangle\varphi$.

Proof. 1. If G can make φ true no matter what agents from $A \setminus G$ simultaneously announce, they can make φ true if all agents from coalition $A \setminus G$ announce $\top$.

2. Trivially by the semantics of the grand coalition.

3. *From left to right.* We prove the contrapositive. Let $(M, w) \models [\![G]\!]\langle A \setminus G\rangle\varphi$ for some arbitrary (M, w). By the semantics of CAL we have that $\forall\psi_G, \exists\chi_{A\setminus G}, \exists\chi'_{A\setminus G}$: $(M, w) \models \psi_G \rightarrow \langle\psi_G \wedge \chi_{A\setminus G}\rangle\langle\chi'_{A\setminus G}\rangle\varphi$. Due to PAL validity $\langle\psi\rangle\langle\chi\rangle\varphi \leftrightarrow \langle\psi \wedge [\psi]\chi\rangle\varphi$ the latter is equivalent to $(M, w) \models \psi_G \rightarrow \langle\psi_G \wedge \chi_{A\setminus G} \wedge [\psi_G \wedge \chi_{A\setminus G}]\chi'_{A\setminus G}\rangle\varphi$. Note that $\chi_{A\setminus G} \wedge [\psi_G \wedge \chi_{A\setminus G}]\chi'_{A\setminus G}$ in the presence of $\psi_G \wedge \chi_{A\setminus G}$ is equivalent to some epistemic formula announced by agents from $A \setminus G$ (see the proof of Proposition 5.13 for details). Hence, by semantics the latter is equivalent to $(M, w) \models [\![G]\!]\varphi$.

   From right to left. Immediate by the fact that $A \setminus G$ can announce $\top_{A\setminus G}$.

4. Immediate by the fact that G can announce $\top_G$ after the coalition announcement.

5. The counterexample is the same as in Proposition 5.8 with $G = \{b\}$. Indeed, $(M_1, s) \models \langle\!\langle \{b\} \rangle\!\rangle\langle\{b\}\rangle\varphi$, which is equivalent to $\exists\psi_b, \forall\chi_a, \exists\psi'_b$: $(M, w) \models \psi_b \wedge [\psi_b \wedge \chi_a]\langle\psi'_b\rangle\varphi$. Let $\psi_b := K_b\top$. Then we have that $\forall\chi_a, \exists\psi'_b$: $(M, w) \models K_b\top \wedge [K_b\top \wedge \chi_a]\langle\psi'_b\rangle\varphi$, or $\forall\chi_a, \exists\psi'_b$: $(M, w) \models [\chi_a]\langle\psi'_b\rangle\varphi$. The rest of the proof follows the one of Proposition 5.8 with substitution of b's simultaneous choice $\{t', s\}$ with the consecutive choice $\{s\}$.

□

Whether CAL operators can be expressed in GAL is an open question (work in progress). The most probable definition of coalition announcements in terms of group announcements is $\langle\!\langle G \rangle\!\rangle\varphi \leftrightarrow \langle G \rangle[A \setminus G]\varphi$. The validity of this formula was stated to be an open question in [van Ditmarsch, 2012; Ågotnes et al., 2016]. We settle this problem by proving one direction and presenting a counterexample to the other direction.

Consider the left-to-right direction of the formula. In the antecedent all agents make a simultaneous announcement, whereas in the consequent agents from $A \setminus G$ know the announcement ψ_G made by G. Thus, in the updated model $(M, w)^{\psi_G}$ the agents in $A \setminus G$ may have learned some *new* epistemic formulas $\chi_{A\setminus G}$ that they did not know before the announcement. However, since ψ_G holds in the initial model, and $\chi_{A\setminus G}$ holds in the updated one, agents from $A \setminus G$ can always make an announcement in the initial model that they know that after the announcement of ψ_G, $\chi_{A\setminus G}$ is true. Informally, Beth can say to Ann: 'I don't know whether *Infinite Jest* is worth reading, but if you say that it is, then I'll learn that you've heard about (and probably read) the book.' This announcement, made simultaneously with the announcement by G, 'models' the effect of announcing $\chi_{A\setminus G}$ later. Moreover, this announcement is equivalent to an announcement of an epistemic formula due to the translation function (Definition 2.25).

Proposition 5.13. $\langle\!\langle G \rangle\!\rangle\varphi \to \langle G \rangle [A \setminus G]\varphi$ is valid.

Proof. Assume that for some pointed model (M, w) it holds that $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi$. By the semantics of CAL this is equivalent to

$$\exists \psi_G, \forall \chi_{A\setminus G} : (M, w) \models \psi_G \land [\psi_G \land \chi_{A\setminus G}]\varphi.$$

Since $\chi_{A\setminus G}$ quantifies over all possible announcements by $A \setminus G$, it also quantifies over a specific subset of these announcements — $K_{A\setminus G}[\psi_G]\chi'_{A\setminus G} := \bigwedge_{a\in A\setminus G} K_a[\psi_G]\chi'_a$ for some ψ_G and for all $\chi'_a \in \mathcal{L}_{EL}$.

Hence $\exists \psi_G, \forall \chi_{A\setminus G}$: $(M, w) \models \psi_G \land [\psi_G \land \chi_{A\setminus G}]\varphi$ implies

$$\exists \psi_G, \forall \chi'_{A\setminus G} : (M, w) \models \psi_G \land [\psi_G \land K_{A\setminus G}[\psi_G]\chi'_{A\setminus G}]\varphi.$$

Note that $K_{A\setminus G}[\psi_G]\chi'_{A\setminus G}$ is not an epistemic formula *per se*. It is equivalent, however, to an epistemic formula of type $K_{A\setminus G}\chi_{A\setminus G}$, where $\chi_{A\setminus G} \in \mathcal{L}_{EL}$, via translation $t(K_{A\setminus G}[\psi_G]\chi'_{A\setminus G})$ (Definition 2.25). Thus we have that

$$\exists \psi_G, \forall \chi'_{A\setminus G} : (M, w) \models \psi_G \land [\psi_G \land t(K_{A\setminus G}[\psi_G]\chi'_{A\setminus G})]\varphi.$$

Let us consider announcement $\psi_G \land t(K_{A\setminus G}[\psi_G]\chi'_{A\setminus G})$. By propositional reasoning it is equivalent to $\psi_G \land (\psi_G \to t(K_{A\setminus G}[\psi_G]\chi'_{A\setminus G}))$. Since ψ_G is an epistemic formula, the latter is equivalent to $\psi_G \land t(\psi_G \to K_{A\setminus G}[\psi_G]\chi'_{A\setminus G})$. Applying the PAL axiom $[\psi]K_a\varphi \leftrightarrow (\psi \to K_a[\psi]\varphi)$, we get $\psi_G \land t([\psi_G]K_{A\setminus G}\chi'_{A\setminus G})$, which is equivalent to $\psi_G \land [\psi_G]K_{A\setminus G}\chi'_{A\setminus G}$. Finally, we have that

$$\exists \psi_G, \forall \chi'_{A\setminus G} : (M, w) \models \psi_G \land [\psi_G \land [\psi_G]K_{A\setminus G}\chi'_{A\setminus G}]\varphi.$$

Using the axiom $[\psi][\chi]\varphi \leftrightarrow [\psi \land [\psi]\chi]\varphi$, we get

$$\exists \psi_G, \forall \chi'_{A\setminus G} : (M, w) \models \psi_G \land [\psi_G][K_{A\setminus G}\chi'_{A\setminus G}]\varphi,$$

where $\chi'_{A\setminus G} \in \mathcal{L}_{EL}$. The latter is equivalent $(M, w) \models \langle G \rangle [A \setminus G]\varphi$ due to validity $\models \psi \land [\psi]\varphi \leftrightarrow \langle \psi \rangle \varphi$ and by the semantics of GAL. □

The converse of Proposition 5.13 is, however, not valid. There are two main points in the intuition behind a counterexample. First, an announcement by G may make some states bisimilar and thus indistinguishable for $A \setminus G$. In such a way, agents from $A \setminus G$ may 'lose' some strategies they had in the original model. And the second point is that an announcement by a group of agents $A \setminus G$ can influence not only epistemic relations of their opponents, but of agents from $A \setminus G$ as well.

Proposition 5.14. $\langle G\rangle[A \setminus G]\varphi \to \langle\!\langle G\rangle\!\rangle\varphi$ is not valid.

Proof. We present a counterexample (Figures 5.14, 5.15, and 5.16) to the contraposition $[\langle G\rangle]\varphi \to [G]\langle A \setminus G\rangle\varphi$. The reader can verify that models M_1 and M_2 are bisimulation contracted.

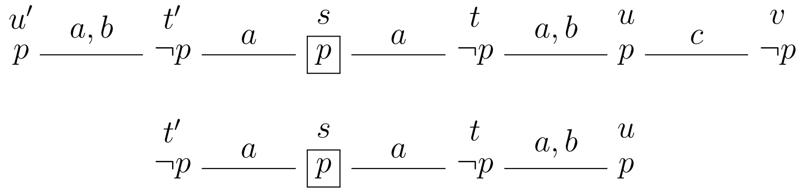

Figure 5.14: Models M_1 and M_2

Let $G = \{a\}$, and $A \setminus G = \{b, c\}$. Also, let $\varphi := \widehat{K}_a K_b \neg p \land \widehat{K}_a(\widehat{K}_b p \land \widehat{K}_b \neg p)$. This formula is a distinguishing formula of state s of model M_2, i.e. φ is true only in (M_2, s) and nowhere else in this proof.

First, we show that $(M_1, s) \models [\langle a\rangle]\varphi$. By the semantics of CAL this means that for every ψ_a, there are χ_b and χ_c such that $(M_1, s) \models \psi_a \to \langle \psi_a \land \chi_b \land \chi_c\rangle\varphi$. In terms of strategies this says that whatever strategy X_a agent a chooses in (M, s), agents $\{b, c\}$ have a strategy $X_{\{b,c\}}$ such that restriction of M to $X_a \cap X_{\{b,c\}}$ results in some M', and $(M', s) \models \varphi$. Agent a has two strategies in (M_1, s): $\top_a$ and $\{u', t', s, t, u\}_a$. On the other hand, due to the fact that intersection of unions of relations of b and c is an identity relation, agents $\{b, c\}$ have all possible submodels of M_1 that include s as strategies. For either of a's strategies, $\top_a$ or $\{u', t', s, t, u\}_a$, agents $\{b, c\}$ announce $\{t', s, t, u\}_{\{b,c\}}$. (Particularly, b announces $\{u', t', s, t, u\}_b$, and c announces $\{t', s, t, u, v\}_c$.) Such a simultaneous joint announcement results in model M_2, and φ holds in (M_2, s).

Now, let us show that $(M_1, s) \not\models [a]\langle\{b, c\}\rangle\varphi$, which is equivalent, by the semantics, to $(M_1, s) \models \langle a\rangle[\{b, c\}]\neg\varphi$. In terms of strategies, this means that agent a has some strategy X_a in (M_1, s) such that whichever strategy agents $\{b, c\}$ choose in $(M_1, s)^{X_a}$, φ does not hold in the resulting model. Let a announce $\{u', t', s, t, u\}$. Such an announcement makes states u and u', and t and t' bisimilar. The resulting model M_3 and its bisimulation contraction are presented in Figure 5.15.

In model $(\|M_3\|, s)$ agents $\{b, c\}$ have the following strategies: $\{s, t, u\}_{\{b,c\}}$, $\{s, t\}_{\{b,c\}}$, and $\{s\}_{\{b,c\}}$. Results of corresponding updates are presented in Figure 5.16.

It is easy to check that none of the models from Figure 5.16 satisfy φ. Hence, $(M_1, s) \not\models [a]\langle\{b, c\}\rangle\varphi$. □

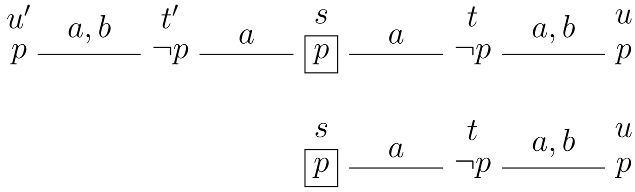

Figure 5.15: Model M_3 and its contraction

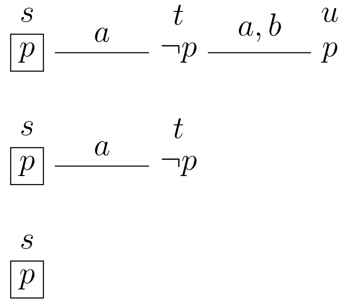

Figure 5.16: Subsequent updates of model $(\|M_3\|, s)$

5.3 Coalition Annoucement Logic Subsumes Coalition Logic

It is known [Ågotnes and van Ditmarsch, 2008] that CAL subsumes CL, i.e. all axioms of CL are valid in CAL, and rules of inference of CL are validity preserving in CAL. However, to the best of our knowledge, a formal proof has not yet been presented.

Proposition 5.15. All of the following are valid and validity preserving in CAL.

$(C0)$ all instantiations of propositional tautologies,
$(C1)$ $\neg\langle\!\langle G\rangle\!\rangle\bot$,
$(C2)$ $\langle\!\langle G\rangle\!\rangle\top$,
$(C3)$ $\neg\langle\!\langle\emptyset\rangle\!\rangle\neg\varphi \to \langle\!\langle A\rangle\!\rangle\varphi$,
$(C4)$ $\langle\!\langle G\rangle\!\rangle(\varphi_1 \wedge \varphi_2) \to \langle\!\langle G\rangle\!\rangle\varphi_1$,
$(C5)$ $\langle\!\langle G\rangle\!\rangle\varphi_1 \wedge \langle\!\langle H\rangle\!\rangle\varphi_2 \to \langle\!\langle G \cup H\rangle\!\rangle(\varphi_1 \wedge \varphi_2)$, if $G \cap H = \emptyset$,
$(R0)$ $\vdash \varphi, \varphi \to \psi \Rightarrow \vdash \psi$,
$(R1)$ $\vdash \varphi \leftrightarrow \psi \Rightarrow \vdash \langle\!\langle G\rangle\!\rangle\varphi \leftrightarrow \langle\!\langle G\rangle\!\rangle\psi$.

Proof. $C0$ and $R0$ are obvious.

$C1$: It holds that $\models \top$, and $\top$ is true in every restriction of a model, i.e. $\models [\psi]\top$. In particular, for some model (M, w) and all true formulas ψ_G and $\chi_{A\setminus G}$: $(M, w) \models \langle\psi_G \wedge \chi_{A\setminus G}\rangle\top$. We can relax the requirement of ψ_G being true by adding the formula as an antecedent. Formally, for all (true and false) ψ_G and some (true) $\chi_{A\setminus G}$: $(M, w) \models \psi_G \to \langle\psi_G \wedge \chi_{A\setminus G}\rangle\top$. The latter is $(M, w) \models [\langle G\rangle]\top$ by the semantics this is equivalent to $(M, w) \models \neg\langle\!\langle G\rangle\!\rangle\bot$ by the duality of the coalition announcement operators.

$C2$: For any pointed model (M, w) and any announcement $\psi_G \wedge \chi_{A\setminus G}$ it holds that $(M, w) \models [\psi_G \wedge \chi_{A\setminus G}]\top$. The latter implies that for some true ψ_G and for all $\chi_{A\setminus G}$: $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\top$, which is $(M, w) \models \langle\!\langle G \rangle\!\rangle\top$ by the semantics.

$C3$: Let $\neg\langle\!\langle \emptyset \rangle\!\rangle\neg\varphi$ be true in some arbitrary pointed model (M, w). This is equivalent to $\exists\psi_A$: $(M, w) \models \neg[\psi_A]\neg\varphi$, which is $(M, w) \models \langle\!\langle A \rangle\!\rangle\varphi$ by the semantics.

$C4$: Suppose that for some (M, w), $(M, w) \models \langle\!\langle G \rangle\!\rangle(\varphi_1 \wedge \varphi_2)$ holds. By the semantics, $\exists\psi_G, \forall\chi_{A\setminus G}$: $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}](\varphi_1 \wedge \varphi_2)$. Then, by axiom of PAL $[\psi](\varphi \wedge \chi) \leftrightarrow [\psi]\varphi \wedge [\psi]\chi$, we have $\exists\psi_G, \forall\chi_{A\setminus G}$: $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi_1 \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi_2$. The latter implies $\exists\psi_G, \forall\chi_{A\setminus G}$: $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi_1$, which is $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi_1$ by the semantics.

$C5$: Assume that for some (M, w) we have that $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi_1 \wedge \langle\!\langle H \rangle\!\rangle\varphi_2$. Let us consider the first conjunct $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi_1$. By the semantics it is equivalent to $\exists\psi_G, \forall\chi_{A\setminus G}$: $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi_1$. Since $G \cap H = \emptyset$, we can split $\chi_{A\setminus G}$ into χ_H and $\chi_{A\setminus G\cup H}$. Thus we have that $\exists\psi_G, \forall\chi_H, \forall\chi_{A\setminus(G\cup H)}$: $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_H \wedge \chi_{A\setminus(G\cup H)}]\varphi_1$. The same holds for the second conjunct: $\exists\psi_H, \forall\chi_G, \forall\chi_{A\setminus(G\cup H)}$: $(M, w) \models \psi_H \wedge [\psi_H \wedge \chi_G \wedge \chi_{A\setminus(G\cup H)}]\varphi_2$. Since χ_H (χ_G) quantifies over all formulas known to H (G), we can substitute χ_H (χ_G) with ψ_H (ψ_G). Hence we have

$$\exists\psi_G, \exists\psi_H, \forall\chi_{A\setminus(G\cup H)}:$$
$$(M, w) \models \psi_G \wedge \psi_H \wedge [\psi_G \wedge \psi_H \wedge \chi_{A\setminus(G\cup H)}]\varphi_1 \wedge [\psi_G \wedge \psi_H \wedge \chi_{A\setminus G\cup H}]\varphi_2.$$

By the axiom of PAL $[\psi](\varphi \wedge \chi) \leftrightarrow [\psi]\varphi \wedge [\psi]\chi$, we have that

$$\exists\psi_G, \exists\psi_H, \forall\chi_{A\setminus(G\cup H)} : (M, w) \models \psi_G \wedge \psi_H \wedge [\psi_G \wedge \psi_H \wedge \chi_{A\setminus(G\cup H)}](\varphi_1 \wedge \varphi_2),$$

and the latter is equivalent to $(M, w) \models \langle\!\langle G \cup H \rangle\!\rangle(\varphi_1 \wedge \varphi_2)$ by the semantics.

$R1$: Assume that $\models \varphi \leftrightarrow \psi$. This means that for any pointed model (M, w) the following holds: $(M, w) \models \varphi$ iff $(M, w) \models \psi$ (1). Now suppose that for some pointed model (M, v) it holds that $(M, v) \models \langle\!\langle G \rangle\!\rangle\varphi$. By the semantics, $\exists\psi_G, \forall\chi_{A\setminus G}$: $(M, v) \models \psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi$, which is equivalent to the following: $(M, v) \models \psi_G$ and $((M, v) \models \psi_G \wedge \chi_{A\setminus G}$ implies $(M^{\psi_G \wedge \chi_{A\setminus G}}, v) \models \varphi)$. By (1) we have that $\exists\psi_G, \forall\chi_{A\setminus G}$: $(M, v) \models \psi_G$ and $((M, v) \models \psi_G \wedge \chi_{A\setminus G}$ implies $(M^{\psi_G \wedge \chi_{A\setminus G}}, v) \models \psi)$, which is $(M, v) \models \langle\!\langle G \rangle\!\rangle\psi$ by the semantics. The same argument holds in the other direction. □

Proposition 5.15 indicates that CAL is indeed a coalition logic in the sense of [Pauly, 2002]. The proof of the proposition is a straightforward application of the semantics of CAL. In [Ågotnes and van Ditmarsch, 2008], however, it was shown that for every epistemic model there is an epistemic coalition model that satisfies exactly the same formulas of the logic. This implies that some of the results presented in the chapter (in particular, Propositions 5.1, 5.8, and 5.15) follow immediately from this fact. Nonetheless, our alternative proofs are interesting in their own right as more ‘direct’ versions of their ‘immediate’ siblings.

Chapter 6

A Logic of Coalition and Relativised Group Announcements

A sound and complete axiomatisation of CAL is an open question. One of the reasons why finding one seems hard is the inherent alternation of quantifiers in the semantics of the coalition announcement operators. In order to mitigate this, we introduce relativised group announcements that allow us to separate a coalition's announcements from counter-announcements by their opponents. The resulting formalism, *Coalition and Relativised Group Announcement Logic* (CoRGAL), is sound and complete. To the best of our knowledge, this is the first axiomatisation of a logic with coalition announcements. CoRGAL is reminiscent of alternating-time temporal dynamic epistemic logic (ATDEL) [de Lima, 2014]. The latter, however, is a more PDL-style logic [Harel et al., 2000] with postconditions and factual change. Moreover, in ATDEL agents are not required to know the formulas they announce. This chapter is based on [Galimullin and Alechina, 2017] (and its corrected version [Galimullin and Alechina, 2018]).

6.1 Syntax, Semantics, and Axiomatisation

6.1.1 Syntax and Semantics

Let P denote a countable set of propositional variables, and A be a finite set of agents.

Definition 6.1 (Language of CoRGAL)**.** The *language of coalition and relativised group announcement logic* $\mathcal{L}_{CoRGAL}$ is as follows:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [\varphi]\varphi \mid [G, \varphi]\varphi \mid [\langle G \rangle]\varphi,$$

where $p \in P$, $a \in A$, $G \subseteq A$, and all the usual abbreviations of propositional logic and conventions for deleting parentheses hold. Diamond version of the operator $[G, \chi]\varphi$ is defined as $\langle G, \chi\rangle\varphi := \neg[G, \chi]\neg\varphi$.

Relativised group announcement $[G, \chi]\varphi$ says that 'given announcement χ, whatever agents from G announce at the same time, they cannot avoid φ.' In this chapter we are interested in χ's that are announced by a (anti-)coalition, i.e. we consider primarily $\chi_{A\setminus G} := \bigwedge_{a\in A\setminus G} K_a\chi_a$ such that $\chi_a \in \mathcal{L}_{EL}$.

Let us recall the definition of necessity forms.

Definition 6.2. (Necessity forms) Let $\varphi \in \mathcal{L}_{CoRGAL}$, then *necessity forms* are inductively defined as follows:

$$\eta(\sharp) ::= \sharp \mid \varphi \to \eta(\sharp) \mid K_a\eta(\sharp) \mid [\varphi]\eta(\sharp).$$

The dual of a necessity form $\eta(\varphi)$ is a *possibility form* $\eta\{\varphi\}$ that is defined as $\eta(\varphi) ::= \neg\eta\{\neg\varphi\}$. The atom $\sharp$ has a unique occurrence in each necessity form. The result of the replacement of $\sharp$ with φ in some $\eta(\sharp)$ is denoted as $\eta(\varphi)$ and is inductively defined as follows:

- $\sharp(\varphi) = \varphi$,
- $(\psi \to \eta)(\varphi) = \psi \to \eta(\varphi)$,
- $(K_a\eta)(\varphi) = K_a\eta(\varphi)$,
- $([\psi]\eta)(\varphi) = [\psi]\eta(\varphi)$.

Formulas of CoRGAL are interpreted on epistemic models introduced in Definition 2.2.

Definition 6.3 (Semantics of CoRGAL)**.** Let a pointed model (M, w) with $M = (W, \sim, V)$, $a \in A$, $G \subseteq A$, $\psi_G \in \mathcal{L}_{EL}$ and $\varphi, \psi, \chi \in \mathcal{L}_{CoRGAL}$ be given. The *semantics of coalition and relativised group announcement logic* is presented below.

$$\begin{array}{lll}
(M,w) \models p & \text{iff} & w \in V(p) \\
(M,w) \models \neg\varphi & \text{iff} & (M,w) \not\models \varphi \\
(M,w) \models \varphi \wedge \psi & \text{iff} & (M,w) \models \varphi \text{ and } (M,w) \models \psi \\
(M,w) \models K_a\varphi & \text{iff} & \text{for all } v \in W : w \sim_a v \text{ implies } (M,v) \models \varphi \\
(M,w) \models [\psi]\varphi & \text{iff} & (M,w) \models \psi \text{ implies } (M,w)^\psi \models \varphi \\
(M,w) \models [G,\chi]\varphi & \text{iff} & (M,w) \models \chi \text{ and } \forall\psi_G : (M,w) \models [\psi_G \wedge \chi]\varphi \\
(M,w) \models [\langle G\rangle]\varphi & \text{iff} & \forall\psi_G, \exists\chi_{A\setminus G} : (M,w) \models \psi_G \\
 & & \text{implies } (M,w) \models \langle\psi_G \wedge \chi_{A\setminus G}\rangle\varphi
\end{array}$$

Note that as in GAL and CAL we restrict the quantification in $[G, \chi]\varphi$ and $[\langle G\rangle]\varphi$ to formulas of epistemic logic. This allows us to avoid circularity in the definition of semantics.

Semantics for duals of relativised group announcements and coalition announcements is as follows.

$$\begin{array}{lll}
(M,w) \models \langle G,\chi\rangle\varphi & \text{iff} & (M,w) \models \chi \text{ implies } \exists\psi_G : (M,w) \models \langle\psi_G \wedge \chi\rangle\varphi \\
(M,w) \models \langle[G]\rangle\varphi & \text{iff} & \exists\psi_G, \forall\chi_{A\setminus G} : (M,w) \models \psi_G \text{ and } (M,w) \models [\psi_G \wedge \chi_{A\setminus G}]\varphi
\end{array}$$

Note that semantics of coalition announcement operators are given in a 'classic' way. An equivalent definition is possible using relativised group announcements.

$$(M,w) \models [\langle G \rangle]\varphi \text{ iff } \forall \psi_G : (M,w) \models \langle A \setminus G, \psi_G \rangle \varphi$$
$$(M,w) \models \langle [G] \rangle\varphi \text{ iff } \exists \psi_G : (M,w) \models [A \setminus G, \psi_G] \varphi$$

6.1.2 Relativised Group Announcements

Relativised group announcements help us to 'split' coalition announcements, and treat the coalition's announcement and anti-coalition responses separately. Note that the syntax of such announcements is very similar to the one for relativised common knowledge.

Next we show some intuitive properties of relativised group announcements.

Proposition 6.1. All of the following are valid:

1. $[G]\varphi \leftrightarrow [G,\top]\varphi$
2. $[\emptyset,\psi]\varphi \leftrightarrow \langle\psi\rangle\varphi$
3. $[A,\psi]\varphi \rightarrow \langle\psi\rangle\varphi$
4. $\neg\chi \rightarrow \langle G,\chi\rangle\varphi$

Proof. 1. Trivial by the definition of semantics (Definition 6.3).

2. By property 2 of Proposition 3.3.
3. By the fact that group A can always announce $\top_A := \bigwedge_{a\in A} K_a\top$.
4. Assume that for some arbitrary pointed model (M,w) we have that $(M,w) \models \neg\chi$. By propositional reasoning we have that $(M,w) \models \neg\chi \rightarrow (\chi \rightarrow \psi)$ for any $\psi \in \mathcal{L}_{CoRGAL}$. In particular, $(M,w) \models \neg\chi \rightarrow (\chi \rightarrow \langle\psi_G \wedge \chi\rangle\varphi)$ for any ψ_G. The latter implies $(M,w) \models \neg\chi \rightarrow \langle G,\chi\rangle\varphi$ by the semantics.

□

The first property states that classic group announcements can be defined using relativised group announcement. Indeed, announcing a tautology in conjunction with an announcement by a group does not have any additional effect on the resulting model. Validities 2 and 3 demonstrate the relation between public announcements and relativised group announcements with empty and grand groups. Note that the property 3 holds only in one direction. A counterexample for the other direction would be a model with two states such that p holds only in one of them, agent's a relation is universal, and agent's b relation is the identity. If $\psi := p \vee \neg p$ and $\varphi := \neg K_a p$, then $\langle p \vee \neg p\rangle \neg K_a p$ is true, and $[\{a,b\}, p \vee \neg p]\neg K_a p$ is false in the p-state (since agent b can announce $K_b p$). Formula 4 says that if a false formula is being announced, we can always add a group announcement such that any φ holds vacuously afterwards.

6.1.3 Axiom System of CoRGAL

In this section we present an axiomatisation of CoRGAL and show its soundness. The axiomatisation based on the axiom system for PAL, and have two additional axioms and four additional rules of inference.

Definition 6.4 (Axiomatisation of CoRGAL)**.** The *axiom system for* CoRGAL is an extension of PAL with a relativised version of GAL and interaction axioms.

$(A0)$	propositional tautologies,
$(A1)$	$K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$,
$(A2)$	$K_a\varphi \to \varphi$,
$(A3)$	$K_a\varphi \to K_aK_a\varphi$,
$(A4)$	$\neg K_a\varphi \to K_a\neg K_a\varphi$,
$(A5)$	$[\varphi]p \leftrightarrow (\varphi \to p)$,
$(A6)$	$[\varphi]\neg\psi \leftrightarrow (\varphi \to \neg[\varphi]\psi)$,
$(A7)$	$[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$,
$(A8)$	$[\varphi]K_a\psi \leftrightarrow (\varphi \to K_a[\varphi]\psi)$,
$(A9)$	$[\varphi][\psi]\chi \leftrightarrow [\varphi \wedge [\varphi]\psi]\chi$,
$(A10)$	$[G,\chi]\varphi \to \chi \wedge [\psi_G \wedge \chi]\varphi$ for any ψ_G,
$(A11)$	$[\langle G\rangle]\varphi \to \langle A \setminus G, \psi_G\rangle\varphi$ for any ψ_G,
$(R0)$	If $\vdash \varphi$ and $\vdash \varphi \to \psi$, then $\vdash \psi$,
$(R1)$	If $\vdash \varphi$, then $\vdash K_a\varphi$,
$(R2)$	If $\vdash \varphi$, then $\vdash [\psi]\varphi$,
$(R3)$	If $\vdash \varphi$, then $\vdash [G,\chi]\varphi$,
$(R4)$	If $\vdash \varphi$, then $\vdash [\langle G\rangle]\varphi$,
$(R5)$	If $\forall\psi_G :\vdash \eta(\chi \wedge [\psi_G \wedge \chi]\varphi)$, then $\vdash \eta([G,\chi]\varphi)$,
$(R6)$	If $\forall\psi_G :\vdash \eta(\langle A \setminus G, \psi_G\rangle\varphi)$, then $\vdash \eta([\langle G\rangle]\varphi)$.

We call CoRGAL the smallest subset of $\mathcal{L}_{CoRGAL}$ that contains all the axioms $A0 - A11$ and is closed under rules of inference $R0 - R6$. Elements of CoRGAL are called *theorems*. Note that $R5$ and $R6$ are infinitary rules: they require an infinite number of premises.

Axiom $A10$ says that if given some χ, agents from G cannot avoid φ no matter what they announce, they cannot avoid φ making any particular announcement in this situation. The property expressed by $A11$ is as follows: if for every announcement by G there is a 'counter-announcement' by $A \setminus G$, then for some particular announcement ψ_G by G there is a 'counter-announcement' by $A \setminus G$. Rules $R3$ and $R4$ are necessitation rules for relativised group announcements and coalition announcements. Informally, they express the fact that if φ is valid, φ remains so no matter what agents announce. Rules $R5$ and $R6$ demonstrate how to infer formulas with relativised group and coalition announcements from an infinite number of premises.

Proposition 6.2. Axioms $A10$ and $A11$ are valid.

Proof. Follows directly from the definition of semantics (Definition 6.3). We just show validity of $(A11)$.

Assume that for some arbitrary pointed model (M, w) is holds that $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi$. By semantics this is equivalent to $\forall\psi_G, \exists\chi_{A\setminus G}$: $(M, w) \models \psi_G \rightarrow \langle\psi_G \land \chi_{A\setminus G}\rangle\varphi$. Since ψ_G quantifies over all epistemic formulas known to G, we can choose any particular ψ_G. Hence, we have that $\exists\chi_{A\setminus G}$: $(M, w) \models \psi_G \rightarrow \langle\psi_G \land \chi_{A\setminus G}\rangle\varphi$, which is equivalent to $(M, w) \models \langle A \setminus G, \psi_G\rangle\varphi$ by semantics. □

Proposition 6.3. *R*5 and *R*6 are truth-preserving.

Proof. The proof is by induction on the construction of necessity forms.

($R5$) *Base case.* If for all ψ_G we have that $(M, w) \models \chi \land [\psi_G \land \chi]\varphi$, then this is equivalent to $(M, w) \models [G, \chi]\varphi$ by the semantics.

Induction Hypothesis. If for some (M, w) it holds that $(M, w) \models \eta(\chi \land [\psi_G \land \chi]\varphi)$ for all ψ_G, then $(M, w) \models \eta([G, \chi]\varphi)$.

Case $\forall\psi_G$: $\tau \rightarrow \eta(\chi \land [\psi_G \land \chi]\varphi)$ for some $\tau \in \mathcal{L}_{CoRGAL}$. This means that $(M, w) \models \neg\tau$ or $(M, w) \models \eta(\chi \land [\psi_G \land \chi]\varphi)$. By Induction Hypothesis we have that $(M, w) \models \neg\tau$ or $(M, w) \models \eta([G, \chi]\varphi)$, which is equivalent to $(M, w) \models \tau \rightarrow \eta([G, \chi]\varphi)$.

Case $\forall\psi_G$: $K_a\eta(\chi \land [\psi_G \land \chi]\varphi)$ for some $a \in A$. By semantics we have that for every $v \in W$: $w \sim_a v$ implies $(M, v) \models \eta(\chi \land [\psi_G \land \chi]\varphi)$. By Induction Hypothesis we conclude that for every $v \in W$: $w \sim_a v$ implies $(M, v) \models \eta([G, \chi]\varphi)$, which is equivalent to $(M, w) \models K_a\eta([G, \chi]\varphi)$.

Case $\forall\psi_G$: $[\tau]\eta(\chi \land [\psi_G \land \chi]\varphi)$ for some $\tau \in \mathcal{L}_{CoRGAL}$. This means that $(M, w) \models \tau$ implies $(M, w)^\tau \models \eta(\chi \land [\psi_G \land \chi]\varphi)$. By Induction Hypothesis we have that $(M, w) \models \tau$ implies $(M, w)^\tau \models \eta([G, \chi]\varphi)$, which is equivalent to $(M, w) \models [\tau]\eta([G, \chi]\varphi)$.

($R6$) *Base case.* If for all ψ_G we have that $(M, w) \models \langle A \setminus G, \psi_G\rangle\varphi$, then this is equivalent to $(M, w) \models \langle\!\langle G \rangle\!\rangle\varphi$ by the semantics.

Induction Hypothesis. If for some (M, w) it holds that $(M, w) \models \eta(\langle A \setminus G, \psi_G\rangle\varphi)$ for all ψ_G, then $(M, w) \models \eta(\langle\!\langle G \rangle\!\rangle\varphi)$.

Case $\forall\psi_G$: $\tau \rightarrow \eta(\langle A \setminus G, \psi_G\rangle\varphi)$ for some $\tau \in \mathcal{L}_{CoRGAL}$. This means that $(M, w) \models \neg\tau$ or $(M, w) \models \eta(\langle A \setminus G, \psi_G\rangle\varphi)$. By Induction Hypothesis we have that $(M, w) \models \neg\tau$ or $(M, w) \models \eta(\langle\!\langle G \rangle\!\rangle\varphi)$, which is equivalent to $(M, w) \models \tau \rightarrow \eta(\langle\!\langle G \rangle\!\rangle\varphi)$.

Case $\forall\psi_G$: $K_a\eta(\langle A \setminus G, \psi_G\rangle\varphi)$ for some $a \in A$. By semantics we have that for every $v \in W$: $w \sim_a v$ implies $(M, v) \models \eta(\langle A \setminus G, \psi_G\rangle\varphi)$. By Induction Hypothesis we conclude that for every $v \in W$: $w \sim_a v$ implies $(M, v) \models \eta(\langle\!\langle G \rangle\!\rangle\varphi)$, which is equivalent to $(M, w) \models K_a\eta(\langle\!\langle G \rangle\!\rangle\varphi)$.

Case $\forall\psi_G$: $[\tau]\eta(\langle A \setminus G, \psi_G\rangle\varphi)$ for some $\tau \in \mathcal{L}_{CoRGAL}$. This means that $(M, w) \models \tau$ implies $(M, w)^\tau \models \eta(\langle A \setminus G, \psi_G\rangle\varphi)$. By Induction Hypothesis we have that $(M, w) \models \tau$ implies $(M, w)^\tau \models \eta(\langle\!\langle G \rangle\!\rangle\varphi)$, which is equivalent to $(M, w) \models [\tau]\eta(\langle\!\langle G \rangle\!\rangle\varphi)$. □

Theorem 6.4 (Soundness)**.** For all $\varphi \in \mathcal{L}_{CoRGAL}$, if $\varphi \in$ CoRGAL, then φ is valid.

Proof. Soundness of *A*0–*A*4, *R*0, and *R*1 is due to soundness of EL. Axioms *A*5–*A*9 and rule of inference *R*2 are sound, since PAL is sound [van Ditmarsch et al., 2008, Chapter 4]. Soundness of $R5$ and $R6$ follows from Proposition 6.3, and validity of $A10$ and $A11$ is shown in Proposition 6.2. So, it is only left to show that $R3$ and $R4$ are sound. Proofs for both of the rules are similar, and we present only the proof for $R4$.

($R4$) Assume $\models \varphi$. Since public announcements preserve validity, we have that for any (M, w) and ψ, $(M, w) \models [\psi]\varphi$. Since ψ is arbitrary, we have that for all ψ_G and $\chi_{A\backslash G}$ $(M, w) \models [\psi_G \wedge \chi_{A\backslash G}]\varphi$. The latter implies that for some true ψ_G it holds that $(M, w) \models \psi_G \wedge [\psi_G \wedge \chi_{A\backslash G}]\varphi$, which is $(M, w) \models [\langle G\rangle]\varphi$ by semantics. Since (M, w) was arbitrary, we conclude that $\models [\langle G\rangle]\varphi$. □

Note that unlike $R3$ and $R4$, rules $R5$ and $R6$ preserve not only validity, but truth as well (Proposition 6.3). This fact is quite important in the proof of the completeness of the logic, since we use closure on these rules (and modus ponens, which is also truth-preserving) to define theories. As was pointed in Section 3.1, there are multiple rules of inference we would like to use to deduce formulas with relativised group and coalition announcements. It is easy to see that rule

$$\text{If } \forall \psi_G :\vdash \chi \wedge [\psi_G \wedge \chi]\varphi, \text{then} \vdash [G, \chi]\varphi$$

is truth preserving. So is, for example,

$$\text{If } \forall \psi_G :\vdash K_a(\chi \wedge [\psi_G \wedge \chi]\varphi), \text{then} \vdash K_a([G, \chi]\varphi),$$

or

$$\text{If } \forall \psi_G :\vdash \tau \rightarrow [\theta](\chi \wedge [\psi_G \wedge \chi]\varphi), \text{then} \vdash \tau \rightarrow [\theta]([G, \chi]\varphi)$$

and so on. Necessity forms succinctly capture such a plethora of rules of inference.

6.2 Completeness

In order to prove completeness of CoRGAL, we expand and modify the completeness proof for APAL [Balbiani et al., 2008; Balbiani and van Ditmarsch, 2015; Balbiani, 2015]. Although the proof is partially based upon the classic canonical model approach, we have to ensure that construction of maximal consistent theories (Proposition 6.8) allows us to include an infinite amount of formulas for cases of coalition announcements. This is possible due to axioms $A10$, $A11$ and rules of inference $R5, R6$. After that we use induction on the complexity of CoRGAL formulas to prove Truth Lemma.

First, we prove a useful auxiliary lemma.

Lemma 6.5. Let $\varphi, \psi \in \mathcal{L}_{CoRGAL}$. If $\varphi \rightarrow \psi$ is a theorem, then $\eta(\varphi) \rightarrow \eta(\psi)$ is a theorem as well.

Proof. Assume that $\varphi \rightarrow \psi$ is a theorem. We prove the lemma by induction on η.

Base case $\eta := \sharp$. Formula $\varphi \rightarrow \psi$ is a theorem by assumption.

Induction Hypothesis. Assume that for some η, $\eta(\varphi) \rightarrow \eta(\psi)$ is a theorem.

Case $(\tau \to \eta(\varphi)) \to (\tau \to \eta(\psi))$ for some $\tau \in \mathcal{L}_{CoRGAL}$. Formula $(\eta(\varphi) \to \eta(\psi)) \to ((\tau \to \eta(\varphi)) \to (\tau \to \eta(\psi)))$ is a propositional tautology, and, hence, a theorem of CoRGAL. Using Induction Hypothesis and $R0$, we have that $(\tau \to \eta(\varphi)) \to (\tau \to \eta(\psi))$ is a theorem.

Case $(K_a\eta(\varphi)) \to (K_a\eta(\psi))$ for some $a \in A$. Since $\eta(\varphi) \to \eta(\psi)$ is a theorem by Induction Hypothesis, $K_a(\eta(\varphi) \to \eta(\psi))$ is also a theorem by $R1$. Next, $K_a(\eta(\varphi) \to \eta(\psi)) \to (K_a\eta(\varphi) \to K_a\eta(\psi))$ is an instance of $A1$, and, hence, a theorem. Finally, using $R0$ we have that $K_a\eta(\varphi) \to K_a\eta(\psi)$ is a theorem.

Case $([\tau]\eta(\varphi)) \to ([\tau]\eta(\psi))$ for some $\tau \in \mathcal{L}_{CoRGAL}$. Formula $[\tau](\eta(\varphi) \to \eta(\psi)) \to ([\tau]\eta(\varphi) \to [\tau]\eta(\psi))$ is a theorem of PAL (see [van Ditmarsch et al., 2008, Chapter 4]), and hence of CoRGAL. Using Induction Hypothesis and $R0$ we conclude that $[\tau]\eta(\varphi) \to [\tau]\eta(\psi)$ is also a theorem of CoRGAL. □

Now, the first part of the proof up to Proposition 6.8 is based on [Balbiani et al., 2008] and [Goldblatt, 1982, Chapter 2]. Here we introduce theories and prove Lindenbaum Lemma.

Definition 6.5 (Theory)**.** A set of formulas x is called a *theory* if and only if it contains CoRGAL, and is closed under $R0, R5$, and $R6$. A theory x is consistent if and only if $\bot \notin x$, and is maximal if and only if for all $\varphi \in \mathcal{L}_{CoRGAL}$ it holds that either $\varphi \in x$ or $\neg\varphi \in x$.

Note that theories are not closed under necessitation rules. The reason for this is that while these rules preserve validity, they do not preserve truth, whereas $R0, R5$, and $R6$ preserve both validity and truth.

Proposition 6.6. Let x be a theory, $\varphi, \psi \in \mathcal{L}_{CoRGAL}$, and $a \in A$. The following are theories: $x + \varphi = \{\psi : \varphi \to \psi \in x\}, K_a x = \{\varphi : K_a\varphi \in x\}$, and $[\varphi]x = \{\psi : [\varphi]\psi \in x\}$.

Proof. Let ψ be a theorem, i.e. $\psi \in$ CoRGAL. Then $\varphi \to \psi$ is also a theorem, since $\psi \to (\varphi \to \psi) \in$ CoRGAL and CoRGAL is closed under $R0$. Moreover, $K_a\psi$ and $[\varphi]\psi$ are theorems as well due to the fact that CoRGAL is closed under $R1$ and $R2$. Therefore, $\psi \in x+\varphi$, $\psi \in K_a x$, and $\psi \in [\varphi]x$, and hence CoRGAL $\subseteq x + \varphi, K_a x, [\varphi]x$.

The rest of the proof is an extension of the one from [Balbiani et al., 2008], where it was shown that $x + \varphi$, $K_a x$, and $[\varphi]x$ are closed under $R0$. We argue that corresponding sets are closed under $R5$ and $R6$.

Case $x + \varphi$. Suppose that $\eta(\chi \wedge [\psi_G \wedge \chi]\tau) \in x + \varphi$ for some given χ, for all ψ_G, and for some $\tau \in \mathcal{L}_{CoRGAL}$. This means that $\varphi \to \eta(\chi \wedge [\psi_G \wedge \chi]\tau) \in x$ for all ψ_G. Since $\varphi \to \eta(\chi \wedge [\psi_G \wedge \chi]\tau)$ is a necessity form, and x is closed under $R5$ (by Definition 6.5), we infer that $\varphi \to \eta([G, \chi]\tau) \in x$, and, consequently, $\eta([G, \chi]\tau) \in x + \varphi$. So, $x + \varphi$ is closed under $R5$.

Now, let $\forall\psi_G$: $\eta(\langle A \setminus G, \psi_G\rangle\tau) \in x + \varphi$. By the definition of $x + \varphi$ this means that $\varphi \to \eta(\langle A \setminus G, \psi_G\rangle\tau) \in x$ for all ψ_G. Since $\varphi \to \eta(\langle A \setminus G, \psi_G\rangle\tau$ is a necessity form and x is closed under $R6$, we infer that $\varphi \to \eta([\langle G\rangle]\tau) \in x$, and, consequently, $\eta([\langle G\rangle]\tau) \in x + \varphi$. So, $x + \varphi$ is closed under $R6$.

Case $K_a x$. Suppose that $\eta(\chi \wedge [\psi_G \wedge \chi]\tau) \in K_a x$ for some given χ, for all ψ_G, and for some $\tau \in \mathcal{L}_{CoRGAL}$. This means that $K_a \eta(\chi \wedge [\psi_G \wedge \chi]\tau) \in x$ for all ψ_G. Since $K_a \eta(\chi \wedge [\psi_G \wedge \chi]\tau)$ is a necessity form, and x is closed under $R5$ (by Definition 6.5), we infer that $K_a \eta([G, \chi]\tau) \in x$, and, consequently, $\eta([G, \chi]\tau) \in K_a x$. So, $K_a x$ is closed under $R5$.

Now, let $\forall \psi_G$: $\eta(\langle A \setminus G, \psi_G \rangle \tau) \in K_a x$. By the definition of $K_a x$ this means that $K_a \eta(\langle A \setminus G, \psi_G \rangle \tau) \in x$ for all ψ_G. Since $K_a \eta(\langle A \setminus G, \psi_G \rangle \tau$ is a necessity form and x is closed under $R6$, we infer that $K_a \eta([\langle G \rangle]\tau) \in x$, and, consequently, $\eta([\langle G \rangle]\tau) \in K_a x$. So, $K_a x$ is closed under $R6$.

Case $[\varphi]x$. Finally, suppose that $\eta(\chi \wedge [\psi_G \wedge \chi]\tau) \in [\varphi]x$ for some given χ, for all ψ_G, and for some $\tau \in \mathcal{L}_{CoRGAL}$. This means that $[\varphi]\eta(\chi \wedge [\psi_G \wedge \chi]\tau) \in x$ for all ψ_G. Since $[\varphi]\eta(\chi \wedge [\psi_G \wedge \chi]\tau)$ is a necessity form, and x is closed under $R5$ (by Definition 6.5), we infer that $[\varphi]\eta([G, \chi]\tau) \in x$, and, consequently, $\eta([G, \chi]\tau) \in [\varphi]x$. So, $[\varphi]x$ is closed under $R5$.

Now, let $\forall \psi_G$: $\eta(\langle A \setminus G, \psi_G \rangle \tau) \in [\varphi]x$. By the definition of $[\varphi]x$ this means that $[\varphi]\eta(\langle A \setminus G, \psi_G \rangle \tau) \in x$ for all ψ_G. Since $[\varphi]\eta(\langle A \setminus G, \psi_G \rangle \tau$ is a necessity form and x is closed under $R6$, we infer that $[\varphi]\eta([\langle G \rangle]\tau) \in x$, and, consequently, $\eta([\langle G \rangle]\tau) \in [\varphi]x$. So, $[\varphi]x$ is closed under $R6$. □

Proposition 6.7. Let $\varphi \in \mathcal{L}_{CoRGAL}$. Then CoRGAL $+ \varphi$ is consistent iff $\neg\varphi \notin$ CoRGAL.

Proof. From left to right. Suppose to the contrary that CoRGAL$+\varphi$ is consistent and $\neg\varphi \in$ CoRGAL. Then having both φ and $\neg\varphi$ means that $\bot \in$ CoRGAL$+\varphi$, which contradicts to CoRGAL $+ \varphi$ being consistent.

From right to left. Let us consider the contrapositive: if CoRGAL $+ \varphi$ is inconsistent, then $\neg\varphi \in$ CoRGAL. Since CoRGAL $+ \varphi$ is inconsistent, $\bot \in$ CoRGAL $+ \varphi$, or, by Proposition 6.6, $\varphi \to \bot \in$ CoRGAL. By consistency of CoRGAL and propositional reasoning, we have that $\neg\varphi \in$ CoRGAL. □

The following proposition is a variation of Lindenbaum Lemma. In order to prove it, we rely heavily on rules of inference $R5$ and $R6$.

Lemma 6.8 (Lindenbaum)**.** Every consistent theory x can be extended to a maximal consistent theory y.

Proof. Let $\psi_0, \psi_1, \ldots$ be an enumeration of formulas of the language, and let $y_0 = x$. Suppose that for some $n \geq 0$, y_n is a consistent theory, and $x \subseteq y_n$. If $y_n + \psi_n$ is consistent (i.e. if $\neg\psi_n \notin y_n$), then $y_{n+1} = y_n + \psi_n$. Otherwise, if ψ_n is not a conclusion of either $R5$ or $R6$, $y_{n+1} = y_n$.

If ψ_n is a conclusion of $R5$, that is if ψ_n is of the form $\eta([G, \chi]\varphi)$, then $y_{n+1} = y_n + \neg\eta(\chi \wedge [\psi_G \wedge \chi]\varphi)$, where $\neg\eta(\chi \wedge [\psi_G \wedge \chi]\varphi)$ is the first formula in the enumeration such that $\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \notin y_n$. Theory y_{n+1} is consistent due to the fact that if $\neg\eta([G, \chi]\varphi) \in y_n$, then there must exist some ψ_G such that $\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \notin y_n$, for otherwise $R5$ would lead to $\eta([G, \chi]\varphi) \in y_n$, which contradicts the assumption that $\eta([G, \chi]\varphi) \notin y_n$ and consistency of y_n. We pick the first formula $\neg\eta(\chi \wedge [\psi_G \wedge \chi]\varphi)$ in the enumeration, and have that

$\neg\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \in y_{n+1}$ and $\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \notin y_{n+1}$. Note that adding such a witness ψ_G corresponds to the semantics of relativised group announcements, i.e. for formula $\eta\{\langle G, \chi\rangle\neg\varphi\}$ we have ψ_G such that $\eta\{\chi \to \langle\psi_G \wedge \chi\rangle\neg\varphi\}$.

If ψ_n is a conclusion of $R6$, that is if ψ_n is of the form $\eta([\langle G\rangle]\varphi)$, then $y_{n+1} = y_n + \neg\eta(\langle A \setminus G, \psi_G\rangle\varphi)$, where $\neg\eta(\langle A \setminus G, \psi_G\rangle\varphi)$ is the first formula in the enumeration such that $\eta(\langle A \setminus G, \psi_G\rangle\varphi) \notin y_n$. Theory y_{n+1} is consistent due to the fact that if $\neg\eta([\langle G\rangle]\varphi) \in y_n$, then there must exist some ψ_G such that $\eta(\langle A \setminus G, \psi_G\rangle\varphi) \notin y_n$, for otherwise $R6$ would lead to $\eta([\langle G\rangle]\varphi) \in y_n$, which contradicts the assumption that $\eta([\langle G\rangle]\varphi) \notin y_n$ and consistency of y_n. We pick the first formula $\neg\eta(\langle A \setminus G, \psi_G\rangle\varphi)$ in the enumeration, and have that $\neg\eta(\langle A\setminus G, \psi_G\rangle\varphi) \in y_{n+1}$ and $\eta(\langle A\setminus G, \psi_G\rangle\varphi) \notin y_{n+1}$. Note that since for all $\chi_{A\setminus G}$: $\eta([A \setminus G, \psi_G]\varphi) \to \eta(\psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\varphi)$ are theorems, they and their contrapositions (due to Proposition 6.5) are already in y_n (because CoRGAL $\subseteq x \subseteq y_n$). Thus, adding $\neg\eta(\langle A\setminus G, \psi_G\rangle\varphi)$ to y_n adds all the $\neg\eta(\psi_G \to \langle\psi_G \wedge \chi_{A\setminus G}\rangle\varphi)$ for $\chi_{A\setminus G}$ as well. This satisfies the semantics of coalition announcements, i.e. for formula $\eta\{\langle\!\langle G\rangle\!\rangle\neg\varphi\}$ we have some ψ_G such that for all $\chi_{A\setminus G}$: $\eta\{\psi_G \wedge [\psi_G \wedge \chi_{A\setminus G}]\neg\varphi\}$.

Next we need to show that $y = \bigcup_{n=0}^{\infty} y_n$ is a maximal consistent theory. First, we argue that y is consistent, i.e. that if $\psi \in y$, then $\neg\psi \notin y$. Suppose towards a contradiction that $\psi, \neg\psi \in y$. This means that there is n, such that $\psi, \neg\psi \in y_n$, which contradicts y_n being a consistent theory.

Now we argue that y is a theory, i.e. CoRGAL $\subset y$ (1), and y is closed under $R0$ (2), $R5$ (3), and $R6$ (4).

1. Since $x \subseteq y$, we have that CoRGAL $\subseteq x \subset y$.

2. Assume that $\psi \to \varphi \in y$, and $\psi \in y$. This means that there is n such that $\psi \to \varphi, \psi \in y_n$. The latter implies that $\varphi \in y_n$, and thus $\varphi \in y$. Therefore, y is closed under $R0$.

3. Let some ψ_n be $\eta([G, \chi]\varphi)$. By the construction of y_{n+1}, if $\neg\eta([G, \chi]\varphi) \notin y_n$, then $\eta([G, \chi]\varphi) \in y_n$ and hence $\eta([G, \chi]\varphi) \in y$. Since CoRGAL $\subset y$ and y is closed under $R0$, this means that $\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \in y$ for all ψ_G. If $\neg\eta([G, \chi]\varphi) \in y_n$, then $\neg\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \in y_{n+1} \in y$, and hence $\neg\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \in y$. By the consistency of y, we have that $\eta(\chi \wedge [\psi_G \wedge \chi]\varphi) \notin y$, and therefore y is closed under $R5$.

4. Let some ψ_n be $\eta([\langle G\rangle]\varphi)$. By the construction of y_{n+1}, if $\neg\eta([\langle G\rangle]\varphi) \notin y_n$, then $\eta([\langle G\rangle]\varphi) \in y_n$ and hence $\eta([\langle G\rangle]\varphi) \in y$. Since CoRGAL $\subset y$ and y is closed under $R0$, this means that $\eta(\langle A \setminus G, \psi_G\rangle\varphi) \in y$ for all ψ_G. If $\neg\eta([\langle G\rangle]\varphi) \in y_n$, then $\neg\eta(\langle A \setminus G, \psi_G\rangle\varphi) \in y_{n+1}$, and hence $\neg\eta(\langle A \setminus G, \psi_G\rangle\varphi) \in y$. By the consistency of y, we have that $\eta(\langle A \setminus G, \psi_G\rangle\varphi) \notin y$, and therefore y is closed under $R6$.

Finally, we show that y is maximal. For any formula ψ_n in the enumeration, we have that either $\neg\psi_n \notin y_n$, and thus $\psi_n \in y_{n+1} \subseteq y$, or we have that $\neg\psi_n \in y_n$ and hence $\neg\psi_n \in y$. Therefore, for any ψ_n we have that either $\psi_n \in y$ or $\neg\psi_n \in y$. □

The rest of the proof is an expansion of the one from [Balbiani and van Ditmarsch, 2015]. It employs induction on complexity of formulae to prove Truth Lemma (Proposition 6.10) and, ultimately, completeness (Proposition 6.11) of CoRGAL.

Definition 6.6 (Size)**.** The *size* of some formula $\varphi \in \mathcal{L}_{CoRGAL}$ is defined as follows:

1. $Size(p) = 1$,
2. $Size(\neg\varphi) = Size(K_a\varphi) = Size([G,\chi]\varphi) = Size([\langle G\rangle]\varphi) = Size(\varphi) + 1$,
3. $Size(\varphi \wedge \psi) = Size(\varphi) + Size(\psi) + 1$,
4. $Size([\psi]\varphi) = Size(\psi) + 3 \cdot Size(\varphi)$.

The $[,]$-*depth* is defined as follows:

1. $d_{[,]}(p) = 0$,
2. $d_{[,]}(\neg\varphi) = d_{[,]}(K_a\varphi) = d_{[,]}([\langle G\rangle]\varphi) = d_{[,]}(\varphi)$,
3. $d_{[,]}(\varphi \wedge \psi) = \max\{d_{[,]}(\varphi), d_{[,]}(\psi)\}$,
4. $d_{[,]}([\psi]\varphi) = d_{[,]}(\psi) + d_{[,]}(\varphi)$,
5. $d_{[,]}([G,\chi]\varphi) = d_{[,]}(\varphi) + d_{[,]}(\chi) + 1$.

The $[\langle\rangle]$-*depth* is the same as $[,]$, with the following exceptions.

1. $d_{[\langle\rangle]}([G,\chi]\varphi) = d_{[\langle\rangle]}(\varphi) + d_{[\langle\rangle]}(\chi)$,
2. $d_{[\langle\rangle]}([\langle G\rangle]\varphi) = d_{[\langle\rangle]}(\varphi) + 1$.

Definition 6.7 (Size Relation)**.** The binary relation $<^{Size}_{[,],[\langle\rangle]}$ between $\varphi, \psi \in \mathcal{L}_{CoRGAL}$ is defined as follows:

$\varphi <^{Size}_{[,],[\langle\rangle]} \psi$ iff $d_{[\langle\rangle]}(\varphi) < d_{[\langle\rangle]}(\psi)$, or, otherwise, $d_{[\langle\rangle]}(\varphi) = d_{[\langle\rangle]}(\psi)$, and either $d_{[,]}(\varphi) < d_{[,]}(\psi)$, or $d_{[,]}(\varphi) = d_{[,]}(\psi)$ and $Size(\varphi) < Size(\psi)$. The relation is a well-founded strict partial order between formulas. Note that for all epistemic formulas ψ we have that $d_{[,]}(\psi) = d_{[\langle\rangle]}(\psi) = 0$.

We need the following size inequalities between formulas in our proof of the Truth Lemma.

Proposition 6.9. Let ψ_G and $G \subseteq A$ be given, and let $\chi, \varphi, \tau \in$ CoRGAL.

1. $\chi \wedge [\psi_G \wedge \chi]\varphi <^{Size}_{[,],[\langle\rangle]} [G,\chi]\varphi$,
2. $[\tau](\chi \wedge [\psi_G \wedge \chi]\varphi) <^{Size}_{[,],[\langle\rangle]} [\tau][G,\chi]\varphi$,
3. $\langle A \setminus G, \psi_G\rangle\varphi <^{Size}_{[,],[\langle\rangle]} [\langle G\rangle]\varphi$,

4. $[\tau]\langle A \setminus G, \psi_G\rangle\varphi <^{Size}_{[,],[\langle\rangle]} [\tau][\langle G\rangle]\varphi$.

Proof. 1. Note that $[\langle\rangle]$-depth for both sides of the inequality is the same and equals $d_{[\langle\rangle]}(\chi) + d_{[\langle\rangle]}(\varphi)$. In particular, we have the following for the left-hand side: $d_{[\langle\rangle]}(\chi \wedge [\psi_G \wedge \chi]\varphi) = \max\{d_{[\langle\rangle]}(\chi), d_{[\langle\rangle]}([\psi_G \wedge \chi]\varphi)\} = d_{[\langle\rangle]}([\psi_G \wedge \chi]\varphi) = d_{[\langle\rangle]}(\psi_G \wedge \chi) + d_{[\langle\rangle]}(\varphi) = \max\{d_{[\langle\rangle]}(\psi_G), d_{[\langle\rangle]}(\chi)\} + d_{[\langle\rangle]}(\varphi) = d_{[\langle\rangle]}(\chi) + d_{[\langle\rangle]}(\varphi)$. For the right-hand side we have that $d_{[\langle\rangle]}([G, \chi]\varphi) = d_{[\langle\rangle]}(\chi) + d_{[\langle\rangle]}(\varphi)$.

   Since $[\langle\rangle]$-depths are the same, we calculate $[,]$-depths. For the left-hand side we have that $d_{[,]}(\chi \wedge [\psi_G \wedge \chi]\varphi) = d_{[,]}(\chi) + d_{[,]}(\varphi)$. In particular, $d_{[,]}(\chi \wedge [\psi_G \wedge \chi]\varphi) = \max\{d_{[,]}(\chi), d_{[,]}([\psi_G \wedge \chi]\varphi)\} = d_{[,]}([\psi_G \wedge \chi]\varphi) = d_{[,]}(\psi_G \wedge \chi) + d_{[,]}(\varphi) = \max\{d_{[,]}(\psi_G), d_{[,]}(\chi)\} + d_{[,]}(\varphi) = d_{[,]}(\chi) + d_{[,]}(\varphi)$. Depth of the right-hand side formula is $d_{[,]}([G, \chi]\varphi) = 1 + d_{[,]}(\varphi) + d_{[,]}(\chi)$. Hence, $\chi \wedge [\psi_G \wedge \chi]\varphi <^{Size}_{[,],[\langle\rangle]} [G, \chi]\varphi$.

2. On the left-hand side we have $d_{[\langle\rangle]}([\tau](\chi \wedge [\psi_G \wedge \chi]\varphi)) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}(\chi \wedge [\psi_G \wedge \chi]\varphi) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}(\chi) + d_{[\langle\rangle]}(\varphi)$. We have the same $[\langle\rangle]$-depth of the right-hand side: $d_{[\langle\rangle]}([\tau][G, \chi]\varphi) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}([G, \chi]\varphi) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}(\chi) + d_{[\langle\rangle]}(\varphi)$. However, $[,]$-depth is different: $d_{[,]}(\tau) + d_{[,]}(\chi) + d_{[,]}(\varphi)$ and $d_{[,]}(\tau) + d_{[,]}(\chi) + d_{[,]}(\varphi) + 1$ correspondingly (see the previous case). Hence, $[\tau](\chi \wedge [\psi_G \wedge \chi]\varphi) <^{Size}_{[,],[\langle\rangle]} [\tau][G, \chi]\varphi$.

3. On the left-hand side we have that $d_{[\langle\rangle]}(\langle A \setminus G, \psi_G\rangle\varphi) = d_{[\langle\rangle]}(\varphi)$, and on the right-hand side the depth is $d_{[\langle\rangle]}[\langle G\rangle]\varphi = d_{[\langle\rangle]}(\varphi) + 1$. Hence, $\langle A \setminus G, \psi_G\rangle\varphi <^{Size}_{[,],[\langle\rangle]} [\langle G\rangle]\varphi$.

4. Again, according to the definition of $[\langle\rangle]$-depth, $d_{[\langle\rangle]}([\tau]\langle A \setminus G, \psi_G\rangle\varphi) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}(\langle A \setminus G, \psi_G\rangle\varphi) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}(\varphi)$, whereas $d_{[\langle\rangle]}([\tau][\langle G\rangle]\varphi) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}([\langle G\rangle]\varphi) = d_{[\langle\rangle]}(\tau) + d_{[\langle\rangle]}(\varphi) + 1$. Thus, $[\tau]\langle A \setminus G, \psi_G\rangle\varphi <^{Size}_{[,],[\langle\rangle]} [\tau][\langle G\rangle]\varphi$.

□

Definition 6.8 (Canonical Model)**.** The *canonical model* is the model $M^C = (W^C, \sim^C, V^C)$, where

- W^C is the set of all maximal consistent theories,
- $\sim^C$ is defined as $x \sim^C_a y$ iff $K_a x = K_a y$,
- $x \in V^C(p)$ iff $p \in x$.

Relation $\sim^C$ is equivalence due to axioms $A2$, $A3$, and $A4$. And by Lindenbaum Lemma W^C is non-empty.

Lemma 6.10 (Truth)**.** For all maximal consistent theories x and all $\varphi \in \mathcal{L}_{CoRGAL}$, we have that $\varphi \in x$ if and only if $(M^C, x) \models \varphi$.

Proof. The proof is by induction on $<^{Size}_{[,],[\langle\rangle]}$-complexity of formulas. The base case, and cases for boolean, knowledge, and public announcement (with p, $\neg\psi$, $\psi \wedge \chi$, $K_a\psi$, $[\psi]\chi$ within its scope) formulas were proved in [Balbiani and van Ditmarsch, 2015]. We prove here only remaining instances involving relativised group and coalition announcements.

Induction Hypothesis. For all x and $\psi <^{Size}_{[,],[\langle\rangle]} \varphi_0$, we have that $\psi \in x$ if and only if $(M^C, x) \models \psi$.

Case $\varphi_0 = [G, \chi]\varphi$. ($\Rightarrow$) Suppose that $[G, \chi]\varphi \in x$. Since x is a theory, and thus contains all theorems and closed under $R0$, by axiom $A10$ we have that $\forall \psi_G$: $\chi \wedge [\psi_G \wedge \chi]\varphi \in x$. By Proposition 6.9 it holds that $\chi \wedge [\psi_G \wedge \chi]\varphi <^{Size}_{[,],[\langle\rangle]} [G, \chi]\varphi$. Using Induction Hypothesis we have $(M^C, x) \models \chi \wedge [\psi_G \wedge \chi]\varphi$ for all ψ_G. The latter is equivalent to $(M^C, x) \models [G, \chi]\varphi$ by the semantics.

($\Leftarrow$) Let $(M^C, x) \models [G, \chi]\varphi$. By the semantics of CoRGAL this is equivalent to $\forall \psi_G$: $(M^C, x) \models \chi \wedge [\psi_G \wedge \chi]\varphi$. By Proposition 6.9 it holds that $\chi \wedge [\psi_G \wedge \chi]\varphi <^{Size}_{[,],[\langle\rangle]} [G, \chi]\varphi$. Using Induction Hypothesis we have that $\forall \psi_G$: $\chi \wedge [\psi_G \wedge \chi]\varphi \in x$. Since x is a maximal consistent theory and closed under $R5$, it holds that $[G, \chi]\varphi \in x$.

Case $\varphi_0 = [\tau][G, \chi]\varphi$. ($\Rightarrow$) Assume that $[\tau][G, \chi]\varphi \in x$. Since x is a maximal consistent theory and hence contains all theorems, $[\tau]([G, \chi]\varphi \to \chi \wedge [\psi \wedge \chi]\varphi) \in x$ for all ψ_G. Using the distributivity of public announcements and the fact that x is closed under $R0$, we conclude that $[\tau](\chi \wedge [\psi_G \wedge \chi]\varphi) \in x$. Next, by Proposition 6.9 it holds that $[\tau](\chi \wedge [\psi_G \wedge \chi]\varphi) <^{Size}_{[,],[\langle\rangle]} [\tau][G, \chi]\varphi$, and thus using Induction Hypothesis we have that $(M^C, x) \models [\tau](\chi \wedge [\psi_G \wedge \chi]\varphi)$ for all ψ_G. The latter amounts to the fact that $(M^C, x) \models \tau$ implies $(M^C, x)^\tau \models \chi \wedge [\psi_G \wedge \chi]\varphi$ for all ψ_G. By the semantics of CoRGAL, we have that $(M^C, x) \models \tau$ implies $(M^C, x)^\tau \models [G, \chi]\varphi$, which is equivalent to $(M^C, x) \models [\tau][G, \chi]\varphi$.

($\Leftarrow$) Let $(M^C, x) \models [\tau][G, \chi]\varphi$. By the semantics of CoRGAL, this means that for any ψ_G: $(M^C, x) \models [\tau](\chi \wedge [\psi_G \wedge \chi])\varphi$. By Proposition 6.9 we have that $[\tau](\chi \wedge [\psi_G \wedge \chi]\varphi) <^{Size}_{[,],[\langle\rangle]} [\tau][G, \chi]\varphi$ for all ψ_G. Using Induction Hypothesis we can conclude that $\forall \psi_G$: $[\tau](\chi \wedge [\psi_G \wedge \chi]\varphi) \in x$. Since $[\tau]\sharp$ is a necessity form, and due to the fact that x is closed under $R5$, we have that $[\tau][G, \chi]\varphi \in x$.

Case $\varphi_0 = [\langle G \rangle]\varphi$. ($\Rightarrow$) Suppose that $[\langle G \rangle]\varphi \in x$. Since x is a theory, and thus contains all theorems and closed under $R0$, by axiom $A11$ we have that $\forall \psi_G$: $\langle A \setminus G, \psi_G \rangle \varphi \in x$. By Proposition 6.9 it holds that $\langle A \setminus G, \psi_G \rangle \varphi <^{Size}_{[,],[\langle\rangle]} [\langle G \rangle]\varphi$. Using Induction Hypothesis we have $(M^C, x) \models \langle A \setminus G, \psi_G \rangle \varphi$ for all ψ_G. The latter is equivalent to $(M^C, x) \models [\langle G \rangle]\varphi$ by the semantics.

($\Leftarrow$) Let $(M^C, x) \models [\langle G \rangle]\varphi$. By the semantics of CoRGAL this is equivalent to $\forall \psi_G$: $(M^C, x) \models \langle A \setminus G, \psi_G \rangle \varphi$. By Proposition 6.9 it holds that $\langle A \setminus G, \psi_G \rangle \varphi <^{Size}_{[,],[\langle\rangle]} [\langle G \rangle]\varphi$. Using Induction Hypothesis we have that $\forall \psi_G$: $\langle A \setminus G, \psi_G \rangle \varphi \in x$. Since x is a maximal consistent theory and closed under $R6$, it holds that $[\langle G \rangle]\varphi \in x$.

Case $\varphi_0 = [\tau][\langle G \rangle]\varphi$. ($\Rightarrow$) Assume that $[\tau][\langle G \rangle]\varphi \in x$. Since x is a maximal consistent theory and hence contains all theorems, $[\tau]([\langle G \rangle]\varphi \to \langle A \setminus G, \psi_G \rangle \varphi) \in x$ for all ψ_G. Using the distributivity of public announcements and the fact that x is closed under $R0$, we conclude that $[\tau]\langle A \setminus G, \psi_G \rangle \varphi \in x$. Next, by Proposition 6.9 it holds that $[\tau]\langle A \setminus G, \psi_G \rangle \varphi <^{Size}_{[,],[\langle\rangle]} [\tau][\langle G \rangle]\varphi$, and thus using Induction Hypothesis we have that $(M^C, x) \models [\tau]\langle A \setminus G, \psi_G \rangle \varphi$ for all ψ_G. The latter amounts to the fact that $(M^C, x) \models \tau$ implies $(M^C, x)^\tau \models \langle A \setminus G, \psi_G \rangle \varphi$ for all ψ_G. By the semantics of CoRGAL, we have that $(M^C, x) \models \tau$ implies $(M^C, x)^\tau \models [\langle G \rangle]\varphi$, which is equivalent to $(M^C, x) \models [\tau][\![\langle G\rangle\!]]\varphi$.

($\Leftarrow$) Let $(M^C, x) \models [\tau][\![\langle G\rangle\!]]\varphi$. By the semantics of CoRGAL, this means that for any ψ_G: $(M^C, x) \models [\tau]\langle A \setminus G, \psi_G\rangle\varphi$. By Proposition 6.9 we have that $[\tau]\langle A \setminus G, \psi_G\rangle\varphi <^{Size}_{[,],[\![\langle\rangle\!]]} [\tau][\![\langle G\rangle\!]]\varphi$ for all ψ_G. Using Induction Hypothesis we can conclude that $\forall\psi_G$: $[\tau]\langle A \setminus G, \psi_G\rangle\varphi \in x$. Since $[\tau]\sharp$ is a necessity form, and due to the fact that x is closed under $R6$, we have that $[\tau][\![\langle G\rangle\!]]\varphi \in x$. □

Finally, we prove the completeness of CoRGAL.

Proposition 6.11. For all $\varphi \in \mathcal{L}_{CoRGAL}$, if φ is valid, then $\varphi \in$ CoRGAL.

Proof. Towards a contradiction, suppose that φ is valid and $\varphi \notin$ CoRGAL. Since CoRGAL is a consistent theory, and by Propositions 6.6 and 6.7, we have that CoRGAL + $\neg\varphi$ is a consistent theory. Then, by Proposition 6.8, there exists a maximal consistent theory $x \supseteq$ CoRGAL + $\neg\varphi$ such that $\neg\varphi \in x$. By Proposition 6.10, this means that $(M^C, x) \not\models \varphi$, which contradicts φ being a validity. □

Chapter 7

Groups Versus Coalitions

Relative expressivity of logics of quantified public announcements is a long-standing open problem. We provide a small contribution towards solving this problem by showing that CAL and APAL are not at least as expressive as GAL. The proof is based on presenting a property and a set of models such that some GAL formula holds only in the models with the property, and for all formulas of CAL and APAL there are models with and without the property where they are true. In Section 7.2 we introduce a mechanism for evaluating GAL and CAL formulas in epistemic models. After that, we present the models and property expressed in GAL (Sections 7.3 and 7.4.1), and show that none of CAL formulas can capture this property (Section 7.4.2). As a nice 'side-effect' we obtain the same result for APAL. We start the chapter, however, with the proof that CAL is not at least as expressive as APAL (Section 7.1).

7.1 APAL $\not\preceq$ CAL

It was proved in [Ågotnes et al., 2010] that APAL $\not\preceq$ GAL. The intuition behind the proof is that agents A can only force A-definable restrictions of a given model, whereas APAL operators may force any restriction of the model up to bisimulation. The same reasoning can be applied to both groups and coalitions. Therefore, we use the proof from [Ågotnes et al., 2010] and modify it to show that APAL $\not\preceq$ CAL.

Theorem 7.1. APAL $\not\preceq$ CAL.

Proof. Consider an APAL formula $\Diamond(K_a p \land \neg K_b K_a p)$. Let us assume that there is an equivalent CAL formula Ψ. Without loss of generality we also assume that propositional atom q does not occur in Ψ, i.e. $q \notin \Theta_\Psi$.

Consider models in Figure 7.1.

These models correspond to various restrictions of M that agents can enforce. We need to show that the APAL formula distinguishes M and M_1^a, and no CAL

This chapter is the result of joint work with Natasha Alechina, Hans van Ditmarsch, and Tim French. In particular, use of formula games (Definition 7.2) is originally Tim's idea, and everyone contributed equally to the proof of Theorem 7.6

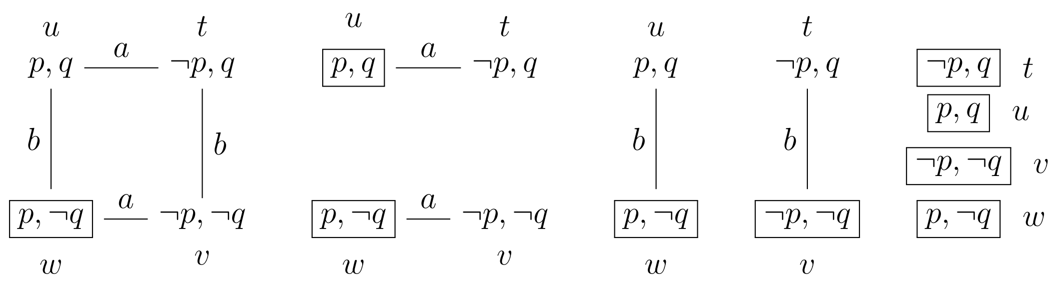

Figure 7.1: From left to right: models M, M_1^a (bottom), M_2^a (top), M_1^b, M_2^b and $M^{a,b}$ (four single-state models)

formula can distinguish them. Formally, it should be the case that $(M, w) \models \Psi$ and $(M_1^a, w) \models \Psi$, while $(M, w) \models \Diamond(K_a p \wedge \neg K_b K_a p)$ and $(M_1^a, w) \not\models \Diamond(K_a p \wedge \neg K_b K_a p)$.

That $(M, w) \models \Diamond(K_a p \wedge \neg K_b K_a p)$ and $(M_1^a, w) \not\models \Diamond(K_a p \wedge \neg K_b K_a p)$ is easy to check. In order to make $K_a p \wedge \neg K_b K_a p$ true, it is required to remove state v and retain states u and t. In model (M, w) announcement of $p \vee q$ would do the trick. And in model (M_1^a, w) either of possible updates, $\top$ or removing state v, does not satisfy the formula.

To prove that $(M, w) \models \Psi$ if and only if $(M_1^a, w) \models \Psi$ we need to show that for all subformulas ψ of Ψ, all states reachable from w, and all updates of M by agents' announcements, the equivalence holds. The extended induction hypothesis of the proof from [Ågotnes et al., 2010] conforms to all these requirements. We consider only cases of coalition announcements, and prove only the first row of the equivalences (proofs for other three rows are similar).

Let $\psi \in \mathcal{L}_{CAL}$, and $q \notin \Theta_\psi$.

Induction Hypothesis:

$$\begin{array}{r}
(M, w) \models \psi \text{ iff } (M_1^a, w) \models \psi \text{ iff } (M, u) \models \psi \text{ iff } (M_2^a, u) \models \psi \\
(M^{a,b}, w) \models \psi \text{ iff } (M_1^b, w) \models \psi \text{ iff } (M_1^b, u) \models \psi \text{ iff } (M^{a,b}, u) \models \psi \\
(M, v) \models \psi \text{ iff } (M_1^a, v) \models \psi \text{ iff } (M, t) \models \psi \text{ iff } (M_2^a, t) \models \psi \\
(M^{a,b}, v) \models \psi \text{ iff } (M_2^b, v) \models \psi \text{ iff } (M_2^b, t) \models \psi \text{ iff } (M^{a,b}, t) \models \psi
\end{array}$$

Case $[\langle\emptyset\rangle]\psi$. Let $(M, w) \models [\langle\emptyset\rangle]\psi$. This is equivalent to the fact that there is a joint $\{a, b\}$-strategy $X_a \cap X_b$ such that $(M, w)^{X_a \cap X_b} \models \psi$. Either agent has two strategies, and all possible intersections are presented in Figure 7.1 (models (M, w), (M_1^a, w), (M_1^b, w), and $(M^{a,b}, w)$). Therefore, at least one of the restrictions should satisfy ψ, i.e. $((M, w) \models \psi$ or $(M_1^a, w) \models \psi$ or $(M_1^b, w) \models \psi$ or $(M^{a,b}, w) \models \psi)$ $(*)$.

According to the induction hypothesis, $(*)$ is equivalent to $(M_1^a, w) \models \psi$ or $(M^{a,b}, w) \models \psi$. The latter means that there is an $\{a, b\}$-strategy in model M^a such that ψ holds in the resulting model, which is equivalent to $(M_1^a, w) \models [\langle\emptyset\rangle]\psi$.

Statement $(*)$ is also equivalent, by the induction hypothesis, to $((M, u) \models \psi$ or $(M_2^a, u) \models \psi$ or $(M_1^b, u) \models \psi$ or $(M^{a,b}, u) \models \psi)$. These are all possible intersections of agents' strategies in (M, u), which is equivalent to $(M, u) \models [\langle \emptyset \rangle] \psi$ by the semantics.

Finally, $(*)$ is equivalent to $((M_2^a, u) \models \psi$ or $(M^{a,b}, u) \models \psi)$, which holds if and only if there is an $\{a, b\}$-strategy in (M_2^a, u) such that ψ is true in the resulting updated model. This is equivalent to $(M_2^a, u) \models [\langle \emptyset \rangle] \psi$ by the semantics.

Case $[\langle \{a, b\} \rangle] \psi$. The same as above with the replacement of 'or' with 'and'.

Case $[\langle \{a\} \rangle] \psi$. Let $(M, w) \models [\langle \{a\} \rangle] \psi$. This is equivalent to the fact that for every a-strategy X_a there is a b-strategy X_b such that $(M, w)^{X_a \cap X_b} \models \psi$. We consider all these intersections: $((M, w) \models \psi$ or $(M_1^b, w) \models \psi)$ and $((M_1^a, w) \models \psi$ or $(M^{a,b}, w) \models \psi)(*)$.

By the induction hypothesis, $(*)$ is equivalent to $((M_1^a, w) \models \psi$ or $(M^{a,b}, w) \models \psi)$, which means that for every a-strategy in model M_1^a there is a b-strategy such that ψ holds in the resulting model. The latter is equivalent to $(M_1^a, w) \models [\langle \{a\} \rangle] \psi$ by the semantics.

According to the induction hypothesis, $(*)$ is equivalent to $((M, u) \models \psi$ or $(M_1^b, u) \models \psi)$ and $((M_2^a, u) \models \psi$ or $(M^{a,b}, u) \models \psi)$. These are all possible combinations of a's strategies and b's responses in (M, u). Hence, it is equivalent to $(M, u) \models [\langle \{a\} \rangle] \psi$ by the semantics.

Finally, $(*)$ is equivalent to $((M_2^a, u) \models \psi$ or $(M^{a,b}, u) \models \psi)$, and the latter holds if and only if for all a-strategies in (M_2^a, u) (agent a has only the trivial strategy in the model), there is a b-strategy such that ψ is true in the resulting updated model, which is equivalent to $(M_2^a, u) \models [\langle \{a\} \rangle] \psi$ by the semantics.

Case $[\langle \{b\} \rangle] \psi$. Similar reasoning as above. □

The rest of the chapter is dedicated to proving GAL $\not\preceq$ CAL.

7.2 Formula Games

The standard approach to comparing expressivity of modal languages is by using formula games [van Ditmarsch et al., 2008, Chapter 8]. In this section we present formula games for CAL and GAL. In order to deal with coalition announcement modalities, we use relativised group announcements (see Section 6.1) that allow us to consider an announcement by a coalition and a counter-announcement by an anti-coalition as separate moves in the game.

Definition 7.1 (NNF)**.** *Negation Normal Form* (NNF) is defined by the following BNF:

$$\varphi ::= \begin{array}{ll} & \top \mid \varphi \wedge \varphi \mid K_a \varphi \mid [G]\varphi \mid [G, \varphi]\varphi \mid [\langle G \rangle]\varphi \\ \mid & \bot \mid p \mid \neg p \mid \varphi \vee \varphi \mid \widehat{K}_a \varphi \mid [\varphi]\varphi \mid \langle G \rangle \varphi \mid \langle G, \varphi \rangle \varphi \mid \langle\!\langle G \rangle\!\rangle \varphi \end{array}.$$

If for some formula φ in NNF the outermost operator is from the top line, then we say that φ is in *Universal Negation Normal Form* (UNNF); and if the outermost operator is from the line below, then φ is in *Existential Negation Normal Form* (ENNF).

Proposition 7.2. Every formula of GAL and CAL is equivalent to a formula in NNF.

Proof. The proof is a straightforward 'pushing' of negations inside of the scope of operators. We use translation function $t : (\mathcal{L}_{GAL} \cup \mathcal{L}_{CAL}) \rightarrow \mathcal{L}_{NNF}$ that is defined as follows:

$$
\begin{array}{lcllcl}
t(\neg p) & = & \neg p & t(p) & = & p \\
t(\neg(\varphi \wedge \psi)) & = & t(\neg\varphi) \vee t(\neg\psi) & t(\varphi \wedge \psi) & = & t(\varphi) \wedge t(\psi) \\
t(\neg K_a \varphi) & = & \widehat{K}_a t(\neg\varphi) & t(K_a \varphi) & = & K_a t(\neg\varphi) \\
t(\neg[\psi]\varphi) & = & t(\psi) \wedge t([\psi]\neg\varphi) & t([\psi]\varphi) & = & [t(\psi)]t(\varphi) \\
t(\neg[G]\varphi) & = & \langle G\rangle t(\neg\varphi) & t([G]\varphi) & = & [G]t(\varphi) \\
t(\neg[\langle G\rangle]\varphi) & = & \langle[G]\rangle t(\neg\varphi) & t([\langle G\rangle]\varphi) & = & [\langle G\rangle]t(\varphi)
\end{array}
$$

□

Note that $\top$, $\bot$, $[G, \psi]\varphi$, and $\langle G, \psi\rangle\varphi$ will not appear in the image of the translation. These formulas, however, play the role of final and intermediate steps in games.

Now we are ready to define formula games.

Definition 7.2 (Formula Games)**.** Let some pointed model (M, w) and φ in NNF be given, and suppose that $\mathcal{M}$ is the set of pointed submodels $(N, w)^X$ of the given model M, where $X \subseteq W$ and $w \in X$. A *formula game for* φ *over* (M, w) is a tuple $\mathcal{G}^{\varphi}_{(M,w)} = (V_\forall, V_\exists, E, s)$, where

- $V_\forall = \{\ulcorner(N, v), \psi\urcorner \mid (N, v) \in \mathcal{M}, \psi \in \mathcal{L}_{UNNF}\} \cup \{\ulcorner(N, v), X, \chi, \psi\urcorner \mid (N, v) \in \mathcal{M}, X \subseteq W, \chi, \psi \in \mathcal{L}_{NNF}\}$ is the set of vertices of the $\forall$-player,
- $V_\exists = \{\ulcorner(N, v), \psi\urcorner \mid (N, v) \in \mathcal{M}, \psi \in \mathcal{L}_{ENNF}\}$ is the set of vertices of the $\exists$-player,
- $E \subset (V_\forall \cup V_\exists) \times (V_\forall \cup V_\exists)$ is the set of edges, where

$$E = \bigcup \left(\begin{array}{l}
\{(\ulcorner(N,v),p\urcorner,\ulcorner(N,v),\top\urcorner),(\ulcorner(N,v),\neg q\urcorner,\ulcorner(N,v),\top\urcorner) \\
\qquad\qquad\qquad\qquad \mid v \in V(p) \text{ and } v \notin V(q)\} \\
\{(\ulcorner(N,v),p\urcorner,\ulcorner(N,v),\bot\urcorner),(\ulcorner(N,v),\neg q\urcorner,\ulcorner(N,v),\bot\urcorner) \\
\qquad\qquad\qquad\qquad \mid v \notin V(p) \text{ and } v \in V(q)\} \\
\{(\ulcorner(N,v),\psi\wedge\chi\urcorner,\ulcorner(N,v),\psi\urcorner),(\ulcorner(N,v),\psi\wedge\chi\urcorner,\ulcorner(N,v),\chi\urcorner\} \\
\{(\ulcorner(N,v),\psi\vee\chi\urcorner,\ulcorner(N,v),\psi\urcorner),(\ulcorner(N,v),\psi\vee\chi\urcorner,\ulcorner(N,v),\chi\urcorner\} \\
\{(\ulcorner(N,v),K_a\psi\urcorner,\ulcorner(N,u),\psi\urcorner) \mid v \sim_a u\} \\
\{(\ulcorner(N,v),\widehat{K}_a\psi\urcorner,\ulcorner(N,u),\psi\urcorner) \mid v \sim_a u\} \\
\{(\ulcorner(N,v),[\chi]\psi\urcorner,\ulcorner(N,v),X,\chi,\psi\urcorner)\} \\
\{(\ulcorner(N,v),X,\chi,\psi\urcorner,\ulcorner(N,u),\chi\urcorner) \mid u \in X\} \\
\{(\ulcorner(N,v),X,\chi,\psi\urcorner,\ulcorner(N,u),t(\neg\chi)\urcorner) \mid u \in W \setminus X\} \\
\{(\ulcorner(N,v),X,\chi,\psi\urcorner,\ulcorner(N,v)^X,\psi\urcorner)\} \\
\{(\ulcorner(N,v),[G]\psi\urcorner,\ulcorner(N,v),[\psi_G]\psi\urcorner)\} \\
\{(\ulcorner(N,v),\langle G\rangle\psi\urcorner,\ulcorner(N,v),[\psi_G]\psi\urcorner)\} \\
\{(\ulcorner(N,v),[\langle G\rangle]\psi\urcorner,\ulcorner(N,v),\langle A\setminus G,\psi_G\rangle\psi\urcorner)\} \\
\{(\ulcorner(N,v),\langle[G]\rangle\psi\urcorner,\ulcorner(N,v),[A\setminus G,\psi_G]\psi\urcorner)\} \\
\{(\ulcorner(N,v),[G,\chi]\psi\urcorner,\ulcorner(N,v),[\psi_G\wedge\chi]\psi\urcorner)\} \\
\{(\ulcorner(N,v),\langle G,\chi\rangle\psi\urcorner,\ulcorner(N,v),[\psi_G\wedge\chi]\psi\urcorner)\}
\end{array} \right),$$

- s is the initial vertex $\ulcorner(M,w),\varphi\urcorner$.

The game is played between the $\forall$-player and the $\exists$-player, and a *play* consists of a sequence of vertices $s, s_1, \ldots, s_n$. The play is built by the players such that for some edge $(s_m, s_{m+1}) \in E$ if $s_m \in V_\forall$, then the universal player chooses s_{m+1}, and if $s_m \in V_\exists$, then the existential player chooses s_{m+1}. If either player is unable to move, i.e. they are in a $\top$-vertex or $\bot$-vertex, then they lose the game.

Consider model (M,w) in Figure 7.2 as an example. The set of pointed submodels of M is $\{(M,w)^{\{w\}},(M,w)^{\{w,v\}},(M,v)^{\{v\}},(M,v)^{\{w,v\}}\}$, and agent b's relation is identity. The formula game for $[p]\widehat{K}_a\neg p$ is presented in Figure 7.3, and the formula game for $\langle[\{b\}]\rangle K_a p$ in partially shown in Figure 7.4.

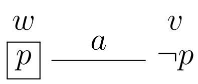

Figure 7.2: Model (M,w)

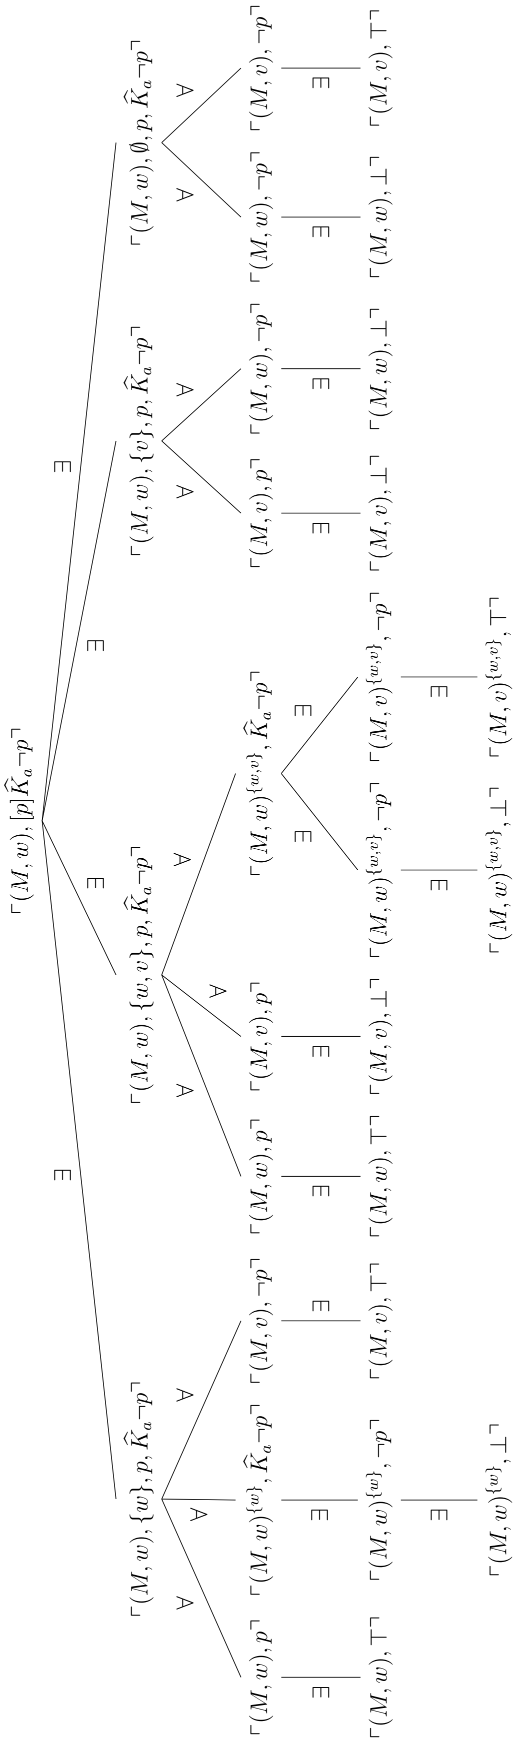

Figure 7.3: Formula game for $[p]\widehat{K}_a\neg p$ over (M, w)

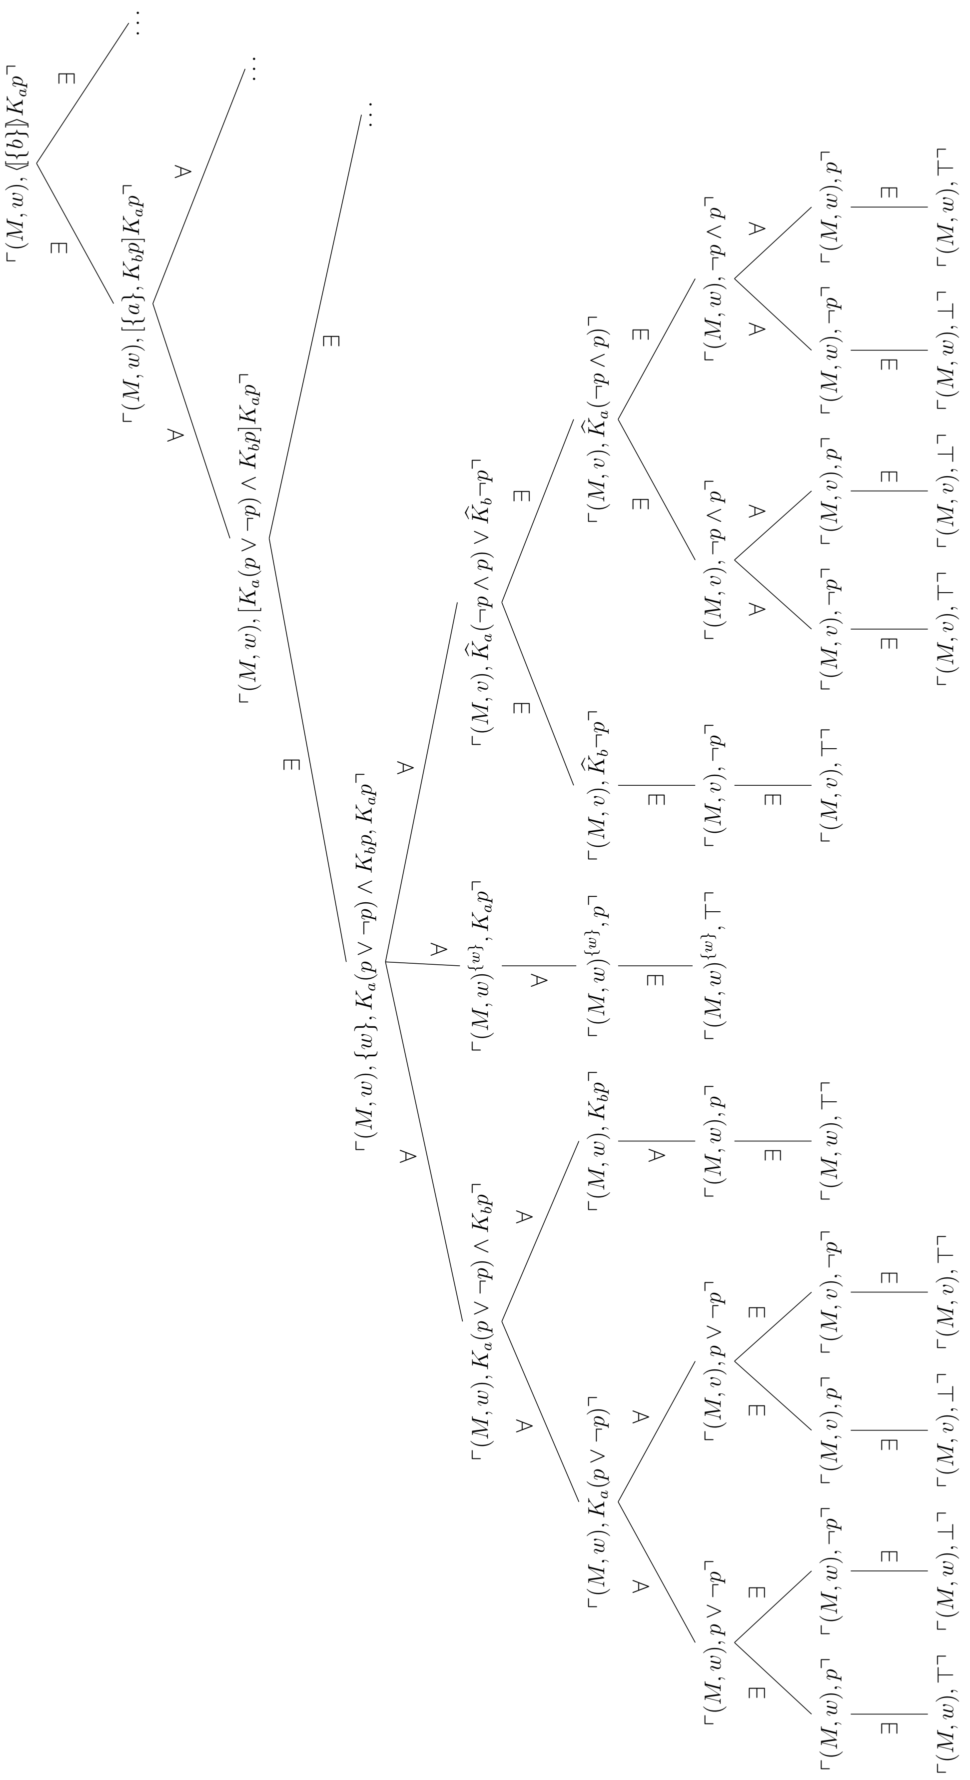

Figure 7.4: Formula game for $\langle\!\langle\{b\}\rangle\!\rangle K_a p$ over (M, w)

Let us ensure that there are no loops in games, i.e. plays of games are finite.

Proposition 7.3. Given formula φ in NNF, pointed model (M, w), and a game $\mathcal{G}^{\varphi}_{(M,w)}$, every play of the game is finite.

Proof. The proof is by induction on subformulas of φ.

Base Case: in the case of a propositional variable there is exactly one step in a play of the game.

Induction Hypothesis: plays of the game for subformulas ψ, $t(\neg\chi)$, and χ are finite on all pointed submodels (N, v) of M.

Propositional and epistemic cases are straightforward, so we omit them. Also note that it means that plays for epistemic formulas are finite.

Case $\ulcorner(N, v), [\chi]\psi\urcorner$: in this node of the game the existential player chooses a subset of the set of states of the given model. Such a choice leads to one of the vertices $\ulcorner(N, v), X, \chi, \psi\urcorner$. Every possible choice of the $\forall$-player from this vertex — $\ulcorner(N, u), \chi\urcorner$, $\ulcorner(N, u), t(\neg\chi)\urcorner$, or $\ulcorner(N, v)^X, \psi\urcorner$ — leads to a vertex with a play for either χ, $t(\neg\chi)$ or ψ that is finite by the Induction Hypothesis. Hence, a play of the game in $\ulcorner(N, v), [\chi]\psi\urcorner$ is finite.

Case $\ulcorner(N, v), [G]\psi\urcorner$: there is just one step from this vertex to some $\ulcorner(N, v), [\psi_G]\psi\urcorner$, and using the Induction Hypothesis and the fact that agents can only announce epistemic formulas, we conclude that the play from this vertex is finite.

Cases $\ulcorner(N, v), \langle G\rangle\psi\urcorner$, $\ulcorner(N, v), [G, \chi]\psi\urcorner$ and $\ulcorner(N, v), \langle G, \chi\rangle\psi\urcorner$ are similar to the previous one.

Case $\ulcorner(N, v), [\langle G\rangle]\psi\urcorner$: from this vertex there is exactly one $\forall$-step to some $\ulcorner\langle A \setminus G, \psi_G\rangle\psi\urcorner$. Since ψ_G is a formula of epistemic logic, and $A \setminus G$ can only announce epistemic formulas, we have that a play from $\ulcorner(N, v), [\langle G\rangle]\psi\urcorner$ is finite.

Case $\ulcorner(N, v), \langle\!\langle G\rangle\!\rangle\psi\urcorner$ is the same as above. □

Next proposition ties together satisfiability and existence of a winning strategy.

Proposition 7.4. The $\exists$-player has a winning strategy in a game $\mathcal{G}^{\varphi}_{(M,w)}$ if and only if $(M, w) \models \varphi$.

Proof. From right to left. The proof is by induction on the complexity c of φ. As a complexity measure we can use the one from Definition 6.6 with negations only in front of propositional variables.

Base Case: Assume that $(M, w) \models p$. Then the corresponding formula game consists only of one $\exists$-step from $\ulcorner(M, w), p\urcorner$ to $\ulcorner(M, w), \top\urcorner$, and the latter is the winning vertex of the existential player. The same argument holds for $\neg p$.

Induction Hypothesis: Assume that for all pointed submodels (N, v) of M and all formulas $t(\psi)$ in NNF such that $c(t(\psi)) < c(\varphi)$, if $(N, v) \models t(\psi)$, then $\ulcorner(N, v), t(\psi)\urcorner$ is a winning position for the $\exists$-player.

Propositional and epistemic cases are straightforward.

Case $(M, w) \models [\psi]\chi$: by the semantics this means that $(M, w) \models \neg\psi$ or $(M, s)^\psi \models \chi$. If the former is the case, then consider $X = [\![\psi]\!]_M$ and $Y = W \setminus [\![\psi]\!]_M$, where X can be an empty set. We have that for all $v \in X$: $(M, v) \models \psi$

and for all $u \in Y$: $(M, u) \models t(\neg\psi)$. By the Induction Hypothesis this implies that $\ulcorner(M, v), \psi\urcorner$ and $\ulcorner(M, u), t(\neg\psi)\urcorner$ are winning positions for the existential player for all $v \in X$ and $u \in Y$. Hence, $\ulcorner(M, w), X, \psi, \chi\urcorner$ is also a winning position for the $\exists$-player that she can choose from $\ulcorner(M, w), [\psi]\chi\urcorner$.

If $(M, w)^\psi \models \chi$, then again we consider $X = [\![\psi]\!]_M$ similarly to the case of $(M, w) \models \neg\psi$. Since $w \in X$, we need to deal with an additional case $(M, w)^X \models \chi$. By the Induction Hypothesis this means that $\ulcorner(M, w)^X, \chi\urcorner$ is a winning position for the $\exists$-player. Hence, $\ulcorner(M, w), X, \psi, \chi\urcorner$ is also a winning position for the $\exists$-player that she can choose from $\ulcorner(M, w), [\psi]\chi\urcorner$.

Case $(M, w) \models \langle G\rangle\psi$: by the semantics $(M, w) \models \langle G\rangle\psi$ is equivalent to $\exists\psi_G$: $(M, w) \models \langle\psi_G\rangle\psi$. The latter implies $(M, w) \models [\psi_G]\psi$. By the Induction Hypothesis, that means that the $\exists$-player can always choose a step in the game that corresponds a winning position $\ulcorner(M, w), [\psi_G]\psi\urcorner$. Thus, $\ulcorner(M, w), \langle G\rangle\psi\urcorner$ is also a winning position for the existential player.

Cases $(M, w) \models [G]\chi$, $(M, w) \models [G, \psi]\chi$, $(M, w) \models \langle G, \psi\rangle\chi$: a similar argument as above.

Case $(M, w) \models \langle\!\langle G\rangle\!\rangle\psi$: by the semantics of relativised group announcements we have that $\exists\psi_G$: $(M, w) \models [A \setminus G, \psi_G]\psi$. By the Induction Hypothesis, node $\ulcorner(M, w), [A \setminus G, \psi_G]\psi\urcorner$ is a winning position for the $\exists$-player. Hence, she can choose that particular ψ_G in $\ulcorner(M, w), \langle\!\langle G\rangle\!\rangle\psi\urcorner$.

Case $(M, w) \models [\![G]\!]\psi$: similar to the previous one.

From left to right. A similar argument as in the opposite direction for the contraposition: if $(M, w) \not\models \varphi$, then the $\forall$-player has a winning strategy in a game $\mathcal{G}^\varphi_{(M,w)}$. $\square$

Let us return to the example in in Figure 7.2. For (M, w) we have that $(M, w) \not\models [p]\widehat{K}_a\neg p$ and $(M, w) \models \langle\!\langle\{b\}\rangle\!\rangle K_a p$. Indeed, in Figure 7.3 the $\exists$-player does not have a winning strategy, and in Figure 7.4 she does (noting that none of a's announcements modify the original model).

7.3 Chain Models

In order to compare relative expressivity of GAL and CAL it is not enough to consider just a pair of models. Suppose that for some $\varphi \in \mathcal{L}_{GAL}$ we have that $(M, w) \models \varphi$ and $(N, v) \not\models \varphi$. In this case, formula φ must have subformulas with group announcement operators, for otherwise φ would be a CAL formula as well. Hence, in φ or its negation there is an existential group announcement operator $\langle G\rangle\psi$. According to the semantics, this implies that we can substitute $\langle G\rangle\psi$ with some $\langle\psi_G\rangle\psi$. And φ with such a substitution for all group announcements operators is a PAL (and hence CAL) formula. In other words, given two finite models group announcement operators $[G]\varphi$ can be 'simulated' via a disjunction with all possible non-equivalent (in the sense of model updates) substitutions ψ_G for $[G]\varphi$.

The same argument can be carried out for any finite set of epistemic models. Therefore, in order to unleash the full potential of group announcement operators we consider two infinite sets of models. After that, we show that there is GAL formula that is true in one such set and false in the other, and no CAL formula can capture the difference between the sets.

Models we are dealing with are called chain models.

Definition 7.3 (Chain Models)**.** A *chain model*, or a *chain*, is an epistemic model $M = (S, \sim, V)$, where

- $S = \{l, l+1, \ldots, r-1, r\} \subset \mathbb{Z}$ is a finite set of consecutive integers,
- $x \sim_a y$ if and only if $y = x + 1$ and x is even,
- $x \sim_b y$ if and only if $y = x + 1$ and x is odd,
- $z - 1 \sim_c z \sim_c z + 1$ if and only if $z \bmod 3 = 1$,
- $V(p) = \{3k, 3k - 1 \in S \mid k \in \mathbb{Z}\}$.

We use pair (l, r) to refer to the corresponding chain model.

In graphical representation of chains we use a solid line for agent a's relation, a dashed line for b's relation, and c cannot distinguish states in the same dotted box. An example of a chain is presented in Figure 7.5.

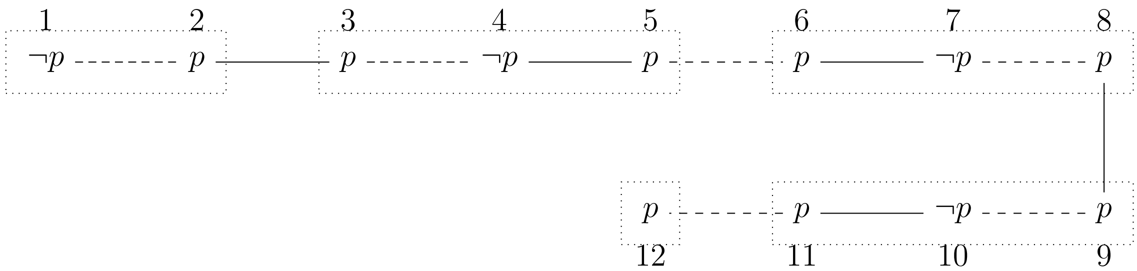

Figure 7.5: A $(1, 12)$-chain

Chain models are regular in their structure, and they may only differ from each other in leftmost and rightmost states. Hence, we can give a classification of chains based on their extremities.

Definition 7.4 (Classification of Chains)**.** Let some chain model (x, y) be given.

- If $x \bmod 6 = 1$ ($y \bmod 6 = 4$), then $(M, x) \models K_a\neg p$ ($(M, y) \models K_a\neg p$).
- If $x \bmod 6 = 2$ ($y \bmod 6 = 3$), then $(M, x) \models K_bK_ap$ ($(M, y) \models K_bK_ap$).
- If $x \bmod 6 = 3$ ($y \bmod 6 = 2$), then $(M, x) \models K_ap$ ($(M, y) \models K_ap$). Note that for such an x (y), chain (x, y) is bisimilar to model $(2x - (y + 1), y)$ ($(x, 2y - x + 1)$) via bisimulation $\{(x + k, x - k - 1) \mid 0 \leq k \leq y - x\}$ ($\{(y + k + 1, y - k) \mid 0 \leq k \leq y - x\}$). See Figure 7.6 for an example.

- If $x \bmod 6 = 4$ ($y \bmod 6 = 1$), then $(M, x) \models K_b\neg p$ ($(M, y) \models K_b\neg p$).
- If $x \bmod 6 = 5$ ($y \bmod 6 = 0$), then $(M, x) \models K_aK_bp$ ($(M, y) \models K_aK_bp$).
- If $x \bmod 6 = 0$ ($y \bmod 6 = 5$), then $(M, x) \models K_bp$ ($(M, y) \models K_bp$). Note that for such an x (y), chain (x, y) is bisimilar to model $(2x - (y + 1), y)$ ($(x, 2y - x + 1)$) via bisimulation $\{(x + k, x - k - 1) \mid 0 \leq k \leq y - x\}$ ($\{(y + k + 1, y - k) \mid 0 \leq k \leq y - x\}$). See Figure 7.6.

Therefore, we can describe the *type* of a chain (x, y) as the pair $[x \bmod 6, y \bmod 6]$.

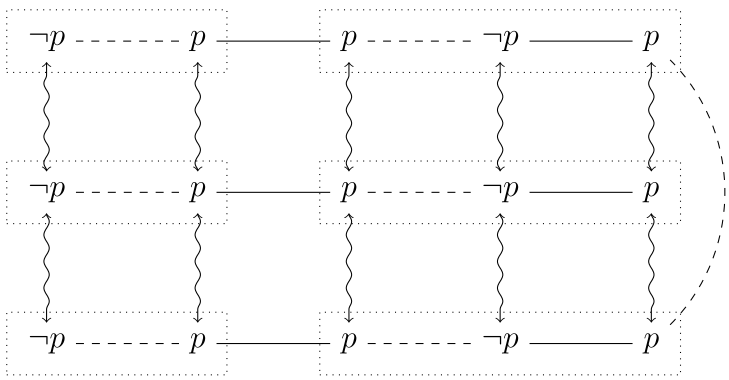

Figure 7.6: A bisimulation (wavy arrows) between chains $(1, 5)$ (in the middle) and $(1, 10)$ (starts at the top and wraps around on the right to the bottom)

In our proof we are primarily interested in models of types $[1, 2]$, $[0, 4]$, and $[0, 2]$, and their examples are depicted in Figures 7.7, 7.8 and 7.9. We also note that models with 'unbroken' c-relations are all bismilar to a $[0, 2]$-chain (see Figure 7.9).

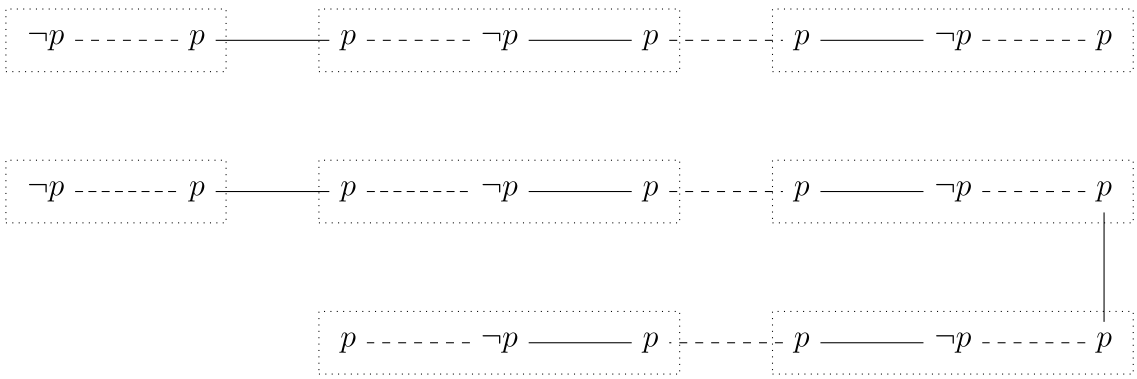

Figure 7.7: $[1, 2]$-models

Definition 7.5 (Terminal State)**.** Given a $[1, 2]$- or $[0, 4]$-chain M, state x of the model is called *terminal*, if $(M, x) \models \Omega$, where $\Omega := K_a\neg p$.

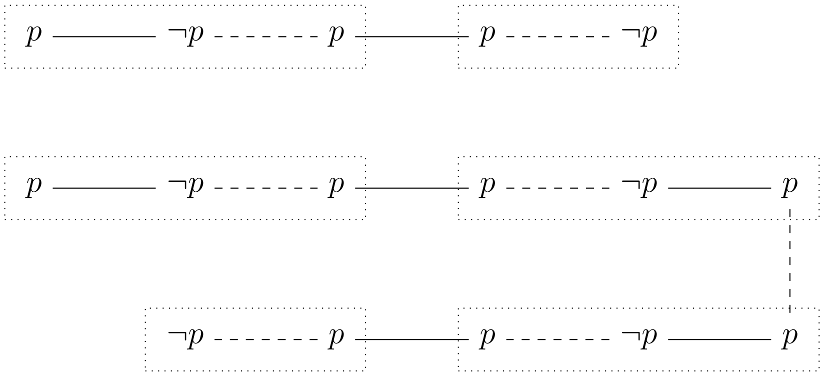

Figure 7.8: $[0, 4]$-models

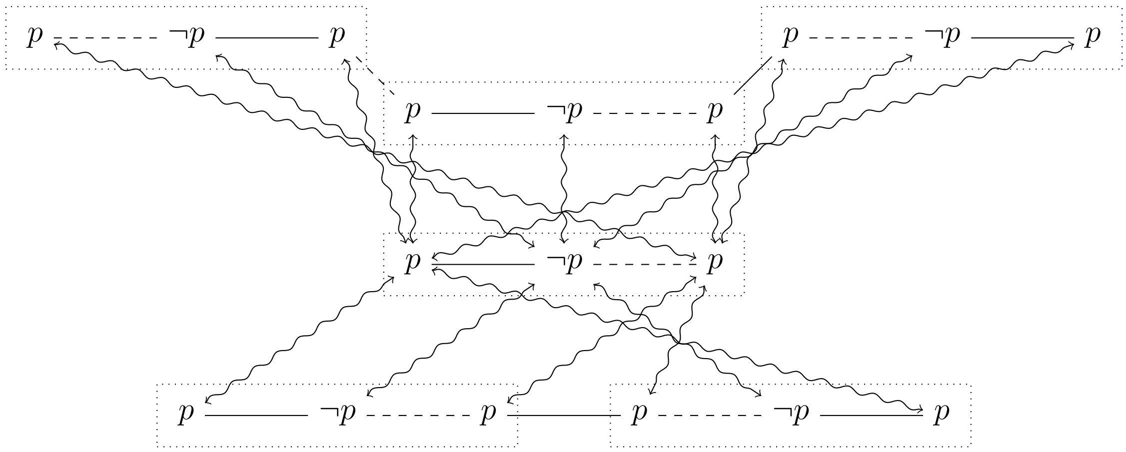

Figure 7.9: The only chain of type $[0, 2]$ up to bisimulation

In Figures 7.7 and 7.8 the terminal state is the leftmost and rightmost state respectively. Note that in such models there is only one terminal state.

We use terminal states to define a property expressible in GAL in the next section. Moreover, terminal states may be used to target other states in the model in order to 'cut' chains. For example, to specify a state that is exactly three steps from the Ω-state we can use the formula:

$$\Omega + 3 := \widehat{K}_b \widehat{K}_a \widehat{K}_b \Omega \wedge K_a K_b K_a \neg \Omega.$$

See Figure 7.10 for representation of the formula.

In the example, if agent b announces, for instance $K_b \neg(\Omega + 3)$, the updated model will be the one without the b-link with the $\Omega + 3$-state (squared). A similar announcement $K_a \neg(\Omega + 3)$ can be made by agent a, and group $\{a, b\}$ can cut any a- and b-links in models with terminal states.

Now let us consider non-terminal rightmost and leftmost states in $[1, 2]$- and $[1, 4]$-chains. They are presented in Figure 7.11. In Definition 7.4 we pointed out that no epistemic formula can distinguish these states from n-bisimilar ones in

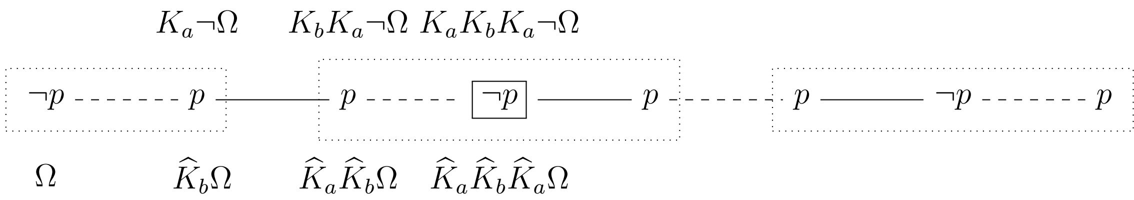

Figure 7.10: Removing states from a model using Ω

larger models. In other words, in order to specify such states, we should refer to the terminal one. Epistemic formulas, however, have a finite size, and hence formulas that refer to the terminal state are true only in chains of some depth, and we can always find a larger chain of the same type such that any given epistemic formula that was true in the smaller model will be false in the greater one.

Therefore, we use formulas of GAL to describe those non-terminal states, and we call these formulas Mid_a and Mid_b. The former is defined as

$$Mid_a := K_a p \land [A](\widehat{K}_b \neg p \to K_a \widehat{K}_b \neg p),$$

and it holds in the rightmost states of $[1, 2]$-models. The latter is defined as

$$Mid_b := K_b p \land [A](\widehat{K}_a \neg p \to K_b \widehat{K}_a \neg p),$$

and it holds in the leftmost states of $[0, 4]$-models.

Figure 7.11: Mid_a and Mid_b

7.4 GAL $\not\leq$ CAL

In this section we define a property of $[1, 2]$-chains expressible in GAL (Section 7.4.1), and show that it is impossible to capture that property in CAL (Section 7.4.2). Throughout the section we assume that for all considered pointed $[1, 2]$-models $((l, r), w)$, distances between l and w and r are sufficiently long.

7.4.1 What GAL Can Express...

We start this section with formulas that are valid on a certain class of chain models. First,

$$T(0, 2) := K_a K_b(\neg p \to [A]((\widehat{K}_a p \land \widehat{K}_b p) \to K_a K_b \neg(K_a p \land K_b p)))$$

describes $[0,2]$-models as there is no announcement from any agent that can make a and b know p without removing all $\neg p$ states.

Formulas for $[1,2]$- and $[0,4]$-models are as follows:

$$T(1,2) := \neg T(0,2) \wedge [K_b \neg \Omega] T(0,2) \wedge [\neg Mid_b \wedge K_b \neg \Omega] T(0,2),$$

$$T(0,4) := \neg T(0,2) \wedge [K_b \neg \Omega] T(0,2) \wedge [\neg Mid_a \wedge K_b \neg \Omega] T(0,2).$$

Intuitively, they mean that $[1,2]$- and $[0,4]$-chains are not bisimilar to $[0,2]$-chains (first conjunct), removing the link with the terminal state makes them bisimilar to a $[0,2]$-chain (second conjunct), and they differ between each other in extreme non-terminal states described by Mid_a and Mid_b (third conjunct). Note that group announcement operators appear only in $T(0,2)$, Mid_a and Mid_b, and none of these formulas mention agent c.

The actual property we are interested in applies to pointed models. Given a pointed model $((l,r),w)$ of type $[1,2]$, is the terminal node in the a direction from w ($((l,r),w)$ is an a-model), or the b direction ($((l,r),w)$ is a b-model)? See Figure 7.12 for a representation of this problem.

Figure 7.12: A $(1,9)$-model, where $w=7$, and $((1,9),w)$ is an a-model

We show that GAL can express whether a given pointed model is an a- or b-model.

The formula that expresses the property of (M,w) being a b-model is

$$b:\Omega = \bigwedge \left(\begin{array}{l} K_a p \to \langle \{c\} \rangle (Mid_a \wedge T(1,2)) \\ \neg p \to K_a (p \to \langle \{c\} \rangle (Mid_b \wedge T(0,4))) \\ K_b p \to [\{c\}](Mid_b \to \neg T(0,4)) \end{array} \right).$$

Formula $a:\Omega$ can be obtained by swapping subscripts a and b, and formulas $T(1,2)$ and $T(0,4)$ in $b:\Omega$:

$$a:\Omega = \bigwedge \left(\begin{array}{l} K_b p \to \langle \{c\} \rangle (Mid_b \wedge T(0,4)) \\ \neg p \to K_b (p \to \langle \{c\} \rangle (Mid_a \wedge T(1,2))) \\ K_a p \to [\{c\}](Mid_a \to \neg T(1,2)) \end{array} \right).$$

We sketch a proof of correctness of formula $a:\Omega$.

Theorem 7.5. Let sets $\mathcal{M}_A$ and $\mathcal{M}_B$ of all a and b pointed $[1,2]$-chains be given. Then $(M,w) \models a:\Omega$ for all $(M,w) \in \mathcal{M}_A$, and $(N,v) \not\models a:\Omega$ for all $(N,v) \in \mathcal{M}_B$.

Proof. The reader is encouraged to use figures from the previous section for reference. Let $(M, w) \models a : \Omega$ for some $[1, 2]$-chain (M, w). Since no conjunction of any two formulas $K_b p$, $\neg p$, or $K_a p$ can be true in a pointed chain, we have that either $(M, w) \models K_b p$, or $(M, w) \models \neg p$, or $(M, w) \models K_a p$.

Case $K_b p$. Let $(M, w) \models K_b p$. By the construction of chain models, this means that b cannot distinguish two p-states in two adjacent c-equivalence classes, and a considers $\neg p$ possible in the current c-equivalence class. Hence, c can cut b's relation making the current state a Mid_b state. Note that the terminal state remains intact, and thus we have that $T(0, 4)$ holds in the updated model. This means that $(M, w) \models \langle\{c\}\rangle(Mid_b \wedge T(0, 4))$.

Assume that $(N, v) \models K_b p$. As (N, v) is a $b : \Omega$-model, every cut by c either cuts the b-relation, and hence cuts the path to the terminal state, or does not satisfy Mid_b (c cannot make the current state to be extreme). Therefore, $(N, v) \models [\{c\}](\neg Mid_b \vee \neg T(0, 4))$.

Case $\neg p$. Let $(M, w) \models \neg p$ and $(M, t) \models p$ for some t such that $w \sim_b t$. By the construction of chain models, this means that a cannot distinguish two p-states in two adjacent c-equivalence classes, and b considers $\neg p$ possible in the current c-equivalence class. Hence, c can cut a's relation making the current state t a Mid_a state. Note that the terminal state remains intact, and thus we have that $T(1, 2)$ holds in the updated model. This means that $(M, t) \models p \wedge \langle\{c\}\rangle(Mid_a \wedge T(1, 2))$ for some $w \sim_b t$. We can make the latter formula less strict so that it holds in $\neg p$ states as well: $(M, t) \models p \rightarrow \langle\{c\}\rangle(Mid_a \wedge T(1, 2))$. By the construction of chains, there are only two states in b-relation with the current one: a p-state and a $\neg p$-state. Thus, $(M, w) \models K_b(p \rightarrow \langle\{c\}\rangle(Mid_a \wedge T(1, 2)))$, and we finally have that $(M, w) \models \neg p \rightarrow K_b(p \rightarrow \langle\{c\}\rangle(Mid_a \wedge T(1, 2)))$.

Assume that $(N, v) \models \neg p$ and $(N, s) \models p$ for some s such that $v \sim_b s$. As (N, v) is a b-model, (N, s) is an a- model. So, every cut by c either cuts the a-relation, and hence cuts the path to the terminal state, or does not satisfy Mid_a (c cannot make the current state to be extreme). Therefore, $(N, s) \models [\{c\}](\neg Mid_a \vee \neg T(1, 2))$.

Case $K_a p$. Let $(M, w) \models K_a p$. By the construction of chain models, this means that a cannot distinguish two p-states in two c-equivalence classes, and b considers $\neg p$ possible in the current c-equivalence class. Hence, if c cuts a-relation making Mid_a true, she also makes the terminal state inaccessible from the current one. On the other hand, if the terminal state is still accessible from the current state, then in this case the current state does not satisfy Mid_a. This means that $(M, w) \models [\{c\}](Mid_a \rightarrow T(1, 2))$.

Assume that $(N, v) \models K_a p$. As (N, v) is a b-model, c has a cut such that Mid_a and $T(1, 2)$ holds. Such a cut 'removes' all c-equivalence classes to the right of the current state, and makes the current state the rightmost state in the updated model. Therefore, $(N, v) \models \neg[\{c\}](\neg Mid_a \vee \neg T(1, 2))$. □

7.4.2 ...And CAL Can Not

In this section we show that no CAL formula can capture the property of a pointed model 'being an a-model.'

An intuition behind the proof is that CAL operators require all agents announce their knowledge formulas simultaneously. For our chain models, intersection of agents' relations is an identity, and hence if it is possible to force some configuration of an a-model, then agents together, whether in the same coalition, or divided, can replicate the same configuration in a b-model. Contrast this to formula $a : \Omega$ in the previous section. The only agent that makes any announcements is c, and her relation is not discerning enough to force isomorphic submodels of some a- and b-models. If c preserves the terminal state in one class of models, she cannot replicate this announcement in the other class such that the resulting updated models are isomorphic (c cannot cut her own equivalence class to make Ω true in the opposite direction).

Theorem 7.6. Let sets $\mathcal{M}_A$ and $\mathcal{M}_B$ of all a and b pointed $[1, 2]$-chains be given. Then for all $\Psi \in \mathcal{L}_{CAL}$, if $(M, w) \models \Psi$ for all $(M, w) \in \mathcal{M}_A$, then there exists some $(N, v) \in \mathcal{M}_B$ such that $(N, v) \models \Psi$.

Proof. Suppose, contrary to our claim, that for all $(M, w) \in \mathcal{M}_A$ there is a formula $\Psi \in \mathcal{L}_{CAL}$ such that $(M, w) \models \Psi$ and for all $(N, v) \in \mathcal{M}_B$ it holds that $(N, v) \not\models \Psi$. The proof proceeds by playing simultaneous formula games over all pointed chains.

We also assume that models in both sets are sufficiently large: for $|\Psi| = n$ we have that models in sets $\mathcal{M}_A$ and $\mathcal{M}_B$ are 2^n-bisimilar to each other. This is to ensure that no epistemic formula can distinguish any two models.

Let us partition the sets of games into $\mathcal{G}_A$ and $\mathcal{G}_B$ where player $\exists$ will have a winning strategy for games in $\mathcal{G}_A$, and player $\forall$ will have a winning strategy for games in $\mathcal{G}_B$.

For all moves in games except moves via coalition announcements, we proceed as follows. If it is an $\exists$-player move, we consider the games played over $\mathcal{M}_A$, and play the move for the $\exists$-player's winning strategy on all models in $\mathcal{M}_A$. We also play the corresponding move over $\mathcal{M}_B$: in the case of disjunction, we choose the same disjunct, and in the case of $\widehat{K}_a$-move in 2^k-bisimilar states, we consider moves equivalent if the chosen states are 2^{k-1}-bisimilar. If it is a $\forall$-player move, we play the move that agrees with the $\forall$ winning strategy in $\mathcal{G}_B$ games, and copy this move in the $\mathcal{G}_A$ games. Thus, we are playing two winning strategies against one another, and the game ends if either $\exists$-player or $\forall$-player cannot move. However, this cannot happen because all pointed models are 2^n-bisimilar.

Therefore, we need show that we can maintain the following *invariant*: after step i of the formula game, there are infinitely many models of the same type in $\mathcal{M}_A$ and $\mathcal{M}_B$ that are still 2^{n-i}-bisimilar. In the final step of the game, we end up with some propositional variable on which both classes of models agree. Hence, we have a contradiction since both players have a winning strategy by the assumption.

Cases of boolean and epistemic formulas are trivial.

Case $[\chi]\psi$. Assume that for some $(M,w) \in \mathcal{M}_A$, $\ulcorner(M,w),[\chi]\psi\urcorner$ is a winning position for the $\exists$-player. This means that there is a subset X such that $\ulcorner(M,w),X,\chi,\psi\urcorner$ is also a winning position. Let for some $(N,v) \in \mathcal{M}_B$ such that (M,w) and (N,v) are 2^{n-i}-bisimilar, $[\![\chi]\!]_N = Y$. We consider two cases.

First, if $(M,w)^\chi$ is 2^{n-i-1} bisimilar to $(N,v)^\chi$, then $\exists$ can play the corresponding move $\ulcorner(N,v),Y,\chi,\psi\urcorner$ in $\mathcal{G}_B$, and the invariant continues to hold.

In the second case, if $(M,w)^\chi$ and $(N,v)^\chi$ are not 2^{n-i-1}-bisimilar, there is some $s \in W^M$ and and some $t \in W^N$ such that (M,s) and (N,t) disagree on the interpretation of χ, and s and t are within the same 2^{n-i-1} steps from w and v respectively. Suppose $(M,s) \models \chi$ and $(N,t) \not\models \chi$. In this case $\exists$ must still play the corresponding move $\ulcorner(N,v),Y,\chi,\psi\urcorner$ in $\mathcal{G}_B$, as any alternative to Y would allow the $\forall$-player to have a winning strategy. The universal player, however, can respond with the moves $\ulcorner(M,s),\chi\urcorner$ in $\mathcal{G}_A$ and $\ulcorner(N,t),t(\neg\chi)\urcorner$ in $\mathcal{G}_B$. According to Proposition 7.4, the $\exists$-player has a winning strategy in $\ulcorner(N,t),t(\neg\chi)\urcorner$, and thus the $\forall$-player has a winning strategy in $\ulcorner(N,t),\chi\urcorner$. Since s and t are the same number of steps away from w and v, we have that (M,s) is an a-model if and only if (N,t) is a b-model. Moreover, since s and t are within 2^{n-i-1} steps and the original models (M,w) and (N,v) are 2^{n-i}-bisimilar, it follows that (M,s) and (N,t) are 2^{n-i-1}-bisimilar. Hence the invariant holds for those models and the proof proceeds in (M,s) and (N,t).

So, we must reach coalition operators. Note that at this point games may cease to be over $[1,2]$-models since prior public announcements may have cut chains in various ways. However, this does not affect the proof as we are interested in agents' announcements rather than in chain types. Moreover, for the coalition cases we do not have to keep the invariant since all these cases lead straight to a contradiction.

We will consider only existential coalition announcement operators $\langle\![G]\!\rangle\psi$, and the corresponding results for $[\![G]\!]\psi$ can be obtained by swapping A to B, and the $\exists$-player to the $\forall$-player.

Case $\langle\![\{a,b,c\}]\!\rangle\psi$. Let $(M,w) \models \langle\![\{a,b,c\}]\!\rangle\psi$. According to Definition 7.2 there is a relativised group announcement by a, b, and c such that $\ulcorner(M,w),[\emptyset,\psi_{\{a,b,c\}}]\psi\urcorner$ is a winning position for the $\exists$-player. For this node there is only one possible $\forall$-step: $\ulcorner(M,w),[\psi_{\{a,b,c\}}]\psi\urcorner$. Since $\mathcal{M}_B$ is infinite, there is a model (N,v) and an announcement $\chi_{\{a,b,c\}}$ by a, b, and c, such that $(M,w)^{\psi_{\{a,b,c\}}}$ is isomorphic to $(N,v)^{\chi_{\{a,b,c\}}}$ (see Figure 7.13 for an example).

Indeed, consider set $W^{\psi_{\{a,b,c\}}}$. We can enumerate states in the set from left to right. Next, let $N^0 = (W^0,\sim^0,V^0)$ be a model such that $W^0 = W^{\psi_{\{a,b,c\}}}$, $w_{n+1-i} \in V^0(p)$ if and only if $w_i \in V(p)$, and $w_{n+1-i} \sim_a^0 w_{n-i}$ if and only if $w_i \sim_a w_{i+1}$ for all $a \in A$. In other words, we flip model $(M,w)^{\psi_{\{a,b,c\}}}$ from left to right. Note that agents' relations are also flipped: if state w was an a-state, it would become a b-state. Moreover, we can always find a b-model (N,v) that has N^0 as a submodel. Since agents can together enforce any configuration of (N,v), they have a joint announcement $\chi_{\{a,b,c\}}$ such that $(N,v)^{\chi_{\{a,b,c\}}}$ is isomorphic to $(M,w)^{\psi_{\{a,b,c\}}}$, where $(N,v)^{\chi_{\{a,b,c\}}} = N^0$. The same argument can be made in other cases of the proof.

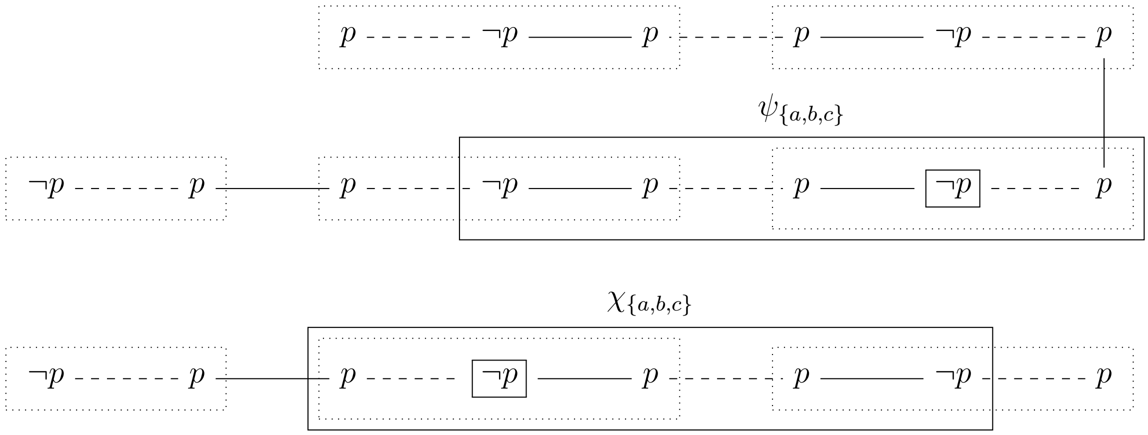

Figure 7.13: An a-model (above) and a b-model (below)

Thus $\ulcorner(N,v),[\chi_{\{a,b,c\}}]\psi\urcorner$ (and hence $\ulcorner(N,v),[\emptyset,\chi_{\{a,b,c\}}]\psi\urcorner$) is a winning position for the $\exists$-player, and she has a winning strategy for a model from $\mathcal{M}_B$. A contradiction. Note that since agents a, b, and c can together enforce any configuration of a model (up to bisimulation), the argument holds for the case of arbitrary public announcements.

Case $\langle\!\langle\{a,b\}\rangle\!\rangle\psi$. Let $(M,w)\models\langle\!\langle\{a,b\}\rangle\!\rangle\psi$. This means that $\ulcorner(M,w),\langle\!\langle\{a,b\}\rangle\!\rangle\psi\urcorner$ is a winning position for the $\exists$-player. Therefore, $\ulcorner(M,w),[\{c\},\psi_{\{a,b\}}]\psi$ is also a winning node for the player. This means that whichever announcement $\psi_{\{c\}}$ by agent c the $\forall$-player chooses, the $\exists$-player is still in a winning position $\ulcorner(M,w),[\psi_{\{a,b\}}\wedge\psi_{\{c\}}]\psi\urcorner$. There is a model $(N,v)\in\mathcal{M}_B$ such that for some announcement $\chi_{\{a,b\}}$ by agents a and b it holds that $(M,w)^{\psi_{\{a,b\}}}$ is isomorphic to $(N,v)^{\chi_{\{a,b\}}}$, and c has an isomorphic set of possible counter-announcements (see Figure 7.13 for an example). This is due to the fact that a and b can together force any configuration of a model. Hence $\ulcorner(N,v),[\{c\},\chi_{\{a,b\}}]\psi\urcorner$ is also a winning position for the existential player, and this leads to a contradiction.

Case $\langle\!\langle\{a,c\}\rangle\!\rangle\psi$. Let $(M,w)\models\langle\!\langle\{a,c\}\rangle\!\rangle\psi$. This means that $\ulcorner(M,w),\langle\!\langle\{a,c\}\rangle\!\rangle\psi\urcorner$ is a winning position for the $\exists$-player. Therefore, $\ulcorner(M,w),[\{b\},\psi_{\{a,c\}}]\psi$ is also a winning node for the player. This means that whichever announcement $\psi_{\{b\}}$ by agent b the $\forall$-player chooses, the $\exists$-player is still in a winning position $\ulcorner(M,w),[\psi_{\{a,c\}}\wedge\psi_{\{b\}}]\psi\urcorner$. Consider a model $(N,v)\in\mathcal{M}_B$. If there is some announcement $\chi_{\{a,c\}}$ by agents a and c such that $(M,w)^{\psi_{\{a,c\}}}$ and $(N,v)^{\chi_{\{a,c\}}}$ are isomorphic, then by the similar reasoning as in the previous case we have a contradiction. See Figure 7.14, where counter-announcements by b are depicted by dashed rectangles.

Note that $\{a,c\}$ sometimes cannot make such an announcement, because the coalition cannot cut a's relations that are within c-equivalence classes, and $(M,w)^{\psi_{\{a,c\}}}$ may contain some extreme state. In other words, this a's relations that a and c cannot cut, may have been cut by a previous public announcement (and hence the corresponding state is the rightmost or the leftmost one). Since our chosen a-model is large enough even after being trimmed by public announcements (i.e. because the invariant holds), there is an a-relation in $(M,w)^{\psi_{\{a,c\}}}$

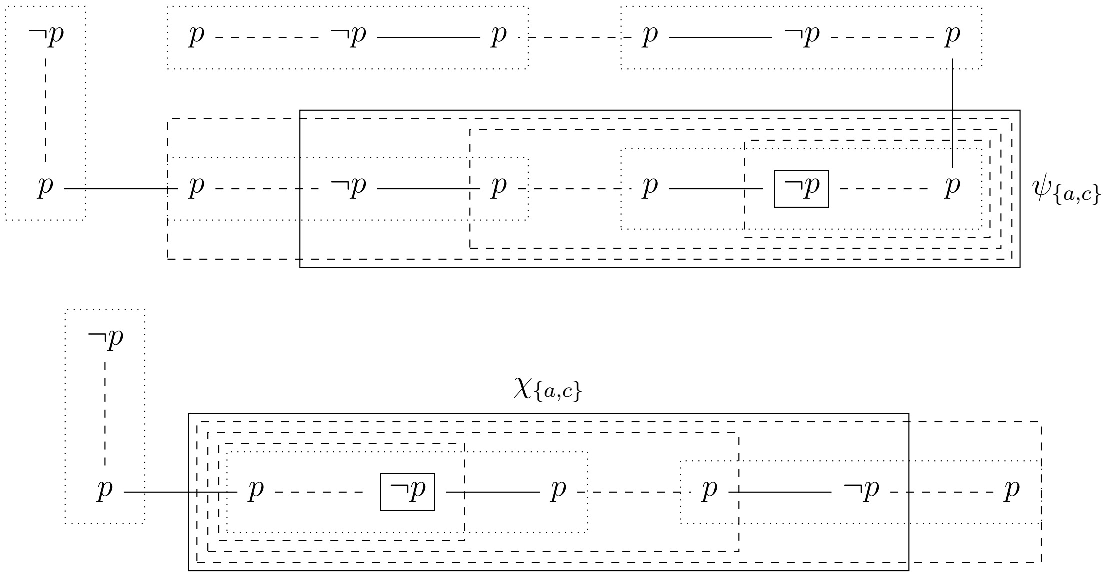

Figure 7.14: An a-model (above) and a b-model (below)

between two c-equivalence classes that b cut. Moreover, a submodel $(M,w)^{\psi'_{\{a,c\}}}$ of $(M,w)^{\psi_{\{a,c\}}}$ that is restricted by that b-cut can also be forced by $\{a,c\}$ (because c can cut this relation as well). Thus, replicating the corresponding move in (N,v) allows the existential player to have a winning strategy in a b-model no matter what agent b announces at the same time, and in this case the set of responses by b will be a subset of those she had in the a-model. This means that $\ulcorner(N,v),[\{b\},\chi_{\{a,c\}}]\psi)$ is also a winning node for the $\exists$-player. Hence, a contradiction. For an example, see Figure 7.15, where the set of counter-announcements by b in a b-model is a subset of the set of counter-announcements by b in an a-model.

Case $\langle\!\langle\{b,c\}\rangle\!\rangle\psi$ is similar to the previous one.

Case $\langle\!\langle\{a\}\rangle\!\rangle\psi$. Similar to the case $\langle\!\langle\{a,c\}\rangle\!\rangle\psi$. If a cannot make (N,v) isomorphic to $(M,w)^{\psi_{\{a\}}}$, then it is enough to cut a b-relation between two c-equivalence classes and 'announce' such a subset of $(M,w)^{\psi_{\{a\}}}$. In this case, it is still an a-announcement in the b-model, as well as it is one of the counter-announcements by $\{b,c\}$ in the a-model (c cuts b's relation). Hence, the set of counter-announcements in the b-model is the subset of counter-announcements in the a-model.

Cases $\langle\!\langle\{b\}\rangle\!\rangle\psi$ and $\langle\!\langle\{c\}\rangle\!\rangle\psi$ are as the previous one.

This completes the proof. □

Combining Theorems 7.5 and 7.6, we obtain the final result.

Theorem 7.7. GAL $\not\preceq$ CAL.

In case $\langle\!\langle\{a,b,c\}\rangle\!\rangle\psi$ of the proof agents a, b, and c can together enforce any configuration of a given model. This is due to the fact that the intersection of the corresponding relations is an identity relation. Hence, for chain models

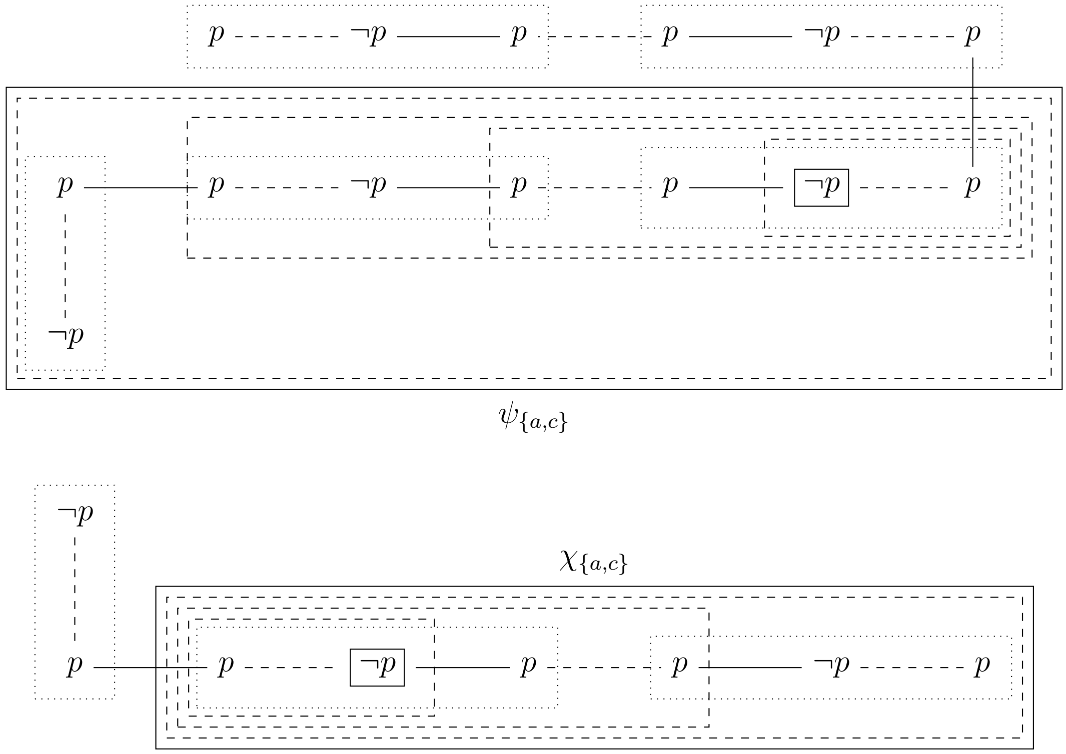

Figure 7.15: An a-model (above) and a b-model (below)

$\langle\!\langle \{a,b,c\} \rangle\!\rangle \psi$ is equivalent to $\Diamond\psi$ (and $[\!\langle \{a,b,c\} \rangle\!]\psi$ is equivalent to $\Box\psi$), and we have the following result:

Corollary 7.8. GAL $\not\preceq$ APAL.

That GAL $\not\preceq$ APAL was conjectured in Ågotnes et al. [2010], where it was also shown that APAL $\not\preceq$ GAL. Now we combine these two results.

Theorem 7.9. APAL and GAL are incomparable.

Chapter 8

Conclusion

The final chapter of the thesis serves two purposes. First, we briefly recapitulate the main results of our work. After that, we present a potpourri of related open research questions we find particularly interesting.

8.1 What Has Been Done

The complexity of the model checking problem for coalition announcement logic was shown to be PSPACE-complete in Chapter 4. To obtain this result we used an alternative equivalent semantics for the coalition announcement operator, where instead of quantification over an infinite set of possible epistemic announcements by agents, we consider a finite number of strategies. The latter are unions of equivalence classes and defined on finite models with the help of distinguishing formulas. We also considered a special case of the model-checking problem when formulas within scopes of coalition operators are positive. In this case the complexity of the model-checking algorithm is in P.

In Chapter 5 we considered validity and non-validity of various formulas of group and coalition announcement logics. Some of those formulas were mentioned as open questions in the literature. The most interesting result is non-validity of $\langle\!\langle G \rangle\!\rangle \varphi \leftrightarrow \langle G \rangle [A \setminus G] \varphi$. We proved it in the left-to-right direction, and presented a counterexample for the right-to-left direction. Among other interesting results is non-validity of the Church-Rosser principle in both GAL and CAL.

Since the introduction of coalition announcement operators in 2008 [Ågotnes and van Ditmarsch, 2008], no complete and sound axiomatisations of logics that include the operators have been given. We present such a logic in Chapter 6. Apart from the PAL base and coalition announcements, our logic includes relativised group announcements. This new operators $[G, \chi]\varphi$ mean that 'given a formula χ, whatever agents from G announce, φ holds after their announcement has been made jointly with χ.' Relativised group announcements allowed us to tame the inherent alternation of quantifiers in the coalition operators.

We show that CAL and APAL are not at least as expressive as GAL in Chapter 7. To achieve this result, we introduce formula games between the universal and the existential players. Relativised group announcements help us once again by separating moves corresponding to coalition announcements and anti-coalition responses. After that we define two classes of pointed models and show that there is a GAL formula that distinguishes them, whereas no CAL or APAL formula can do that. We also modify the existing proof to demonstrate that CAL is not at least as expressive as APAL.

8.2 What Is Yet To Be Done

Axiomatisation of CAL. First and foremost, a sound and complete axiomatisation of coalition announcement logic has been an open question since [Ågotnes and van Ditmarsch, 2008]. Moreover, the completeness proof of such an axiomatisation would be different from the proofs in [Balbiani et al., 2008] and [Balbiani and van Ditmarsch, 2015]. In particular, double quantification in coalition announcements does not fit well with the proof of the Lindenbaum lemma, where we find witnesses for negations of a quantified announcement operator (e.g. the GAL operator). In the case of CAL, we would be required to add an infinite number of witnesses in a single step. Hence, if an axiomatisation of CAL is found, we expect its completeness proof to differ significantly from those of APAL, GAL, and CoRGAL.

Finding finitary axiomatisations. Axiomatisations of all logics of quantified public announcements discussed in the thesis – APAL, GAL, CAL, CoRGAL – include infinitary rules of inference. It is unknown whether finitary axiomatisations of the logics exist. The non-existence of the one for a somewhat related logic – arbitrary arrow update logic with common knowledge (AAULC) – was demonstrated in [Kuijer, 2017]. Applying the proof for AAULC to a logic with quantified announcements and without common knowledge seems exciting and challenging at the same time. Another way of approaching the problem of finitary axiomatisations has been presented in [van Ditmarsch and French, 2017], where the authors considered a restriction of APAL to arbitrary announcements of boolean formulas only. The resulting logic, boolean arbitrary public announcement logic (BAPAL), has a complete finitary axiomatisation. Thus, a more viable open problem would be providing finitary axiomatisations for restricted versions of GAL and CAL, e.g. restricting agents' group and coalition announcements to announcements of boolean formulas only.

Finding decidable fragments. The satisfiability problem for all of APAL, GAL, and CAL is undecidable [Ågotnes et al., 2016]. Hence, finding decidable fragments of those logics is an open problem. And there are already some results in this direction. For example, BAPAL, which is a quite well-behaved logic, is decidable. The same can be said about another restriction of APAL – $APAL^+$ [van Ditmarsch et al., 2018]– where quantification is only over positive EL formulas. Testing similar restrictions of GAL and CAL is an intriguing avenue of further research.

Expressivity. All the known results on the expressivity of logics of quantified public announcements, including those that were presented in Chapter 7, are depicted in Figure 8.1.

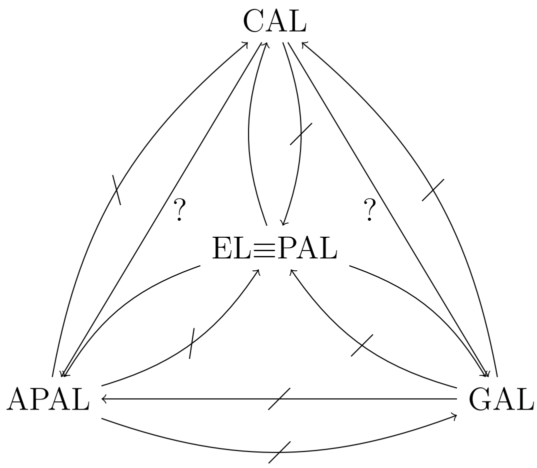

Figure 8.1: Relative expressivity of logics of quantified public announcements

In the figure an arrow $\rightarrow$ from logic L_1 to L_2 means that $L_1 \preceq L_2$. A struck-out arrow $\not\rightarrow$ denotes that $L_1 \not\preceq L_2$, and an arrow $\overset{?}{\rightarrow}$ labelled with a question mark indicates an open problem. Hence, whether CAL $\not\preceq$ APAL and CAL $\not\preceq$ GAL is yet to be answered. We conjecture that all logics of quantified public announcements are incomparable with one another.

Adding common and distributed knowledge. To the best of our knowledge, extensions of quantified announcement logics with common or distributed knowledge have not been considered. Probably the closest work that touches upon that problem is [Kuijer, 2017], where the author shows that AAULC does not have a finitary axiomatisation. Common knowledge plays an integral role in establishing that result. An axiomatisation of AAULC is, however, an open question.

Epistemic planning. Another perspective on GAL and CAL may be offered by the problem of epistemic planning [Bolander and Andersen, 2011]. In such a setting the task would be to find a sequence of group or coalition announcements such that for some initial model and a joint goal formula, this sequence makes the formula true. A similar problem was mentioned as an open question in [Bolander, 2017]. Since in both GAL and CAL agents do not communicate within their groups, considering implicitly coordinated plans [Engesser et al., 2017], where agents are not allowed to negotiate in advance, seems to be yet another promising area of further research.

Beyond announcements. Public announcements can be considered as a special case of more general structures — action models [Baltag et al., 1998]. Action model logic (AML) allows us to reason about more complex epistemic events such as suspicion, private announcements, etc. A quantification over action models was considered in [Hales, 2013], where an axiomatisation of arbitrary action model logic (AAML) was presented. The axiom system is based on reduction axioms, and hence AAML is as expressive as K. Note that AAML is not based on EL, and extending the basis of the logic beyond K is an open question.

Another way of modifying AAML is to consider only actions executed by groups and coalitions of agents. In that setting we may require agents to know preconditions of a certain epistemic action in the same way we require agents to know formulas they announce in GAL and CAL. However, as it was pointed out, AAML does not require the accessibility relation to be an equivalence. In this sense, AAML is not epistemic. Hence, if we would like agents to know the preconditions of actions, we should take this into account. A somewhat similar logic with the epistemic basis was presented in [Bozzelli et al., 2014].

Bibliography

Ågotnes, T. and Alechina, N. (2012). Epistemic coalition logic: completeness and complexity. In van der Hoek, W., Padgham, L., Conitzer, V., and Winikoff, M., editors, *Proceedings of AAMAS 2012*, pages 1099–1106. IFAAMAS.

Ågotnes, T., Balbiani, P., van Ditmarsch, H., and Seban, P. (2010). Group announcement logic. *Journal of Applied Logic*, 8(1):62–81.

Ågotnes, T. and van Ditmarsch, H. (2008). Coalitions and announcements. In Padgham, L., Parkes, D. C., Müller, J. P., and Parsons, S., editors, *Proceedings of AAMAS 2008*, pages 673–680. IFAAMAS.

Ågotnes, T. and van Ditmarsch, H. (2011). What will they say? - public announcement games. *Synthese*, 179(1):57–85.

Ågotnes, T. and van Ditmarsch, H. (2014). Knowledge games and coalitional abilities. In Baltag, A. and Smets, S., editors, *Johan van Benthem on Logic and Information Dynamics*, pages 451–485. Springer.

Ågotnes, T., van Ditmarsch, H., and French, T. (2016). The undecidability of quantified announcements. *Studia Logica*, 104(4):597–640.

Ågotnes, T., Goranko, V., Jamroga, W., and Wooldridge, M. (2015). Knowledge and ability. In van Ditmarsch, H., Halpern, J. Y., van der Hoek, W., and Kooi, B., editors, *Handbook of Epistemic Logic*, pages 543–589. College Publications.

Ågotnes, T. and Wáng, Y. N. (2017). Resolving distributed knowledge. *Artificial Intelligence*, 252:1–21.

Alur, R., Henzinger, T. A., and Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713.

Andréka, H., van Benthem, J., and Németi, I. (1995). Back and forth between modal logic and classical logic. *Logic Journal of the IGPL*, 3(5):685–720.

Balbiani, P. (2015). Putting right the wording and the proof of the truth lemma for *APAL*. *Journal of Applied Non-Classical Logics*, 25(1):2–19.

Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., and de Lima, T. (2008). 'Knowable' as 'known after an announcement'. *Review of Symbolic Logic*, 1(3):305–334.

Balbiani, P. and van Ditmarsch, H. (2015). A simple proof of the completeness of APAL. *Studies in Logic*, 8(1):65–78.

Baltag, A., Moss, L. S., and Solecki, S. (1998). The logic of public announcements, common knowledge, and private suspicions. In Gilboa, I., editor, *Proceedings of TARK 1998*, pages 43–56. Morgan Kaufmann.

van Benthem, J. (2001). Games in dynamic-epistemic logic. *Bulletin of Economic Research*, 53(4):219–248.

van Benthem, J. (2010). *Modal Logic for Open Minds*, volume 199 of *CSLI Lecture Notes*. CSLI Publications.

van Benthem, J. (2014). *Logic in Games*. MIT Press.

van Benthem, J., van Eijck, J., and Kooi, B. (2006). Logics of communication and change. *Information and computation*, 204(11):1620–1662.

Blackburn, P., van Benthem, J., and Wolter, F., editors (2006). *Handbook of modal logic*. Elsevier.

Blackburn, P., de Rijke, M., and Venema, Y. (2001). *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press.

Bolander, T. (2017). A gentle introduction to epistemic planning: The DEL approach. In Ghosh, S. and Ramanujam, R., editors, *Proceedings of Methods for Modalities 2017*, volume 243 of *EPTCS*, pages 1–22.

Bolander, T. and Andersen, M. B. (2011). Epistemic planning for single and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34.

Bozzelli, L., van Ditmarsch, H., French, T., Hales, J., and Pinchinat, S. (2014). Refinement modal logic. *Information and Computation*, 239:303–339.

van Ditmarsch, H. (2012). Quantifying notes. In *Proceedings of WoLLIC 2012*, pages 89–109.

van Ditmarsch, H., van Eijck, J., Sietsma, F., and Wang, Y. (2012). On the logic of lying. In van Eijck, J. and Verbrugge, R., editors, *Games, Actions and Social Software*, volume 7010 of *LNCS*, pages 41–72. Springer.

van Ditmarsch, H., Fernández-Duque, D., and van der Hoek, W. (2014). On the definability of simulation and bisimulation in epistemic logic. *Journal of Logic and Computation*, 24(6):1209–1227.

van Ditmarsch, H. and French, T. (2017). Quantifying over boolean announcements. *CoRR*, abs/1712.05310.

van Ditmarsch, H., French, T., and Hales, J. (2018). Positive announcements. *CoRR*, abs/1803.01696.

van Ditmarsch, H., van der Hoek, W., and Kooi, B. (2008). *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer.

van Ditmarsch, H., van der Hoek, W., and Kuijer, L. B. (2016). Fully arbitrary public announcements. In Beklemishev, L. D., Demri, S., and Maté, A., editors, *Proceedings of AiML 2016*, pages 252–267. College Publications.

van Ditmarsch, H. and Kooi, B. (2006). The secret of my success. *Synthese*, 153(2):339–339.

van Ditmarsch, H. and Kooi, B. (2008). Semantic results for ontic and epistemic change. In Bonanno, G., van der Hoek, W., and Wooldridge, M., editors, *Proceedings of LOFT 7*, volume 3 of *Texts in Logic and Games*, pages 87–118. Amsterdam University Press.

Engesser, T., Bolander, T., Mattmüller, R., and Nebel, B. (2017). Cooperative epistemic multi-agent planning for implicit coordination. In Ghosh, S. and Ramanujam, R., editors, *Proceedings of Methods for Modalities 2017*, volume 243 of *EPTCS*, pages 75–90.

Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. (2004). *Reasoning About Knowledge*. MIT press.

French, T., van der Hoek, W., Iliev, P., and Kooi, B. (2013). On the succinctness of some modal logics. *Artificial Intelligence*, 197:56–85.

Galimullin, R. and Alechina, N. (2017). Coalition and group announcement logic. In Lang, J., editor, *Proceedings of TARK 2017*, volume 251 of *EPTCS*, pages 207–220.

Galimullin, R. and Alechina, N. (2018). Coalition and group announcement logic. *CoRR*, abs/1810.02769.

Galimullin, R., Alechina, N., and van Ditmarsch, H. (2018). Model checking for coalition announcement logic. In Trollmann, F. and Turhan, A.-Y., editors, *KI 2018: Advances in Artificial Intelligence*, volume 11117 of *LNCS*, pages 11–23. Springer.

Goldblatt, R. (1982). *Axiomatising the Logic of Computer Programming*, volume 130 of *LNCS*. Springer.

Hales, J. (2013). Arbitrary action model logic and action model synthesis. In *Proceedings of LICS 2013*, pages 253–262. IEEE Computer Society.

Halpern, J. Y. and Moses, Y. (1992). A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial intelligence*, 54(3):319–379.

Harel, D., Kozen, D., and Tiuryn, J. (2000). *Dynamic Logic*. MIT Press.

Hintikka, J. (1962). *Knowledge and Belief. An Introduction to the Logic of the Two Notions*. Cornell University Press.

Hoshi, T. (2010). Merging DEL and ETL. *Journal of Logic, Language and Information*, 19(4):413–430.

Kuijer, L. B. (2017). Arbitrary arrow update logic with common knowledge is neither RE nor co-RE. In Lang, J., editor, *Proceedings of TARK 2017*, volume 251 of *EPTCS*, pages 373–381.

de Lima, T. (2014). Alternating-time temporal dynamic epistemic logic. *Journal of Logic and Computation*, 24(6):1145–1178.

Lutz, C. (2006). Complexity and succinctness of public announcement logic. In Nakashima, H., Wellman, M. P., Weiss, G., and Stone, P., editors, *Proceedings of AAMAS 2006*, pages 137–143. ACM.

Moss, L. (2015). Dynamic epistemic logic. In van Ditmarsch, H., Halpern, J. Y., van der Hoek, W., and Kooi, B., editors, *Handbook of Epistemic Logic*, pages 261–312. College Publications.

Pacuit, E. (2017). *Neighborhood Semantics for Modal Logic*. Springer.

Paige, R. and Tarjan, R. E. (1987). Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989.

Pauly, M. (2002). A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166.

Plaza, J. (2007). Logics of public communications (reprint of 1989's paper). *Synthese*, 158(2):165–179.

Renne, B., Sack, J., and Yap, A. (2016). Logics of temporal-epistemic actions. *Synthese*, 193(3):813–849.

Wang, Y. and Cao, Q. (2013). On axiomatizations of public announcement logic. *Synthese*, 190(1):103–134.