Refinement Quantifiers for Logics of Belief and Knowledge

James Hales

This report is submitted as partial fulfilment of the requirements for the Honours Programme of the School of Computer Science and Software Engineering, The University of Western Australia, 2014

Acknowledgements

First and foremost I would like to thank my supervisors, Tim French and Rowan Davies. Tim and Rowan have had a remarkable level of involvement in my honours project, supporting me greatly in the last year with frequent meetings and proof-readings of draft papers. It's thanks to them that, in my honours year I submitted a paper to the Methods for Modalities conference in Osuna, Spain, which was ultimately accepted for publication, an experience that has further motivated me to pursue a career in academia. Tim and Rowan must have reviewed dozens of draft copies of my paper and thesis, providing helpful feedback each time, and I am truly grateful for the investment of time that they have both made in supervising my honours project. I would also like to acknowledge the financial support that Tim and Rowan have granted me, which has allowed me to attend the conference in Osuna.

I would like to thank John Slaney and Rajeev Goré of the Australian National University, for the Logic Summer School that I attended in 2010. John Slaney organises the Logic Summer School at ANU each year, and it was thanks to him and ANU that I received a scholarship to attend the summer school, where I received my first in-depth introduction to logic. Rajeev Goré in particular gave me my first course in modal logic, from which I gained a lot, and Rajeev also offered me advice about the topic of my honours thesis, which I ultimately followed.

I would also like to thank Hans van Ditmarsch, and the anonymous reviewers of the Methods for Modalities conference that I submitted a paper to this year. Much of the recommendations that I received from the reviewers, and also from Hans, has made it into this thesis.

Finally I would like to acknowledge the financial support of the Hackett Foundation for the Hacket Foundation Alumni Honours Scholarship, which has made my life easier in the last year, and has partly funded my trip to Osuna for the Methods for Modalities conference.

Abstract

Modal logics are extensions of propositional logic, with which one can qualify the truth of statements with operators known as modalities. Epistemic logic is a variant of modal logic, commonly known as the logic of knowledge. Modalities in epistemic logic qualify statements by saying that a particular agent knows a statement to be true. Thus epistemic logic can be used to reason about the knowledge of a collection of agents. Doxastic logic is a similar logic, used to reason about beliefs rather than knowledge.

Informative updates are events that provide agents in modal settings with additional information. The effect of an informative update may be to give new knowledge or beliefs to the agents in the system. Informative updates in an epistemic setting can be modelled by the refinements of a Kripke model in modal logic.

Refinement quantifiers are introduced to variants of modal logic to produce refinement quantified modal logics. The refinement quantifiers are operators that quantify over the refinements of a Kripke model, and as these refinements model informative updates, this quantification can be said to be equivalent to quantifying over the informative updates that are possible in a Kripke model. Recent work by van Ditmarsch, French and Pinchinat [23] has presented an axiomatisation and decidability results for the single-agent refinement quantified modal logic. We extend these results to apply to the single-agent doxastic and epistemic logics, and to the multi-agent modal and doxastic logics, by providing sound and complete axiomatisations, and decidability and expressivity results for each of these logics.

Keywords: Epistemic Logic, Refinement Quantification, Axiomatizations, Decision Procedures.
CR Categories: F.4, I.2.4

Contents

CHAPTER 1

Introduction

Modal logics are extensions of propositional logic, used to reason about the properties of relational structures known as *Kripke models*, by qualifying the truth of statements with operators known as *modalities*. *Epistemic logic* is a variant of modal logic, commonly known as the logic of knowledge, where the Kripke models that are considered encode the knowledge that a collection of agents hold about the state of the world, and the effect of modalities is to say that an agent *knows* a statement to be true. *Doxastic logic* is a similar variant of modal logic, the logic of belief, that considers the beliefs that a collection of agents hold. The difference between the notions of knowledge and belief used in epistemic and doxastic logics is simply that anything that an agent knows must be true, whereas it is permissible for an agent to believe something that is false.

Dynamic epistemic logic and *dynamic doxastic logic* are fields that consider how the knowledge or beliefs of a collection of agents may change in response to *informative updates*. Informative updates are events, such as messages or announcements, that communicate new information to agents. Informative updates change the knowledge or beliefs of agents, whilst leaving information directly about the state of the world unchanged. These fields consider questions such as what knowledge or belief states can result from specific informative updates, and how or whether a specific knowledge or belief state can be achieved from a series of arbitrary informative updates, perhaps with restrictions on what kinds of informative updates are allowed.

A classic example is that of the muddy children puzzle, explained by van Ditmarsch, van der Hoek and Kooi [24]:

> "A group of children has been playing outside and are called back into the house by their father. The children gather round him. As one may imagine, some of them have become dirty from the play and in particular: they may have mud on their forehead. Children can only see whether other children are muddy, and not if there is any mud on their own forehead. All this is commonly known, and the

> children are, obviously, perfect logicians. Father now says: At least one of you has mud on his or her forehead. And then: Will those who know whether they are muddy please step forward. If nobody steps forward, father keeps repeating the request."

The puzzle is whether each of the children can eventually determine whether they are muddy or not. The children are not allowed to communicate with one another, except by stepping forward when they know whether they are muddy or not. It can be shown that if there are m muddy children, those children will step forward after the father has asked his question m times, at which point the remaining children will know that they are not muddy, and so they too will step forward the next time the father asks his question. The fact that this is the case can be shown by modelling the knowledge of the children using a Kripke model, and the father's initial announcement, and each round of the children stepping forward or not stepping forward, in response to the father's requests, as informative updates.

The muddy children puzzle demonstrates how agents in an epistemic system can gain knowledge about the state of the world indirectly. In the case of m muddy children, the first $m-1$ times that the father asks his question, no child steps forward, and so the only information that is given is that none of the children knows whether or not they are muddy. Despite being given the same information over and over, that does not directly provide any information about whether the children are muddy or not, this is enough for each child to eventually determine whether or not they are muddy, and this fact is surprising.

Problems of this form motivate our research into dynamic epistemic logic and dynamic doxastic logic; we wish to know whether it is possible for agents to achieve certain knowledge or belief after some sequence of informative updates. This is an interesting question because, as we have already demonstrated with the muddy children puzzle, it is not always obvious whether or not certain knowledge or belief states are attainable. In the particular example of the muddy children puzzle, the fact that the desired knowledge is attainable is easy to deduce, since a specific sequence of informative updates is given. More generally however we are interested in situations where, rather than performing a specific sequence of informative updates, any informative update may be performed. This has applications in areas such as games with imperfect information, where one may wish to know whether certain knowledge can be deduced from given information, or in security protocols, where one may desire assurances that certain knowledge cannot be indirectly deduced from seemingly innocuous communication.

Previous work has considered a number of different representations of informative updates. In the context of epistemic logic, two of the most notable are *public* *announcements*, and *action models*. Public announcements are a relatively simple and restrictive form of informative update, whereas actions models are more general and expressive. We avoid discussing representations in doxastic logic for the time being, as the most notable, belief revision, is conceptually different from the epistemic notions of information change that the present work is based on.

As we are interested in questions about what knowledge or belief states can result from informative updates, we consider extensions of epistemic or doxastic logics that allow us to reason about this. There are two main directions that we can take with such extensions. The first is logics with which we can reason about the results of *specific* informative updates; in such a logic we can say of some specific informative update that "after the informative update α, the statement ϕ is true in the resulting knowledge state". The second is logics with which we can reason about the results of *arbitrary* informative updates, by quantifying over them; in such a logic we can say that "there exists an informative update after which ϕ is true in the resulting knowledge state".

The *public announcement logic*, introduced by Plaza [15], and also independently by Gerbrandy and Groenvald [11] is an extension of epistemic logic that introduces an operator for reasoning about the results of a specific public announcement. A similar logic, the *action model logic*, introduced by Baltag and Moss [5], allows reasoning about the results of executing a specific action model. Balbiani, et al. [4] then explored the *arbitrary public announcement logic*, an extension of the public announcement logic that adds an operator for reasoning about the results of arbitrary public announcements. It was later shown by van Ditmarsch and French [3] that the arbitrary public announcement logic is undecidable. It has not been considered whether a similar extension of action model logic would be decidable or not, but the effect of the undecidability result in arbitrary public announcement logic has been to encourage research into weaker versions of this logic that are more likely to be decidable and do not rely on the relatively complicated semantics of action models.

The *future event logic*, introduced by van Ditmarsch and French [22] is an extension of epistemic logic that introduces an operator for quantifying over the *refinements* of a Kripke model. The finite refinements of a Kripke model can be shown to correspond to the results of executing arbitrary action models on that Kripke model, and so the effect of quantifying over the refinements of a Kripke model can reasonably be said to be equivalent to quantifying over arbitrary informative updates. van Ditmarsch, French and Pinchinat [23] gave an axiomatisation and decidability results for a simplified version of this logic, in the context of single-agent modal logic, as opposed to multi-agent epistemic logic.

The present work extends the future event logic to the setting for which it was originally intended: epistemic logic. We will refer to the future event logic using the more general name of *refinement quantified modal logics*. We consider refinement quantified versions of modal, doxastic and epistemic logics. We provide axiomatisations for the refinement quantified epistemic and doxastic logics in the single-agent case, and for the refinement quantified modal and doxastic logics in the multi-agent case. Our axiomatisations provide translations from the refinement quantified modal logics into the corresponding basic modal logics; this gives us decidability and expressivity results for these logics.

CHAPTER 2

Literature review

We review previous techniques for modelling information change in epistemic logic and doxastic logic, and previous work in the refinement quantified modal logic, that motivates the present work.

2.1 Information change in epistemic logics

Dynamic epistemic logic is a general term for the study of information change in modal epistemic systems. Information change is performed by informative updates, events that provide agents with additional information, whilst leaving information directly about the state of the world unchanged. Notable representations for informative updates in dynamic epistemic logic include *public announcements* and *action models*. I will briefly discuss these representations of informative updates, and the logics that have been devised to reason about them.

2.1.1 Public announcements

A public announcement is a simple form of informative update, where a true statement is announced to all agents in a system at once [24]. A public announcement takes the form of an epistemic formula α. As α is announced to all agents at once, one result of the announcement is that afterwards every agent knows that α is true. In addition to this, as the announcement was made publicly, it becomes *common knowledge*. Common knowledge means that every agent knows α, every agent knows that every agent knows that α, and so on ad infinitum. In terms of the Kripke model representation of the epistemic system, a public announcement has the effect of restricting the states of the Kripke model to those states that are consistent with α, as given the additional information that α is the case, those states that are inconsistent with α are now considered impossible by every agent.

The public announcement logic was axiomatised by Plaza [15], and also independently by by Gerbrandy and Groenvald [11]. The logic introduces an operator $[\alpha]$, where α is an epistemic formula, that can be used to reason about the results of a specific public announcement. The formula $[\alpha]\phi$ means that after the formula α is publicly announced, ϕ holds in the resulting epistemic state. This allows one to reason about the consequences of specific public announcements in the epistemic system.

Public announcements are a very limited form of informative update, because the information contained in a public announcement necessarily becomes common knowledge to all agents in the system. Public announcements cannot for example model informative updates that provide information privately to only some of the agents in the system. Public announcements are however suited to some interesting problems: for example the muddy children puzzle that was discussed in the introduction to this thesis can be modelled in terms of public announcements.

2.1.2 Action models

Action models capture a more general notion of informative updates than public announcements. Compared to public announcements, action models are able to represent informative updates that provide information to only a subset of the agents, or where an agent may be aware that one of several possible informative updates has occurred, but is uncertain as to which update actually occurred.

For example, suppose that Alice and Bob have made a bet on the flip of a coin. Alice flips the coin, and holds it on the back of her hand so that neither person can see what it is initially. Neither Alice nor Bob know whether the coin landed on heads or tails, and each of them knows that the other is ignorant of the result of the flip (i.e. it is common knowledge). Alice then lifts her hand, and without letting Bob see the coin, sees that the coin landed on heads. Bob saw Alice do this, but he didn't see what the coin landed on. The act of Alice looking at the coin is an informative update, one result of which is that Alice now knows that the coin landed on heads, whilst Bob still doesn't know. Thus the effect of this informative update was to provide different information to Alice than to Bob. Furthermore, Bob saw Alice look at the coin, so he's aware that an informative update has occurred, but is uncertain as to which update actually occurred. From the point of view of Bob, Alice may have learned that the coin landed on heads, or that it landed on tails; Bob cannot distinguish between these two possibilities. Thus Bob now knows that either Alice knows that the coin landed on heads, or that Alice knows that the coin landed on tails, but doesn't know which is actually the case.

Action models are represented by relational structures, similar to Kripke models [24]. We write (M, s) to denote a *pointed action model*, where M is the action model itself, and s is the node in the action model that represents the actual informative update that takes place. An action model is executed on a Kripke model, resulting in a new Kripke model. If the action model is said to represent an informative update, then the new Kripke model represents the resulting knowledge state after the informative update has been performed.

The logic of action models was introduced by Baltag, Solecki and Moss [5]. Similar to the public announcement logic, an operator $[\mathrm{M}, \mathrm{s}]$ is introduced that has the effect of performing an informative update, this time in the form of executing an action model on the Kripke model that represents the current epistemic state. The formula $[\mathrm{M}, \mathrm{s}]\phi$ means that after the action model (M, s) is executed in the current epistemic state, ϕ holds in the resulting epistemic state.

Epistemic actions are a similar representation of informative updates. van Ditmarsch [20, 21, 24] introduced the logic of epistemic actions, a logic similar to the logic of action models. The difference between epistemic actions and action models is that an epistemic action is represented as a formula instead of as a relational structure. The formula may contain operators whose purpose is to limit the effect of informative updates to a certain subset of agents, or to give the effect that one of several informative updates may have taken place, but from the point of view of some of the agents, which update has actually occurred is unclear. This can be compared to a public announcement, which is also represented as a formula.

2.1.3 Arbitrary public announcement logic

Balbiani, Baltag, van Ditmarsch, et al. [4] introduced the arbitrary public announcement logic, an extension of the public announcement logic, that provides an operator for quantifying over arbitrary public announcements. It introduces an operator, $\Box\psi$, which means that after any arbitrary public announcement, ψ holds. Its dual operator, $\Diamond\psi$ means that after some public announcement, ψ holds.

The same paper briefly discusses a possible arbitrary action model logic, a generalisation of arbitrary public announcement logic allowing for any kind of informative update, modelled as the execution of an action model. French and van Ditmarsch [3] later showed that the arbitrary public announcement was undecidable. Whilst the arbitrary action model logic has not been considered in depth, the undecidability result in arbitrary public announcement logic has encouraged research into weaker versions of this logic instead, as they are more likely to be decidable and do not rely on the relatively complicated semantics of action models.

2.2 Information change in doxastic logics

In dynamic epistemic logic, we consider how knowledge changes in response to informative updates that provide agents with additional information. A key property of epistemic systems is that everything that an agent knows must be true, and any information that an agent is provided with through an informative update must also be true. A consequence of this is that once agents have gained certain information, they do not lose this information; further informative updates cannot cause the agent to forget past information, or to reconsider the truth of previous statements. By contrast, in doxastic systems, the beliefs that an agent holds do not necessarily have to be true, and the information that agents are given also do not have to be true. Therefore it is reasonable for an agent to reconsider its beliefs, and possibly reject past information, if it is provided with new information that contradicts its old beliefs. We will briefly discuss *belief revision*, which models the revision and possible removal of old beliefs, and *action models*, which can be generalised from epistemic systems to doxastic systems.

2.2.1 Belief revision

The AGM approach to belief revision, named for Alchourrn, Gärdenfors and Makinson [2], who did much of the initial work in this area, considers the beliefs of an agent as a set of propositional formulae, called the belief set, that the agent believes to be true. Informative updates are represented by an operation on the belief set, called a *revision*, wherein a new propositional formula is incorporated into the belief set, and the existing beliefs are revised to accommodate it. The process of belief revision involves adding the new information to the belief set, whilst rejecting existing beliefs that may contradict with the new information. Revision can be defined as an operation over a set of propositional formulae, and so reasoning about belief revision can be done directly in these terms. Dynamic doxastic logic is a logic similar to the previously discussed public announcement logics, and action model logics, which introduces an operator for reasoning about the result of an informative update in terms of belief revision [24]. Variants of dynamic doxastic logic use different operators for belief revision. Segerberg [17] provided axiomatisation for some variants of dynamic doxastic logics.

2.2.2 Action models

The same system of action models for modelling informative updates in epistemic models can be generalised to apply to doxastic models [24]. Unlike action models applied in the epistemic setting, action models applied in the doxastic setting do not necessarily have to represent a truthful informative update, so it is possible for agents to come to believe statements that are not actually true. By contrast to belief revision, it is not possible for action models to cause an agent to reject past information. In effect, once an agent has been provided with information, it is committed to believe that information forever.

It is also possible for an action model to cause a belief that was once founded to become unfounded. For example, in the previous coin flipping example, if Alice had taken a peek at the coin without Bob's seeing her do so, then the result of this is that Alice now believes that the coin landed on heads, whilst Bob continues to believe (incorrectly) that Alice does not believe either way that the coin landed on heads or tails. Bob continues to hold his initial belief because he has not been exposed to any information that would cause that belief to change. Such a change is not permissible in an epistemic setting, where knowledge must be founded, but it is permissible in a doxastic setting, where beliefs are allowed to be unfounded.

2.3 Other related logics

2.3.1 Group announcement logics

The group announcement logic, introduced by Ågotnes et al. [1] is an extension of public announcement logic that allows one to reason about whether a group of agents within an epistemic system are able to cooperate in order to bring about some knowledge state. The agents in the group are able to cooperate by publicly announcing statements that the agents know, but are unable to announce statements that they do not know. The logic introduces an operator, $< G >$, where G is a subset of the agents, and the formula $< G > \phi$ means that there is some set of formulae, such every formula is known by at least one of the agents in G, and such that after each agent announces the formulae that they know, ϕ is true in the resulting knowledge state. Ågotnes et al. [1] provide a sound and complete axiomatisation of the group announcement logic, show that it is decidable, and give complexity and expressivity results for the logic. Notably, the group announcement logic is not as expressive as the arbitrary public announcement logic, and it is conjectured that the arbitrary public announcement is not as expressive as the group announcement logic either. The difference between group announcement logic and the arbitrary public announcement logic is that the group announcement logic essentially quantifies over the public announcements that the agents in the group know, whereas the arbitrary public announcement logic quantifies over all possible public announcements.

2.3.2 Propositional dynamic logics

Propositional dynamic logic is a variant of dynamic logic, first introduced by Fischer and Ladner [8]. Propositional dynamic logic is a logic that introduces a concept of performing an action, called a program, on a modal state. The programs in propositional dynamic logic are symbolic in nature, and do not have any inherent effect on the model that it is performed on. Programs can be combined, by composition, repetition, concurrent execution or by including a test, where an action is performed only if a particular formula is satisfied at the current state. Actions are performed using modal operators that are labelled with the actions themselves; hence there is a modal operator for every possible action. What differentiates propositional dynamic logics from standard modal logics is that the models under consideration have algebraic constraints on them, which preserve certain properties about the composition of programs, amongst other properties. Propositional dynamic logic is decidable, and has decision procedures which have an exponential running time in the worst case [16].

van Benthem, van Eijck and Kooi [18] showed that the action model logic can be translated to propositional dynamic logic, and hence propositional dynamic logic is at least as expressive as the action model logic. Miller and Moss however showed that adding quantification over action models, as in arbitrary action model logic, makes the logic undecidable [14].

2.3.3 Description logics

Description logics are logics used for knowledge representation, often used in applications involving ontologies, such as artificial intelligence or the Semantic Web. Description logics are similar to modal logics, in the sense that they are designed to be decidable fragments of first-order logic, and in fact many description logics are simply syntactic variations of modal logics [7]. Dynamic description logics, a hybrid of propositional dynamic logics and description logics have been considered by Wolter and Zakharyaschev [25]. Wolter and Zakharyaschev extended the description logic $\mathcal{ALC}$, which is a syntactic variation of the modal logic $\mathcal{L}^{\mathbf{K}}$, with propositional dynamic operators, provided an axiomatisation for the resulting logic and showed that it was decidable.

2.3.4 Bisimulation quantified modal logics

Bisimulations are a relationship between Kripke models, that preserves the truth of modal formulae interpreted over the related Kripke models. Bisimulations are defined formally in Section 3.2. Bisimulation quantified modal logics extend modal logics with an operator for quantifying over the models that are bisimilar to the Kripke model under consideration, but where a particular propositional atom may vary in its interpretation. In effect it is quantifying over the truth value of a propositional variable, in models bisimilar to the current Kripke model. French [9] considered a number of bisimulation quantified modal logics, and provided a translation to the modal μ-calculus, and decidability or undecidability results for several such logics. The notion of a bisimulation is related to the notion of a refinement, on which the refinement quantified modal logics are based.

2.4 Refinement quantified modal logics

The future event logic, introduced by van Ditmarsch and French [22] logic, is an extension of epistemic logic that provides an operator for quantification over arbitrary refinements of Kripke models. The semantics of the logic were provided by van Ditmarsch and French, along with several results and comparisons to justify the future event logic as accurately capturing the notion of quantifying over arbitrary informative updates. The semantics of the future event logic was also compared to previously defined logics, in particular to the bisimulation quantified epistemic logic, and to the proposed arbitrary action model logic. In this thesis we refer to the future event logic with the more general term of refinement quantified modal logics, to reflect the fact that refinement quantifiers are applicable in non-epistemic settings.

In their paper, van Ditmarsch and French [22] relate the refinement quantified modal logics to the bisimulation quantified modal logics. The two logics are closely related due to the relationships between bisimulations and refinements. The semantics of bisimulation quantifiers and refinement quantifiers vary significantly, as the bisimulation quantifiers quantify over bisimilar models, and also bind a propositional variable, whereas refinement quantifiers quantify over refinements and do not bind any variables. However van Ditmarsch and French show that because the refinements of a Kripke model correspond to domain restrictions of bisimilar Kripke models, the refinement quantifiers can be represented by bisimulation quantifiers that quantify over variables that do not appear elsewhere in the formula. This shows that the bisimulation quantified modal logics are at least as expressive as the refinement quantified modal logics. On the other hand, the axiomatisation of the bisimulation quantified modal logics was completed using a translation to the modal μ-calculus, whereas corresponding results in the refinement quantified modal logics so far have relied on translations to basic modal logics. This leads to more desireable qualities, such as better model-checking and decision procedures, and simpler axiomatisations of the refinement quantified modal logics.

A comparison with the refinement quantified modal logics was also given by van Ditmarsch and French to the proposed arbitrary action model logic [22]. The main difference between the two logics is that the refinement quantified epistemic logic does not include an operator for reasoning about the results of a specific informative update. van Ditmarsch and French [22] show that if such an operator is added to the refinement quantified epistemic logic, then the resulting logic is equivalent to the arbitrary action model logic. The semantics of the refinement quantifier in the refinement quantified epistemic logic are much simpler than the semantics of the action model quantifier in arbitrary action model logic, as they do not rely on the mechanics of action models.

Although not considered in detail by van Ditmarsch and French, we compare the related refinement quantified doxastic logic to previously defined notions of information change in doxastic logics. The refinement quantified epistemic logic quantifies over the refinements of Kripke models, which are equivalent to the execution of arbitrary action models on the Kripke model [22]. A consequence of this in the doxastic setting is that the notion of information change that the refinement quantified doxastic logic captures is closer to that of the action models that we have discussed earlier, rather than belief revision. In particular, this means that agents are committed to believing information after it has been provided to them, and that it is not possible for agents to reject information that they have previously been provided.

The refinement quantified modal logic is considered in more depth by van Ditmarsch, French and Pinchinat [23]. They consider the single-agent modal variant of the logic, rather than the epistemic variant. An axiomatisation of the logic is provided, and is proven sound and complete. The completeness proof for the axiomatisation is performed via a provably correct translation to single-agent modal logic. This translation shows that the single-agent refinement quantified modal logic is expressively equivalent to the single-agent modal logic. This expressivity result gives us several results about the single-agent refinement quantified modal logic from properties of basic modal logic, in particular that the single-agent refinement quantified modal logic is decidable. An explicit decision procedure was also given by van Ditmarsch, French and Pinchinat, that runs in 2EXP time using a tableau method, which gives an upper bound for the complexity of the logic. The refinement quantified modal logic is also shown to be exponentially more succinct than the basic modal logic, a result that hints that the decision procedure that was given is likely to be optimal. Extending the refinement quantified modal logic to the multi-agent cases, and to epistemic and doxastic variants, are left as future work.

CHAPTER 3

Technical preliminaries

In this chapter we will introduce the technical preliminaries used throughout this thesis. We begin by giving the definitions of modal, doxastic and epistemic logics that we use. We then introduce the concept of a refinement, and then define the syntax and semantics of the refinement quantified modal logics themselves.

3.1 Modal, doxastic and epistemic logics

In this section we recall the definitions of modal, doxastic and epistemic logics, and the notation that we use for these logics. We refer the reader to Appendix A for a more thorough introduction to modal logics.

Let A be a non-empty, finite set of agents, and let P be a non-empty, countable set of propositional atoms.

Definition 3.1.1 (Kripke model)**.** A *Kripke model* $M = (S, R, V)$ consists of a *domain* S, which is a set of states (or worlds), *accessibility* $R : A \to \mathcal{P}(S \times S)$, and a *valuation* $V : P \to \mathcal{P}(S)$.

The class of all Kripke models is called K. We write $M \in K$ to denote that M is a Kripke model.

For $R(a)$, we write R_a. We write sR_a for $\{t \mid (s, t) \in R_a\}$ and we write $R_a t$ for $\{s \mid (s, t) \in R_a\}$. As we will be required to discuss several models at once, we will use the convention that $M = (S^M, R^M, V^M)$, $N = (S^N, R^N, V^N)$, and so on. For $s \in S^M$ we will let M_s refer to the pair (M, s), also known as the pointed Kripke model of M at state s.

Definition 3.1.2 (Doxastic model)**.** A *doxastic model* is a Kripke model $M = (S, R, V)$ such that the relation R_a is serial, transitive, and Euclidean for all $a \in A$. The class of all doxastic models is called *KD45*. We write $M \in$ *KD45* to denote that M is a doxastic model.

Definition 3.1.3 (Epistemic model)**.** An *epistemic model* is a Kripke model $M = (S, R, V)$ such that the relation R_a is an equivalence relation for all $a \in A$. The class of all epistemic models is called *S5*. We write $M \in \mathit{S5}$ to denote that M is an epistemic model.

Throughout this thesis we will be presenting results in both doxastic and epistemic logic. As such, when we are discussing doxastic logic, we will assume that all Kripke models are implicitly doxastic models, and likewise when we are discussing epistemic logic, we will assume that all Kripke models are implicitly epistemic models. When we are discussing results in general modal logic, we will not assume any restrictions on the Kripke models.

Definition 3.1.4 (Language $\mathcal{L}$)**.** Given a non-empty, finite set of agents A and a non-empty, countable set of propositional atoms P, the language of $\mathcal{L}$ is defined by the following abstract syntax:

$$\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \Box_a \alpha$$

where $p \in P$, $a \in A$ and $\alpha \in \mathcal{L}$.

Standard abbreviations include: $\top ::= p \vee \neg p$ for some $p \in P$; $\bot ::= \neg\top$; $\phi \vee \psi ::= \neg(\neg\phi \wedge \neg\psi)$; $\phi \to \psi ::= \neg\phi \vee \psi$; and $\Diamond_a \phi ::= \neg\Box_a\neg\phi$.

We also use the cover operator $\nabla_a \Gamma$, where Γ is a finite set of formulae, which is an abbreviation for $\nabla_a \Gamma ::= \Box_a \bigvee_{\gamma \in \Gamma} \gamma \wedge \bigwedge_{\gamma \in \Gamma} \Diamond_a \gamma$. An axiomatisation of the modal μ-calculus, using the cover operator, was given by Bilkova, Palmigiano and Venema [6]. The cover operator is relied on for our axiomatisation, in much the same way it is relied on for the axiomatisation of $\mathcal{L}^{\mathbf{K}}_{\triangleright(1)}$ presented by van Ditmarsch, French and Pinchinat [23].

We note that the basic modalities $\Box_a$ and $\Diamond_a$ can be expressed in terms of the cover operator, a fact that we will rely upon for all of our axiomatisations. We note that $\Box_a \phi \leftrightarrow \nabla_a\{\phi\} \vee \nabla_a \emptyset$ and $\Diamond_a \phi \leftrightarrow \nabla_a\{\phi, \top\}$ [6].

Definition 3.1.5 (Semantics of $\mathcal{L}^{\mathbf{C}}$)**.** Let C be a class of Kripke models, and let $M = (S, R, V) \in C$ be a Kripke model taken from C. The interpretation of ϕ is defined inductively:

$$\begin{array}{lll}
M_s & \models & p \text{ iff } s \in V_p \\
M_s & \models & \neg\phi \text{ iff } M_s \not\models \phi \\
M_s & \models & \phi \wedge \psi \text{ iff } M_s \models \phi \text{ and } M_s \models \psi \\
M_s & \models & \Box_a \phi \text{ if for all } t \in S : (s,t) \in R_a \text{ implies } M_t \models \phi
\end{array}$$

We say that a formula ϕ is *satisfied* by a pointed Kripke model $M_s \in C$ if and only if $M_s \vDash \phi$. We say that ϕ is satisfied by a Kripke model $M \in C$ if and only if $M_s \vDash \phi$ for some $s \in S^M$. We say that ϕ is *satisfied* by a class of Kripke models C if and only if it is satisfied by every Kripke model $M \in C$. We say that ϕ is *valid* in a Kripke model $M \in C$ if and only if $M_s \vDash \phi$ for every $s \in S^M$. We write $M \vDash \phi$. We say that ϕ is *valid* in a class of Kripke models C if and only if $M \vDash \phi$ for every $M \in C$. We write $C \vDash \phi$.

The logics $\mathcal{L}^{\mathbf{K}}$, $\mathcal{L}^{\mathbf{KD45}}$ and $\mathcal{L}^{\mathbf{S5}}$ are instances of $\mathcal{L}^{\mathbf{C}}$ with classes K, *KD45* and *S5* respectively. It should be noted that $\mathcal{L}^{\mathbf{KD45}}$ is a conservative extension of $\mathcal{L}^{\mathbf{K}}$, and $\mathcal{L}^{\mathbf{S5}}$ is a conservative extension of $\mathcal{L}^{\mathbf{KD45}}$ (and also of $\mathcal{L}^{\mathbf{K}}$). This means that every valid formula in $\mathcal{L}^{\mathbf{K}}$ is also valid in $\mathcal{L}^{\mathbf{KD45}}$, and likewise for $\mathcal{L}^{\mathbf{KD45}}$ and $\mathcal{L}^{\mathbf{S5}}$. This is because any formula that is valid with respect to a particular class of Kripke models is also valid for any subclass of those Kripke models.

3.2 Bisimulation, simulation and refinements

In this section we introduce refinements, the concept that is central to the refinement quantified modal logics. We first introduce the related notion of a bisimulation, and define simulations and refinements in terms of the properties that define a bisimulation. We finally remark on the properties of refinements that make them suitable for representing informative updates in the refinement quantified modal logics.

Definition 3.2.1 (Bisimulation)**.** Let $M = (S, R, V)$ and $M' = (S', R', V')$ be Kripke models. A non-empty relation $\mathcal{R} \subseteq S \times S'$ is a *bisimulation* if and only if for all $s \in S$ and $s' \in S'$, with $(s, s') \in \mathcal{R}$, for all $a \in A$:

atoms $s \in V(p)$ if and only if $s' \in V'(p)$ for all $p \in P$

forth-a for all $t \in S$, if $R_a(s, t)$, then there is a $t' \in S'$ such that $R'_a(s', t')$ and $(t, t') \in \mathcal{R}$

back-a for all $t' \in S'$, if $R'_a(s', t')$, then there is a $t \in S$ such that $R_a(s, t)$ and $(t, t') \in \mathcal{R}$.

We call M_s and $M'_{s'}$ bisimilar if there is a bisimulation between M and M' linking s and s', and we write $M_s \leftrightarroweq M'_{s'}$ to denote this.

We note that bisimulation is an equivalence relation [7].

Bisimulations relate Kripke models that are in a sense indistinguishable to modal logics. If two Kripke models are bisimilar, then they each satisfy the same set of modal formulae. This is a property of modal logics that is called bisimulation invariance.

Lemma 3.2.1 (Bisimulation invariance)**.** *Let $M_s, M'_{s'} \in K$ be Kripke models such that $M_s \leftrightarrows M'_{s'}$, and let $\phi \in \mathcal{L}$. Then $M_s \vDash \phi$ if and only if $M'_{s'} \vDash \phi$.*

This is shown by Blackburn, de Rijke and Venema [7]. We note that this result holds equally if the models M_s and $M'_{s'}$ are taken from subclasses of K.

The notion of bisimulation invariance is an essential property of modal logics. Indeed, the modal logic **K** may be defined as the bisimulation invariant fragment of first order logic. Blackburn, de Rijke and Venema [7] discuss this property in a chapter of their book.

If we relax the properties of bisimulation slightly, we get a related notion known as a *simulation*, or its reverse relation, a *refinement*.

Definition 3.2.2 (Simulation and refinement)**.** Let M and M' be Kripke models. A non-empty relation $\mathcal{R} \subseteq S \times S'$ is a *simulation* if and only if it satisfies **atoms** and **forth-**a for every $a \in A$.

We call $M'_{s'}$ a simulation of M_s and we call M_s a refinement of $M'_{s'}$. We write $M'_{s'} \succeq M_s$ to denote this, or alternatively, $M_s \preceq M'_{s'}$.

A relation that satisfies **atoms** and **forth-**b for every $b \in A$, and satisfies **back-**b for every $b \in A - \{a\}$, for some $a \in A$, is an *a-simulation*.

We call $M'_{s'}$ an a-simulation of M_s, and we call M_s an a-refinement of $M'_{s'}$. We write $M'_{s'} \succeq_a M_s$ to denote this, or alternatively, $M_s \preceq_a M'_{s'}$.

Lemma 3.2.2. *The relation $\preceq_a$ is reflexive and transitive, and satisfies the Church-Rosser property over the class of K models.*

This is shown by van Ditmarsch and French [22]

It should be clear that bisimulations and refinements are closely related notions, as their definitions reuse the properties of **atoms**, **forth** and **back**. Bisimulations require all three properties, whilst refinements relax the requirement for the **back** property. As a consequence, refinements satisfy weaker versions of the properties that bisimulations satisfy. Whilst the bisimulation relation is an equivalence relation, the refinement relation is a pre-order relation; moreover, we note that the bisimulation relation is the equivalence relation that is naturally induced from the pre-order relation of refinement. Furthermore, whilst bisimulations of models preserve the truth of all modal formulae, refinements of models preserve a restricted class of modal formulae, known as positive formulae, which we define shortly.

As we have mentioned previously, refinement quantified modal logics quantify over informative updates, by quantifying over the refinements of a Kripke model. We will briefly introduce two properties of refinements to illustrate how refinements correspond to the notion of an informative update.

Definition 3.2.3 (Positive formulae)**.** A positive formula is defined by the following abstract syntax:

$$\alpha ::= p \mid \neg p \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \Box_a \alpha$$

where $p \in P$ and $a \in A$.

Proposition 3.2.3. *Let M_s and $M'_{s'}$ be Kripke models such that $M'_{s'} \preceq M_s$, and let ϕ be a positive formula. If $M_s \models \phi$ then $M'_{s'} \models \phi$.*

This is shown by van Ditmarsch and French [22]

The significance of positive formulae is that they are the formulae that should be preserved in the result of any informative update.

Using an epistemic interpretation, one can reasonably say that an informative update should only cause an agent to gain additional information, and cannot cause an agent to *forget* any information that it has previously been told. Thus anything that an agent knows before an informative update, the agent should continue to know after an informative update. This is not quite accurate, as a formula that refers to a lack of knowledge of a particular agent may be invalidated by an informative update, if that agent is provided with information about what it previously did not know. For example, suppose that we have a situation where $\Box_a \neg \Box_b p$, i.e. a knows that b doesn't know that p. Then it is reasonable that an informative update could tell b that p is true, and so after the informative update, $\Box_b p$ becomes true. This causes $\neg \Box_b p$ to become false, and since it is now false, it cannot possibly be known, since only true statements can be known. Thus not all knowledge is necessarily retained by an informative update, only positive knowledge is guaranteed to be retained. The fact that refinements preserve positive knowledge is an important property for refinements in modelling informative updates in an epistemic setting.

We can compare this epistemic interpretation to a doxastic interpretation. Informative updates for beliefs may carry different connotations than informative updates for knowledge; specifically, in some cases beliefs may be revised. For example, if we have that $\Box_a p$, i.e. a believes that p is true, then depending on how you choose to interpret informative updates for belief, it may be reasonable for an informative update to cause a to reconsider its beliefs. For example, a may find evidence that casts doubt on the truth of p, leading to a situation where a no longer believes that p is true, but considers both p and $\neg p$ to be possible, i.e. $\Diamond_a p \wedge \Diamond_a \neg p$. In another example, a may find compelling evidence that $\neg p$ is actually the case, leading to a situation where $\Box_a \neg p$. In both of these situations, we initially have a positive formula $\Box_a p$, but the truth of this formula is not preserved in an informative update that allows beliefs to be revised.

This thesis focuses on quantification over the refinements of Kripke models, and thus the informative updates that these refinements correspond to necessarily preserve the truth of positive formulae. This prohibits the revision of knowledge or beliefs; once an agent knows or believes a positive formula to be true, it is committed to that knowledge or belief forever. In the epistemic setting this is perfectly intuitive, but in the doxastic setting this is not always how informative updates are modelled. We note how this property is similar to the properties of executing action models in a doxastic setting; as we remarked in Section 2.2.2, action models are not capable of revising beliefs. In fact, as we will see next, refinements and action models are very closely related.

Proposition 3.2.4. *On finite epistemic models, every finite refinement is equivalent to the execution of an action model.*

This is shown by van Ditmarsch and French [22].

The property that every finite refinements is equivalent to the execution of an action model is the main justification for refinements as representations of informative updates. Although we do not explicitly define action models in this thesis, it should be noted that action models are inherently finite structures, whereas the refinements of even a finite model can potentially be infinite, and therefore the above property does not hold for the infinite refinements of epistemic models. However van Ditmarsch and French [22] show that if the refinement quantified epistemic logic is extended with the operator for reasoning about the results of specific action models, then the resulting logic is equivalent to the arbitrary action model logic. Hence quantifying over the refinements of a model is equivalent to quantifying over the informative updates.

3.3 Refinement quantified modal logics

In this section we introduce the syntax and semantics of the refinement quantified modal logics. The definitions that we give are general, in the sense that the refinement quantified modal, doxastic and epistemic logics are all instances of the definition that we give. We also provide some examples and basic results that have previously been given by van Ditmarsch, French and Pinchinat [23].

We begin with a definition of the language, $\mathcal{L}_{\triangleright}$, of the refinement quantified modal logics, along with a definition of its semantics over a general class of frames.

Definition 3.3.1 (Language of $\mathcal{L}_{\triangleright}$)**.** Given a finite set of agents A and a set of propositional atoms P, the language of $\mathcal{L}_{\triangleright}$ is defined by the following abstract syntax:

$$\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \Box_a\alpha \mid \blacktriangleright_a\alpha$$

where $p \in P$, $a \in A$ and $\alpha \in \mathcal{L}_{\triangleright}$.

We use the standard abbreviations from basic modal logic, along with an abbreviation for the dual of the $\blacktriangleright_a$ operator, $\triangleright_a\phi ::= \neg\blacktriangleright_a\neg\phi$. For $\blacktriangleright_a\phi$ we say that for every a-refinement, ϕ holds. For $\triangleright_a\phi$ we say that there is an a-refinement such that ϕ holds.

Definition 3.3.2 (Semantics of $\mathcal{L}_{\triangleright}^{\mathbf{C}}$)**.** Let C be a class of Kripke models, and let $M = (S, R, V) \in C$ be a Kripke model taken from C. The interpretation of ϕ is defined inductively.

$$\begin{array}{lcl}
M_s & \models & p \text{ iff } s \in V_p \\
M_s & \models & \neg\phi \text{ iff } M_s \not\models \phi \\
M_s & \models & \phi \wedge \psi \text{ iff } M_s \models \phi \text{ and } M_s \models \psi \\
M_s & \models & \Box_a\phi \text{ if for all } t \in S : (s,t) \in R_a \text{ implies } M_t \models \phi \\
M_s & \models & \blacktriangleright_a\phi \text{ iff for all } M'_{s'} \in C : M_s \succeq_a M'_{s'} \text{ implies } M'_{s'} \models \phi
\end{array}$$

We use the same terminology and notation for satisfiability and validity as is used for basic modal logics. As we sometimes discuss basic modal logics and refinement quantified modal logics at the same time, we may add a subscript to the turnstile operator, e.g. $\models_{\triangleright} \phi$, to make it explicit that we are working in the refinement quantified version of the logic.

The logics $\mathcal{L}_{\triangleright}^{\mathbf{K}}$, $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$ and $\mathcal{L}_{\triangleright}^{\mathbf{S5}}$ are instances of $\mathcal{L}_{\triangleright}^{\mathbf{C}}$, with classes K, $KD45$ and $S5$ respectively.

It should be emphasised that the interpretation of the refinement quantifier, $\blacktriangleright_a$, varies for each logic, as the refinements considered in the interpretation of each logic must be taken from the appropriate class of Kripke models. It is for this reason that $\mathcal{L}_{\triangleright}^{\mathbf{S5}}$ and $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$ are not conservative extensions of $\mathcal{L}_{\triangleright}^{\mathbf{K}}$. For example, $\triangleright_a\Box_a\bot$ is valid in $\mathcal{L}_{\triangleright}^{\mathbf{K}}$, but not in $\mathcal{L}_{\triangleright}^{\mathbf{S5}}$ or $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$. This is because given any pointed model in K, one can construct an a-refinement from that model by deleting the

a-edges starting at the designated state; in this resulting a-refinement, $\Box_a\bot$ is satisfied, and hence $\triangleright_a\Box_a\bot$ is satisfied in the original model. However because of the serial property of *S5* and *KD45* models, $\Box_a\bot$ is not even satisfiable in $\mathcal{L}_{\triangleright}^{\mathbf{S5}}$ or $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$, and hence $\triangleright_a\Box_a\bot$ is not satisfiable either, as the refinement required must be taken from the classes of *S5* or *KD45*.

We also note that for a class of frames C, the refinement quantified modal logic $\mathcal{L}_{\triangleright}^{\mathbf{C}}$ is a conservative extension of the basic modal logic $\mathcal{L}^{\mathbf{C}}$. This is because the interpretation of any formula that does not contain a $\blacktriangleright_a$ operator is the same as the interpretation of that formula in basic modal logic.

Lemma 3.3.1. *The logic $\mathcal{L}_{\triangleright}^{\mathbf{K}}$ is bisimulation invariant.*

This is proven by van Ditmarsch, French and Pinchinat [23].

Lemma 3.3.2. *The logics $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$ and $\mathcal{L}_{\triangleright}^{\mathbf{S5}}$ are bisimulation invariant.*

The proof for bisimulation invariance in $\mathcal{L}_{\triangleright}^{\mathbf{K}}$, given by van Ditmarsch, French and Pinchinat [23] applies to $\mathcal{L}_{\triangleright}^{\mathbf{S5}}$ and $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$.

Example 3.3.1. We recall the coin-flipping example previously given in Section 2.1.2. Suppose that Alice has flipped a coin, and initially Alice and Bob do not know whether the coin landed on heads or tails, and that both Alice and Bob are aware of each other's ignorance of the result. We represent Alice by the symbol a, Bob by the symbol b, the result that the coin landed on heads by the propositional atom p, and the result that the coin landed on tails by the negated atom, $\neg p$. This situation can be represented by the epistemic model in Figure 3.1.

Suppose that the coin actually landed on heads. Then the world where p is true is the actual world. We ask whether it is possible for Alice to learn that the coin landed on heads whilst Bob continues to be ignorant of the result. We can represent this question by the following statement in refinement quantified epistemic logic:

$$\triangleright_a(\Box_a p \wedge \neg\Box_b p)$$

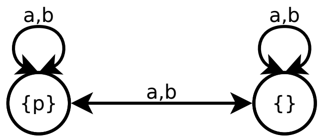

Figure 3.1: Initially Alice and Bob cannot distinguish between the worlds where the coin lands on heads or on tails.

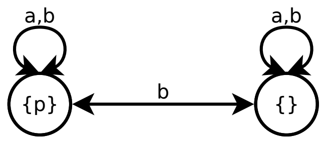

Figure 3.2: After Alice looks at the coin, she can distinguish between the worlds where the coin lands on heads or tails, but Bob still cannot.

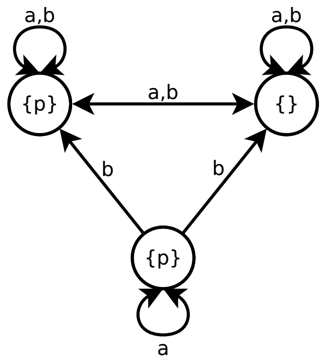

Figure 3.3: After Alice looks at the coin, she can distinguish between the worlds where the coin lands on heads or tails, but Bob still cannot distinguish between these worlds, and neither does he know that Alice can. The actual world is the world on the bottom.

We can show that this statement is satisfied in the actual world of the above epistemic model, by giving the a-refinement in Figure 3.2

Example 3.3.2. We recall the coin-flipping example again, but this time in the setting of doxastic logic. We ask whether it is possible for Alice to come to believe that the coin landed on heads whilst Bob continues to believe that Alice is ignorant of the result. We can represent this question by the following statement in refinement quantified doxastic logic:

$$\triangleright_a(\Box_a p \wedge \Box_b(\neg\Box_a p \wedge \neg\Box_a \neg p))$$

We can show that this statement is satisfied by the model given in Figure 3.1 by giving the a-refinement of the model in Figure 3.3.

We note that the statement is not satisfied in the setting of epistemic logic. This should be clear, as in an epistemic setting, if Alice knows that the coin landed on heads, then Bob cannot know that this isn't the case, because in an epistemic setting agents can only know statements that are actually true.

We briefly list some properties of the refinement quantified modal logic, to give some intuition about the logic.

Proposition 3.3.3. *We have the following validities:*

1. $\models_{\mathcal{L}_\triangleright^{\mathbf{K}}} \blacktriangleright_a(\phi \to \psi) \to \blacktriangleright_a\phi \to \blacktriangleright_a\psi$
2. $\models_{\mathcal{L}_\triangleright^{\mathbf{K}}} \blacktriangleright_a\phi \to \phi$
3. $\models_{\mathcal{L}_\triangleright^{\mathbf{K}}} \blacktriangleright_a\phi \to \blacktriangleright_a\blacktriangleright_a\phi$
4. $\models_{\mathcal{L}_\triangleright^{\mathbf{K}}} \phi$ *implies* $\models_{\mathcal{L}_\triangleright^{\mathbf{K}}} \blacktriangleright_a\phi$
5. $\models_{\mathcal{L}_\triangleright^{\mathbf{K}}} \Box_a\blacktriangleright_a\phi \to \blacktriangleright_a\Box_a\phi$

These properties were shown by van Ditmarsch and French [22]. We note that these properties are also valid in $\mathcal{L}_\triangleright^{\mathbf{S5}}$ and $\mathcal{L}_\triangleright^{\mathbf{KD45}}$.

In previous work, van Ditmarsch, French and Pinchinat [23] have considered the refinement quantified modal logic, $\mathcal{L}_\triangleright^{\mathbf{K}}$, based in the class of models K. They provided a sound and complete axiomatisation for the single-agent case, showed that it was expressively equivalent to the single-agent modal logic, gave a decision procedure for the single-agent case that runs in 2EXP time, and showed that the refinement quantified modal logic is exponentially more succinct than the refinement quantified modal logic.

In this thesis we move towards other variants of refinement quantified modal logics, including the single-agent refinement quantified epistemic and doxastic logics, and the multi-agent refinement quantified modal and doxastic logics. We provide sound and complete axiomatisations for each of these logics, and show that they are expressively equivalent to the modal logics that they are based on. The expressivity result in particular allows us to show several results for each of these logics, particularly that they are all decidable.

CHAPTER 4

Single-agent refinement quantified epistemic and doxastic logics

In this chapter we will provide a sound and complete axiomatisation of the single-agent refinement quantified epistemic and doxastic logics. Our efforts focus mainly on the epistemic variant; as we will see, the axiomatisation for the doxastic variant is very similar, and can be derived using similar reasoning. We do not go into much detail for the doxastic variant, as in Chapter 6 we provide an axiomatisation for the multi-agent doxastic variant, a much more general result.

When discussing the single-agent variants of these logics, we will use a subscript (1) to denote the single-agent logics, e.g. $\mathcal{L}^{\mathbf{S5}}_{(1)}$, $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$, and so on. We also drop the superfluous subscripts denoting agents in the syntax of the logics, e.g. we write $\Box$ instead of $\Box_a$.

4.1 Technical preliminaries

In previous work, van Ditmarsch, French and Pinchinat [23] gave an axiomatisation of the single-agent refinement quantified modal logic. The axiomatisation was expressed in terms of the cover operator, ∇, and the proof of completeness consisted of a translation from the refinement quantified logic to the basic modal logic, a translation that relied on a disjunctive normal form for modal formulae, defined in terms of the cover operator. The cover operator allows us to collect all of the modalities that are true at a particular state, and consider them all at once. This simplified the soundness and completeness proofs of the modal refinement quantified logic. As van Ditmarsch, French and Pinchinat [23] remarked, a proof directly in the setting of basic modal logic, without the cover operator, is possible, but considerably longer.

Our axiomatisations of the single-agent epistemic and doxastic variants also rely on the cover operator, and on a normal form similar to the disjunctive normal form used by van Ditmarsch, French and Pinchinat [23]. The normal form that we introduce is called the cover logic prenex normal form, which is similar to the disjunctive normal form, except that it prohibits nested modalities; thus modalities may only be applied to propositional formulae. As we will see, every single-agent epistemic and doxastic formula is equivalent to a formula in this prenex normal form. The prenex normal form allows us to avoid complications that would arise due to the transitive and symmetric properties of *S5*, or the transitive and Euclidean properties of *KD45*, that would arise when nested modalities are present.

We define the basic prenex normal form first, and then give the definition of the cover logic version. We show that every single-agent epistemic or doxastic formula is equivalent to a formula in these normal forms.

Definition 4.1.1 (Prenex normal form)**.** A formula in prenex normal form is specified by the following abstract syntax, where α is a formula in prenex normal form:

$$\begin{array}{lcl} \alpha & ::= & \delta \mid \alpha \vee \alpha \\ \delta & ::= & \pi \mid \Box\pi \mid \Diamond\pi \mid \delta \wedge \delta \end{array}$$

where π stands for a propositional formula.

Lemma 4.1.1. *Every formula of $\mathcal{L}_{(1)}$ is equivalent to a formula in prenex normal form, under the semantics of $\mathcal{L}^{\mathbf{S5}}_{(1)}$.*

This is shown by Meyer and van der Hoek [19].

Lemma 4.1.2. *Every formula of $\mathcal{L}_{(1)}$ is equivalent to a formula in prenex normal form, under the semantics of $\mathcal{L}^{\mathbf{KD45}}_{(1)}$.*

Proof. The proof given by Meyer and van der Hoek [19] for Lemma 4.1.1 applies also to $\mathcal{L}^{\mathbf{KD45}}_{(1)}$.

Meyer and van der Hoek remarked that the only use of the reflexivity axiom of $\mathcal{L}^{\mathbf{S5}}$, $\mathbf{T}$, in the proof, is in the form of the theorems $\vdash \Box\Box\phi \to \Box\phi$, and $\vdash \Box\neg\Box\phi \to \neg\Box\phi$. Therefore the proof holds for any logic which replaces $\mathbf{T}$ with axioms entailing both of these properties. Both of these properties are obviously valid in $\mathcal{L}^{\mathbf{KD45}}_{(1)}$, and therefore the proof of Lemma 4.1.1 applies to this result. $\Box$

We note that, similar to a disjunctive normal form for propositional logic, conversion to the prenex normal form can result in a formula that is exponentially larger than the original formula in the worst case.

Given this result, we define our cover logic version of the prenex normal form, and show that every epistemic or doxastic formula is equivalent to a formula in cover logic prenex normal form.

Definition 4.1.2 (Cover logic prenex normal form)**.** A formula in cover logic prenex normal form is specified by the following abstract syntax:

$$\alpha ::= \pi \land \nabla\Gamma \mid \alpha \lor \alpha$$

where π is a propositional formula, and Γ is a set of propositional formulae.

Lemma 4.1.3. *Every formula of $\mathcal{L}_{(1)}$ is equivalent to a formula in cover logic prenex normal form, under both the semantics of $\mathcal{L}^{\mathbf{S5}}_{(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{(1)}$.*

Proof. Without loss of generality, we may assume that our given formula is in prenex normal form (by Lemma 4.1.1 for $\mathcal{L}^{\mathbf{S5}}_{(1)}$, or by Lemma 4.1.2 for $\mathcal{L}^{\mathbf{KD45}}_{(1)}$).

Given a formula in prenex normal form, we consider each disjunct separately. We can convert each term $\Box\gamma$ or $\Diamond\gamma$ into an equivalent term using the cover operator, using the equivalences $\Box\gamma \equiv \nabla\{\gamma\}$ and $\Diamond\gamma \equiv \nabla\{\gamma, \top\}$ Note that each resulting term contains a cover operator applied only to a set of propositional formulae.

An inductive argument can be used to show that we can collapse the resulting conjunction of cover operators into a single term containing one cover operator applied to a set of propositional formulae. We use the following equivalence to achieve this, and note that at each stage this equivalence preserves the property that the cover operator is only applied to a set of propositional formulae.

$$\nabla\Gamma \land \nabla\Gamma' \equiv \nabla\big(\{\gamma \land \bigvee_{\gamma' \in \Gamma'} \gamma' \mid \gamma \in \Gamma\} \cup \{\gamma' \land \bigvee_{\gamma \in \Gamma} \gamma \mid \gamma' \in \Gamma'\}\big)$$

Repeating this for each disjunct in our original formula leaves us with a formula in cover logic prenex normal form. $\square$

4.2 Axiomatisation

We provide an axiomatisation of the single-agent refinement quantified epistemic and doxastic logics, $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$, and prove their soundness and completeness.

As we are defining several similar axiomatisations at once, we will define their common elements as $\mathbf{RML}_{(1)}$, and define the axiomatisations of $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$ and $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$, which we call $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$ and $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ respectively, as extensions of $\mathbf{RML}_{(1)}$. We will also define the axiomatisation, $\mathbf{RML}^{\mathbf{K}}_{(1)}$ of single-agent refinement quantified modal logic, so that we can compare our axiomatisations.

Definition 4.2.1. The axiomatisation $\mathbf{RML}_{(1)}$ is a substitution schema of the following axioms:

P	All tautologies of propositional logic.
K	$\Box(\phi \to \psi) \to \Box\phi \to \Box\psi$
R	$\blacktriangleright(\phi \to \psi) \to \blacktriangleright\phi \to \blacktriangleright\psi$
RP	$\blacktriangleright\alpha \leftrightarrow \alpha$, where α is a propositional formula.

Along with the rules:

MP	From $\vdash \phi \to \psi$ and $\vdash \phi$ infer $\vdash \psi$.
NecK	From $\vdash \phi$ infer $\vdash \Box\phi$.
NecR	From $\vdash \phi$ infer $\vdash \blacktriangleright\phi$.

The axioms and rules in $\mathbf{RML}_{(1)}$ were previously presented by van Ditmarsch, French and Pinchinat as axioms of $\mathbf{RML}^{\mathbf{K}}_{(1)}$, which we will define now.

Definition 4.2.2. The axiomatisation $\mathbf{RML}^{\mathbf{K}}_{(1)}$ is a substitution schema consisting of the axioms and rules of $\mathbf{RML}_{(1)}$, along with the additional axiom:

RK $\quad \triangleright\nabla\Gamma \leftrightarrow \bigwedge_{\gamma\in\Gamma} \Diamond\triangleright\gamma$

The axiomatisation $\mathbf{RML}^{\mathbf{K}}_{(1)}$ shares axioms and rules from the axiomatisation for basic modal logic, $\mathbf{K}$.

Definition 4.2.3. The axiomatisation $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ is a substitution schema consisting of the axioms and rules of $\mathbf{RML}_{(1)}$, along with the additional axioms:

T	$\Box\phi \to \phi$
5	$\Diamond\phi \to \Box\Diamond\phi$
RS5	$\triangleright\nabla\Gamma \leftrightarrow \left(\bigvee_{\gamma\in\Gamma} \gamma\right) \wedge \left(\bigwedge_{\gamma\in\Gamma} \Diamond\gamma\right)$

where Γ is a set of propositional formulae.

Definition 4.2.4. The axiomatisation $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$ is a substitution schema consisting of the axioms and rules of $\mathbf{RML}_{(1)}$, along with the additional axioms:

D	$\Box\phi \to \Diamond\phi$
4	$\Box\phi \to \Box\Box\phi$
5	$\Diamond\phi \to \Box\Diamond\phi$
RKD45	$\triangleright\nabla\Gamma \leftrightarrow \bigwedge_{\gamma\in\Gamma} \Diamond\gamma$

where Γ is a set of propositional formulae.

We note that many of the axioms from $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ and $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$ are also axioms from **S5**, **KD45** and $\mathbf{RML}^{\mathbf{K}}_{(1)}$. The differences between these two axiomatisations, and $\mathbf{RML}^{\mathbf{K}}_{(1)}$, is that $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ and $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$ include the additional **S5** and **KD45** axioms respectively, and include a different axiom in place of **RK**. The axioms **RS5** and **RKD45** are similar in form to **RK**, with the exception that the sets Γ may only contain propositional formulae in the case of **RS5** and **RKD45**, and that the axiom **RS5** contains an additional disjunction on the right hand side of the equivalence. We note that the disjunction corresponds to an application of the **T** axiom, but in cover logic notation; it ensures that the reflexive state may be considered as one of the successors considered by the $\Diamond\gamma$ subformulae.

Example 4.2.1. We give an example derivation using $\mathbf{RML}^{\mathbf{S5}}_{(1)}$, showing that $\vdash_{\mathbf{RML}^{\mathbf{S5}}_{(1)}} p \to \triangleright\Box p$.

$$
\begin{array}{ll}
\vdash \Box\neg p \to \neg p & (\mathbf{T}) \\
\vdash p \to p \wedge \neg\Box\neg p & (\mathbf{P}) \\
\vdash p \to p \wedge \Diamond p & (\text{Definition of } \Diamond) \\
\vdash p \to p \wedge \Diamond\neg\neg p & (\mathbf{P}) \\
\vdash p \to p \wedge \Diamond\neg\blacktriangleright\neg p & (\mathbf{RP}) \\
\vdash p \to p \wedge \Diamond\triangleright p & (\text{Definition of } \triangleright) \\
\vdash p \to \triangleright\nabla\{p\} & (\mathbf{RS5}) \\
\vdash p \to \triangleright\Box p & (\text{Definition of } \nabla)
\end{array}
$$

Example 4.2.2. We give an example derivation using $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$, showing that $\vdash_{\mathbf{RML}^{\mathbf{KD45}}_{(1)}} \Diamond p \to \triangleright\Box p$.

$$
\begin{array}{ll}
\vdash \Diamond p \to \Diamond p & (\mathbf{P}) \\
\vdash \Diamond p \to \Diamond\neg\neg p & (\mathbf{P}) \\
\vdash \Diamond p \to \Diamond\neg\blacktriangleright\neg p & (\mathbf{RP}) \\
\vdash \Diamond p \to \Diamond\triangleright p & (\text{Definition of } \triangleright) \\
\vdash \Diamond p \to \triangleright\nabla\{p\} & (\mathbf{RKD45}) \\
\vdash \Diamond p \to \triangleright\Box p & (\text{Definition of } \nabla)
\end{array}
$$

We will now show the soundness of the axiomatisations $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ and $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$.

Lemma 4.2.1. *The axiomatisation* $\mathbf{RML}_{(1)}$ *is sound with respect to the semantic classes of* $K_{(1)}$*,* $S5_{(1)}$ *and* $KD45_{(1)}$*.*

Proof. The soundness of the axioms **P** and **K**, and the rules **MP** and **NecK** can be shown by the same reasoning used to show that they are sound in basic modal logic. The axioms **R** and **RP**, and the rule **NecR** can be shown to be sound using the same reasoning applied in the single-agent refinement quantified modal logic [23]. □

Lemma 4.2.2. *The axiomatisation* $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ *is sound with respect to the semantic class* $S5_{(1)}$*.*

Proof. Soundness of the axioms **P**, **K**, **R**, and **RP**, and the rules **MP**, **NecK** and **NecR** are shown above. Soundness of the axioms **T** and **5** can be shown by the same reasoning used to show that they are sound in basic epistemic logic.

All that is to be shown is the soundness of **RS5**.

Let Γ be a finite set of propositional formulae, and let M_s be a model in *S5* such that $M_s \models \bigvee_{\gamma\in\Gamma} \gamma \wedge \bigwedge_{\gamma\in\Gamma} \Diamond\gamma$.

We need to show that $M_s \models \triangleright\nabla\Gamma$. To do this we will construct a model $N_s \in S5$, construct a simulation from N_s to M_s to show that $N_s \preceq M_s$, and finally show that $N_s \models \nabla\Gamma$.

We begin by constructing the model N_s. We note that for each $\gamma \in \Gamma$, there is some successor $s^\gamma \in sR^M$ such that $M_{s^\gamma} \models \gamma$. We can construct the model N such that $S^N = \{s\} \cup \{s^\gamma \mid \gamma \in \Gamma\}$, $R^N = S^N \times S^N$ and for all $p \in P$, $V^N(p) = V^M(p) \cap S^N$. This model is clearly in *S5*.

Furthermore we have that $M'_s \preceq M_s$ by the relation $\mathcal{R} = \{(s, s)\} \cup \{(t^\gamma, t^\gamma) \mid \gamma \in \Gamma\}$ (**atoms** and **forth** are satisfied).

By construction, for each $\gamma \in \Gamma$, there is a successor $s^\gamma \in sR^N$ such that $N_{s^\gamma} \models \gamma$. For each successor $s \in sR^N$ we have that $N_s \models \bigvee_{\gamma\in\Gamma} \gamma$, as each successor is either one of the s^γ for some $\gamma \in \Gamma$, in which case $N_{s^\gamma} \models \gamma$, or it is our initial state s, in which case $N_s \models \bigvee_{\gamma\in\Gamma} \gamma$ follows from our hypothesis that $M_s \models \bigvee_{\gamma\in\Gamma} \gamma$. Therefore $N_s \models \nabla\Gamma$.

Therefore $M_s \models \triangleright\nabla\Gamma$.

Conversely, let Γ be a finite set of propositional formulae, and let M_s be a model in *S5* such that $M_s \models \triangleright\nabla\Gamma$. Then there exists some model $N_t \preceq M_s$ in *S5*, via some simulation $\mathcal{R} \subseteq S' \times S$, such that $N_t \models \nabla\Gamma$.

From the definition of the cover operator, $N_t \models \Box \bigvee_{\gamma\in\Gamma} \gamma \wedge \bigwedge_{\gamma\in\Gamma} \Diamond\gamma$.

As $N \in S5$, we know that $t \in tR^N$, and so it follows from $N_t \vDash \Box \bigvee_{\gamma\in\Gamma} \gamma$ that $N_t \vDash \bigvee_{\gamma\in\Gamma} \gamma$. As we know that $(t, s) \in \mathcal{R}$, from **atoms** we know that M_s and N_t are equivalent for propositional formulae. As each $\gamma \in \Gamma$ is propositional, it follows that $M_s \vDash \bigvee_{\gamma\in\Gamma} \gamma$.

Furthermore, from $N_t \vDash \bigwedge_{\gamma\in\Gamma} \Diamond\gamma$, we know that for every $\gamma \in \Gamma$, there exists some $t^\gamma \in tR^{M'}$ such that $N_{t^\gamma} \vDash \gamma$. It then follows from **forth** that there exists some $s^\gamma \in S^M$ such that $s^\gamma \in sR^M$ and $(t^\gamma, s^\gamma) \in \mathcal{R}$. From **atoms** we know that M_{s^γ} and N_{t^γ} are equivalent for propositional formulae. As γ is propositional, it follows that $M_{s^\gamma} \vDash \gamma$ and therefore $M_s \vDash \Diamond\gamma$ for each $\gamma \in \Gamma$. Therefore $M_s \vDash \bigwedge_{\gamma\in\Gamma} \Diamond\gamma$, and so $M_s \vDash \bigvee_{\gamma\in\Gamma} \gamma \wedge \bigwedge_{\gamma\in\Gamma} \Diamond\gamma$.

Therefore **RS5** is sound, and so $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ is sound for the logic $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$. □

Lemma 4.2.3. *The axiomatisation* $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$ *is sound with respect to the semantic class* $KD45_{(1)}$.

Proof. The proof is similar to the proof for Lemma 4.2.2. Instead of showing soundness for the $\mathbf{S5}_{(1)}$ axioms, we must show that the $\mathbf{KD45}_{(1)}$ axioms are sound, and this follows from their soundness in $\mathcal{L}^{\mathbf{KD45}}_{(1)}$. The main difference in the proof of soundness is for **RKD45** as compared to the proof for **RS5**, is that in the right to left direction, we do not have to show that $M_s \vDash \bigvee_{\gamma\in\Gamma} \gamma$; as doxastic models do not have to be reflexive, there is no requirement for s to be in the possible worlds for the constructed refinement. For the left to right direction of the proof, the refinement N_t is a KD45 model instead of an S5 model, but this has no bearing on the rest of the proof. □

We show the completeness of the axiomatisations $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ and $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$ by provably correct translations from $\mathcal{L}_{\triangleright(1)}$ to $\mathcal{L}_{(1)}$, using the axioms and rules of $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ and $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$. A provably correct translation uses the axioms of our respective axiomatisations in order to translate our given formula to an equivalent formula in a different form. In our case, we are translating our formulae into a form that does not include any refinement quantifiers (i.e. $\triangleright$-free formulae). As the refinement quantified versions of our logics are conservative extensions of the respective basic modal logics, the interpretation of a $\triangleright$-free formula is equivalent to the interpretation of that same formula in the corresponding basic modal logic. Therefore we can prove theorems in the refinement quantified modal logic by translating the formula to an equivalent $\triangleright$-free form and then proving that theorem in basic modal logic. Given this, the completeness of our axiomatisations then follows from the completeness of the respective basic modal logics.

This is the same strategy used by van Ditmarsch, French and Pinchinat [23] to show the completeness of the single-agent refinement quantified modal logic, and is the same general strategy that we will use for the other logics that we consider in this paper.

We will begin by providing a provably correct translation from $\mathcal{L}_{\triangleright(1)}$ to $\mathcal{L}_{(1)}$, and follow with some results that show that completeness follows from this translation.

Lemma 4.2.4. *Every formula of $\mathcal{L}_{\triangleright(1)}$ is provably equivalent to a formula of $\mathcal{L}_{(1)}$ with the axiomatisation $\mathbf{RML}^{\mathbf{S5}}_{(1)}$.*

Proof. Let $\alpha \in \mathcal{L}_{\triangleright(1)}$. We assume without loss of generality that all $\blacktriangleright$ operators are expressed as $\triangleright$ operators, by the equivalence $\blacktriangleright\phi \leftrightarrow \neg\triangleright\neg\phi$. We prove by induction on the number of occurrences of $\triangleright$ in α that α is equivalent to a $\triangleright$-free formula, and therefore to a formula in $\mathcal{L}_{(1)}$. The base case where α contains no $\triangleright$ is trivial, as a $\triangleright$-free formula is a formula in $\mathcal{L}_{(1)}$. Suppose instead that α contains $n+1$ $\triangleright$ operators, and assume that any formula with n $\triangleright$ operators is provably equivalent to a formula in $\mathcal{L}_{(1)}$. We use the axioms of $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ to show that α is provably equivalent to a formula with n $\triangleright$ operators, and that therefore by the induction hypothesis it is provably equivalent to a formula in $\mathcal{L}_{(1)}$.

Given any subformula from α of type $\triangleright\beta$, such that β is a $\triangleright$-free formula. Without loss of generality, by Lemma 4.1.3 we may assume that β is in disjunctive normal form. We prove by induction on the structure of β that $\triangleright\beta$ is provably equivalent to a formula $\chi \in \mathcal{L}_{(1)}$. The induction hypothesis is that for any proper subformula ϕ of β that $\triangleright\phi$ is equivalent to a formula in $\mathcal{L}_{(1)}$.

The base case is when β is a propositional formula. In this case, from **P** and **RP**, we have that $\triangleright\beta \leftrightarrow \beta$, and therefore $\triangleright\beta$ is equivalent to a formula in $\mathcal{L}_{(1)}$.

The inductive case is when $\beta = \phi \vee \psi$, or when $\beta = \pi \wedge \nabla\Gamma$.

Suppose that $\beta = \phi \vee \psi$. Then $\triangleright(\phi \vee \psi) \leftrightarrow \triangleright\phi \vee \triangleright\psi$ is derivable from **P** and **R**. By the induction hypothesis, $\triangleright\phi$ and $\triangleright\psi$ are equivalent to some $\phi', \psi' \in \mathcal{L}_{(1)}$. Therefore $\triangleright\beta \leftrightarrow \phi' \vee \psi'$, and so $\triangleright\beta$ is equivalent to a formula in $\mathcal{L}_{(1)}$.

Suppose that $\beta = \pi \wedge \nabla\Gamma$ where Γ is a set of propositional formulae. Then $\triangleright(\pi \wedge \nabla\Gamma) \leftrightarrow \pi \wedge \triangleright\nabla\Gamma$ is derivable from **P**, **R** and **RP**. Moreover, $\pi \wedge \triangleright\nabla\Gamma \leftrightarrow \pi \wedge \bigvee_{\gamma\in\Gamma} \gamma \wedge \bigwedge_{\gamma\in\Gamma} \Diamond\gamma$ is derivable from **RS5**. As each $\gamma \in \Gamma$ is propositional, $\triangleright\beta$ is equivalent to a formula in $\mathcal{L}_{(1)}$.

Therefore by the induction, $\triangleright\beta$ is equivalent to a formula $\chi \in \mathcal{L}_{(1)}$ for every $\beta \in \mathcal{L}_{\triangleright(1)}$.

Hence replacing $\triangleright\beta$ in α with χ gives an equivalent formula that contains only n $\triangleright$ operators.

Therefore by the induction, α is equivalent to a formula in $\mathcal{L}_{(1)}$. □

Lemma 4.2.5. *Every formula of $\mathcal{L}_{\triangleright(1)}$ is provably equivalent to a formula of $\mathcal{L}_{(1)}$ with the axiomatisation $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$.*

The proof is similar to the proof for Lemma 4.2.4, with the only difference being that the axiom **RKD45** is used in place of **RS5** in the induction over the structure of β.

The rest of the completeness proof is to show that, given the above translations into $\mathcal{L}_{(1)}$, we can show completeness by using these translations along with the completeness of $\mathcal{L}^{\mathbf{S5}}_{(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{(1)}$.

Corollary. *Let $\phi \in \mathcal{L}_{\triangleright(1)}$ be given and $\psi \in \mathcal{L}_{(1)}$ be equivalent to ϕ under the semantics of $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$. If ψ is a theorem in $\mathcal{L}^{\mathbf{S5}}_{(1)}$, then ϕ is a theorem in $\mathbf{RML}^{\mathbf{S5}}_{(1)}$.*

Proof. Let $\phi \in \mathcal{L}_{\triangleright(1)}$ and let $\psi \in \mathcal{L}_{(1)}$ be semantically equivalent to ϕ. By Lemma 4.2.4, we can obtain some $\phi' \in \mathcal{L}_{(1)}$ that is semantically equivalent to ϕ (and thus also to ψ) by following the given translation steps. We can extend a derivation of ψ to a derivation of ϕ' as the two are semantically equivalent under $\mathcal{L}^{\mathbf{S5}}_{(1)}$, and by the completeness of $\mathcal{L}^{\mathbf{S5}}_{(1)}$ this equivalence is derivable. As $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ is a conservative extension of $\mathcal{L}^{\mathbf{S5}}_{(1)}$, this equivalence is therefore also derivable in $\mathbf{RML}^{\mathbf{S5}}_{(1)}$. The derivation can be further extended to ϕ by observing that all of the reduction steps in Lemma 4.2.4 are provable equivalences in $\mathbf{RML}^{\mathbf{S5}}_{(1)}$. Therefore ϕ is a theorem in $\mathbf{RML}^{\mathbf{S5}}_{(1)}$. □

Lemma 4.2.6. *The axiom schema $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ is complete for the logic $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$.*

Proof. Let $\phi \in \mathcal{L}_{\triangleright(1)}$ such that $\models_{\triangleright} \phi$. Then by Lemma 4.2.4, there exists a semantically equivalent formula $\psi \in \mathcal{L}_{(1)}$ which is $\triangleright$-free. As $\models_{\triangleright} \phi$ and $\phi \leftrightarrow \psi$, then $\models_{\triangleright} \psi$. As ψ is $\triangleright$-free, then it follows that $\models \psi$, and by the completeness of $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ it follows that $\vdash \psi$. Therefore by Corollary 4.2 we have that $\vdash_{\triangleright} \phi$. □

Theorem 4.2.7. *The axiomatisation $\mathbf{RML}^{\mathbf{S5}}_{(1)}$ is sound and complete with respect to the semantic class $S5_{(1)}$.*

Proof. The soundness proof is given in Lemma 4.2.2 and the completeness proof is given in Lemma 4.2.6. □

Theorem 4.2.8. *The axiomatisation $\mathbf{RML}^{\mathbf{KD45}}_{(1)}$ is sound and complete with respect to the semantic class $KD45_{(1)}$.*

Proof. The soundness proof is given in Lemma 4.2.3, and we note that similar results to Corollary 4.2 and Lemma 4.2.6 can be shown with minor modifications to their proofs, which gives us completeness. □

The completeness proofs above were performed with a provably correct translation from $\mathcal{L}_{\triangleright(1)}$ to $\mathcal{L}_{(1)}$, under the semantics of $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$. This shows that $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$ are expressively equivalent to $\mathcal{L}^{\mathbf{S5}}_{(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{(1)}$ respectively. This allows us to show that $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$ share properties of $\mathcal{L}^{\mathbf{S5}}_{(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{(1)}$, via this property. In particular, we will show that $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$ are decidable; that is, that there is a decision procedure for determining whether any formula in $\mathcal{L}_{\triangleright(1)}$ is valid in $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ or $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$.

Theorem 4.2.9. *The logics $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ and $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$ are decidable.*

Proof. Given a formula ϕ in $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$, we can find an equivalent ψ in $\mathcal{L}^{\mathbf{S5}}_{(1)}$ (from Lemma 4.2.4). We can therefore determine whether ψ is satisfiable using a decision procedure designed for $\mathcal{L}^{\mathbf{S5}}_{(1)}$. The decidability for $\mathcal{L}^{\mathbf{S5}}_{\triangleright(1)}$ therefore follows from the decidability of $\mathcal{L}^{\mathbf{S5}}_{(1)}$ [7].

The proof for $\mathcal{L}^{\mathbf{KD45}}_{\triangleright(1)}$ is the same, but relying on Lemma 4.2.5 for the translation, and on the decidability of $\mathcal{L}^{\mathbf{KD45}}_{(1)}$ [7]. $\square$

We note that other results from single-agent epistemic and doxastic logics can be shown in the setting of the refinement quantified versions by using similar reasoning.

CHAPTER 5

Multi-agent refinement quantified modal logic

In this chapter we will provide a sound and complete axiomatisation of the multi-agent refinement quantified modal logic, over the class of K.

5.1 Technical preliminaries

The axiomatisation of the single-agent refinement quantified modal logic given by van Ditmarsch, French and Pinchinat [23] relied on a disjunctive normal form defined in terms of the cover operator, ∇. We introduce the same disjunctive normal form for our result in multi-agent refinement quantified modal logic, as we will be using a similar provably correct translation in order to show the completeness of our axiomatisation.

Definition 5.1.1 (Disjunctive normal form)**.** A formula in disjunctive normal form is defined by the following abstract syntax:

$$\alpha ::= \pi \wedge \bigwedge_{a \in B} \nabla_a \Gamma_a \mid \alpha \vee \alpha$$

where π stands for a propositional formula, $B \subseteq A$, and for $a \in B$, Γ_a stands for a finite set of formulae in disjunctive normal form.

To show that every $\mathcal{L}$ formula is equivalent to a disjunctive normal formula, we first introduce the negation normal form and a corresponding lemma for that form.

Definition 5.1.2 (Negation normal form)**.** A formula in negation normal form is defined by the following abstract syntax:

$$\alpha ::= p \mid \neg p \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \Box_a \alpha \mid \Diamond_a \alpha$$

where $p \in P$ and $a \in A$.

Lemma 5.1.1. *Every formula of $\mathcal{L}$ is equivalent to a formula in negation normal form, under the semantics of $\mathcal{L}^{\mathbf{K}}$.*

Proof. Similar to negation normal forms in propositional logic, we can recursively push the negations inwards using the following equivalences:

$$\begin{array}{rcl} \neg\neg\phi & \leftrightarrow & \phi \\ \neg(\phi \wedge \psi) & \leftrightarrow & \neg\phi \vee \neg\psi \\ \neg\Box_a\phi & \leftrightarrow & \Diamond_a\neg\phi \end{array}$$

□

Lemma 5.1.2. *Every formula of $\mathcal{L}$ is equivalent to a formula in disjunctive normal form, under the semantics of $\mathcal{L}^{\mathbf{K}}$.*

Proof. Let $\alpha \in \mathcal{L}$. Without loss of generality, by Lemma 5.1.1, we may assume that α is in negation normal form. We prove by induction over the structure of α that α is equivalent to a formula in disjunctive normal form. The induction hypothesis is that every strict subformula of α has an equivalent in disjunctive normal form.

The base case is when $\alpha = p$ or $\alpha = \neg p$ for some $p \in P$, in which case we are done.

Suppose that $\alpha = \phi \vee \psi$. By the induction hypothesis, there are formulae ϕ' and ψ' in disjunctive normal form that are equivalent to ϕ and ψ respectively. Then $\phi \vee \psi \leftrightarrow \phi' \vee \psi'$, which is in disjunctive normal form.

Suppose that $\alpha = \Box_a\phi$. By the induction hypothesis, there is a formula ϕ' in disjunctive normal form that is equivalent to ϕ. Then $\Box_a\phi \leftrightarrow \nabla_a\{\phi\} \vee \nabla_a\emptyset$, which is in disjunctive normal form.

Suppose that $\alpha = \Diamond_a\phi$. By the induction hypothesis, there is a formula ϕ' in disjunctive normal form that is equivalent to ϕ. Then $\Diamond_a\phi \leftrightarrow \nabla_a\{\phi, \top\}$, which is in disjunctive normal form.

Suppose that $\alpha = \phi \wedge \psi$. By the induction hypothesis, there are formulae ϕ' and ψ' in disjunctive normal form that are equivalent to ϕ and ψ respectively. Then $\phi \wedge \psi \leftrightarrow \phi' \wedge \psi'$. As ϕ' and ψ' are in disjunctive normal form, then $\phi' = \delta_1 \vee \cdots \vee \delta_m$ and $\psi' = \gamma_1 \vee \cdots \vee \gamma_m$ for some $m, n \geq 0$, where each of the δ_i and γ_i are terms of the form $\pi \wedge \bigwedge_{a \in B \subseteq A} \nabla_a\Gamma_a$. Then we can rewrite α as a disjunction of conjunctions, by the following equivalence:

$$\phi' \wedge \psi' \leftrightarrow \bigvee_{i \leq m, j \leq n} \delta_i \wedge \gamma_j$$

For each $i \leq m$ and $j \leq n$, we have that $\delta_i = \pi \wedge \bigwedge_{a \in B \subseteq A} \nabla_a \Gamma_a$, and $\gamma_j = \rho \wedge \bigwedge_{a \in C \subseteq A} \nabla_a \Gamma'_a$, where π and ρ are propositional formulae, and each Γ_a and Γ'_a is a set of disjunctive normal formulae. Then we can write each conjunction as:

$$\delta_i \wedge \gamma_i \leftrightarrow (\pi \wedge \rho) \wedge \bigwedge_{a \in B \subseteq A} \nabla_a \Gamma_a \wedge \bigwedge_{a \in C \subseteq A} \nabla_a \Gamma'_a$$

We note that the sets of agents B and C may intersect, and hence the same agent may appear in each of those sets, possibly with different sets of formulae Γ_a and Γ'_a. We can combine the two sets of formulae into one, so that each agent appears only once, using the following equivalence:

$$\nabla_a \Gamma \wedge \nabla_a \Gamma' \equiv \nabla_a\big(\{\gamma \wedge \bigvee_{\gamma' \in \Gamma'} \gamma' \mid \gamma \in \Gamma\} \cup \{\gamma' \wedge \bigvee_{\gamma \in \Gamma} \gamma \mid \gamma' \in \Gamma'\}\big)$$

We note that as each $\gamma \in \Gamma$ and $\gamma' \in \Gamma'$ are assumed to be disjunctive normal formulae, that applying a disjunction over each of these sets yields a disjunctive normal formula. Conjoining two disjunctive normal formulae does not yield a disjunctive normal formula, however an inductive argument can be used to show that recursively applying the same translation described here, to each of these conjunctions, yields a disjunctive normal formula.

Repeating this for each disjunct in our original formula leaves us with a formula in cover logic disjunctive normal form.

Therefore every formula of $\mathcal{L}$ is equivalent to a formula in disjunctive normal form. □

We note that, similar to disjunctive normal forms in propositional logic, and to the prenex normal form introduced for the single-agent doxastic and epistemic logics, conversion into disjunctive normal form in modal logic can result in a formula that is exponentially larger than the original formula in the worst case.

5.2 Axiomatisation

We provide an axiomatisation of the multi-agent refinement quantified modal logic, $\mathcal{L}^{\mathbf{K}}_{\triangleright}$, and prove its soundness and completeness.

Definition 5.2.1 (Axiomatisation $\mathbf{RML^K}$)**.** The axiomatisation $\mathbf{RML^K}$ is a substitution schema consisting of the following axioms:

P	All propositional tautologies
K	$\Box(\phi \to \psi) \to \Box\phi \to \Box\psi$
R	$\blacktriangleright_a(\phi \to \psi) \to \blacktriangleright_a\phi \to \blacktriangleright_a\psi$
RP	$\blacktriangleright_a\alpha \leftrightarrow \alpha$ where α is a propositional formula
RComm	$\triangleright_a\nabla_b\Gamma \leftrightarrow \nabla_b\{\triangleright_a\gamma \mid \gamma \in \Gamma\}$ where $a \neq b$
RDist	$\bigwedge_{b\in B} \triangleright_a\nabla_b\Gamma_b \to \triangleright_a \bigwedge_{b\in B} \nabla_b\Gamma_b$ where $B \subseteq A$
RK	$\triangleright_a\nabla_a\Gamma \leftrightarrow \bigwedge_{\gamma\in\Gamma} \Diamond_a\triangleright_a\gamma$

Along with the rules:

MP	From $\vdash \phi \to \psi$ and $\vdash \phi$, infer $\vdash \psi$
NecK	From $\vdash \phi$ infer $\vdash \Box_a\phi$
NecR	From $\vdash \phi$ infer $\vdash \blacktriangleright_a\phi$

The axiomatisation $\mathbf{RML^K}$ shares many of the axioms and rules of the axiomatisation from the single-agent case. The axioms **P**, **K**, **R**, **RP** and **RK**, and the rules **MP**, **NecK** and **NecR** are essentially the same as the axioms that van Ditmarsch, French and Pinchinat [23] used in the single-agent case. The differences are that $\mathbf{RML^K}$ contains axioms for handling the interaction between multiple agents. The axioms **RComm** and **RDist** are novel axioms used to handle the situation where a refinement quantifier is applied to a cover operator of a different agent, and where a refinement quantifier is applied to a conjunction of cover operators belonging to different agents. We call these axioms **RComm** and **RDist** because they correspond to properties that are visually similar to commutativity or distributivity of the $\triangleright$ opereator over the ∇ operators for different agents.

Example 5.2.1. We give an example derivation using $\mathbf{RML^K}$, based on the coin-flipping example given in Example 3.3.1. We suppose that neither Alice nor Bob know whether the coin has landed on heads or tails, and we show that it is possible for Alice to learn that the coin landed on heads, whilst Bob continues to not know.

Thus we give a derivation of $(\neg\Box_a p \wedge \neg\Box_a\neg p \wedge \neg\Box_b p \wedge \neg\Box_b\neg p) \rightarrow \triangleright_a(\Box_a p \wedge \neg\Box_b p)$. Let $\phi = \neg\Box_a p \wedge \neg\Box_a\neg p \wedge \neg\Box_b p \wedge \neg\Box_b\neg p$.

$\vdash \phi \rightarrow \neg\Box_a\neg p \wedge \neg\Box_b p$	(**P**)
$\vdash \phi \rightarrow \Diamond_a p \wedge \Diamond_b\neg p$	(Definition of $\Diamond$)
$\vdash \phi \rightarrow \Diamond_a p \wedge \nabla_b\{\neg p, \top\}$	(Definition of ∇)
$\vdash \phi \rightarrow \Diamond_a\neg\neg p \wedge \nabla_b\{\neg\neg\neg p, \neg\neg\top\}$	(**P**)
$\vdash \phi \rightarrow \Diamond_a\neg\blacktriangleright_a\neg p \wedge \nabla_b\{\neg\blacktriangleright_a\neg\neg p, \neg\blacktriangleright_a\neg\top\}$	(**RP**)
$\vdash \phi \rightarrow \Diamond_a\triangleright_a p \wedge \nabla_b\{\triangleright_a\neg p, \triangleright_a\top\}$	(Definition of $\triangleright$)
$\vdash \phi \rightarrow \triangleright_a\nabla_a\{p\} \wedge \nabla_b\{\triangleright_a\neg p, \triangleright_a\top\}$	(**RK**)
$\vdash \phi \rightarrow \triangleright_a\nabla_a\{p\} \wedge \triangleright_a\nabla_b\{\neg p, \top\}$	(**RComm**)
$\vdash \phi \rightarrow \triangleright_a(\nabla_a\{p\} \wedge \nabla_b\{\neg p, \top\})$	(**RDist**)
$\vdash \phi \rightarrow \triangleright_a(\Box_a p \wedge \Diamond_b\neg p)$	(Definition of ∇)
$\vdash \phi \rightarrow \triangleright_a(\Box_a p \wedge \Box_b p)$	(Definition of $\Diamond$)

We will now show that the axiomatisation is sound with respect to K models.

Lemma 5.2.1. *The axiomatisation* $\mathbf{RML^K}$ *is sound with respect to the semantic class K.*

Proof. The soundness of the axioms **P** and **K**, and the rules **MP** and **NecK** can be shown by the same reasoning used to show that they are sound in basic modal logic. As the axioms **RP** and **R**, and the rule **NecR** involve only a single agent, their soundness can be shown by the same reasoning used to show that they are sound in the single-agent refinement quantified modal logic [23].

All that remains to be shown is the soundness of **RK**, **RComm**, and **RDist**.

RK Suppose that $M_s \in K$ is a Kripke model such that $M_s \models \bigwedge_{\gamma\in\Gamma} \Diamond_a\triangleright_a\gamma$.

We need to show that $M_s \models \triangleright_a\nabla_a\Gamma$. To do this we will construct a model $N_t \in K$, construct an a-simulation from N_t to M_s to show that $N_t \preceq_a M_s$, and finally show that $N_t \models \nabla_a\Gamma$.

We begin by constructing the model N_t. Consider $\gamma \in \Gamma$. From $M_s \models \Diamond_a\triangleright_a\gamma$, there exists a state $s^\gamma \in sR^M$ such that $M_{s^\gamma} \models \triangleright_a\gamma$. Therefore there exists a Kripke model $N^\gamma_{t^\gamma} \preceq_a M_{s^\gamma}$, via some a-simulation $\mathcal{R}^\gamma$, such that $N^\gamma_{t^\gamma} \models \gamma$. Without loss of generality we assume that the N^γ are disjoint.

Let t be a state not in S^M or any of the S^{N^γ}. Then we construct a Kripke model $N = (S^N, R^N, V^N)$ where:

$$\begin{aligned}
S^N &= \{t\} \cup S^M \cup \bigcup_{\gamma\in\Gamma} S^{N^\gamma} \\
R_a^N &= \{(t, t^\gamma) \mid \gamma \in \Gamma\} \cup R_a^M \cup \bigcup_{\gamma\in\Gamma} R_a^{N^\gamma} \\
R_b^N &= \{(t, t') \mid t' \in sR_b^M\} \cup R_b^M \cup \bigcup_{\gamma\in\Gamma} R_b^{N^\gamma} \text{ for } b \in A - \{a\} \\
V^N(p) &= \begin{cases} \{t\} \cup V^M(p) \cup \bigcup_{\gamma\in\Gamma} V^{N^\gamma}(p) & \text{if } s \in V^M(p) \\ V^M(p) \cup \bigcup_{\gamma\in\Gamma} V^{N^\gamma}(p) & \text{otherwise} \end{cases} \quad \text{for } p \in P
\end{aligned}$$

A representation of this model is pictured in Figure 5.1.

We construct an a-simulation $\mathcal{R}$ from N_t to M_s, where:

$$\mathcal{R} = \{(t, s)\} \cup \{(s', s') \mid s' \in S^M\} \cup \bigcup_{\gamma\in\Gamma} \mathcal{R}^\gamma$$

We must show that $\mathcal{R}$ satisfies **atoms**, **forth-**b for every $b \in A$, and **back-**b for every $b \in A - \{a\}$.

atoms We note that, by construction, the valuation of N matches the valuation of its corresponding states in M and each N^γ, and the valuation of N_t matches that of M_s. Therefore $\mathcal{R}$ satisfies **atoms**.

forth We next show that $\mathcal{R}$ satisfies **forth-**b for every $b \in A$. Let $b \in A$ and let $(u, v) \in \mathcal{R}$.

Suppose that $(u, v) \in \mathcal{R}^\gamma$ for some $\gamma \in \Gamma$. Then as $\mathcal{R}^\gamma$ is an a-simulation, it satisfies **forth-**b for every $b \in A$. Hence for every $u' \in uR_b^{N^\gamma} = uR_b^N$, there exists some $v' \in vR_b^M$ such that $(u', v') \in \mathcal{R}^\gamma \subseteq \mathcal{R}$.

Suppose instead that $(u, v) = (s', s')$ for some $s' \in S^M$. Then we note that $s'R_b^N = s'R_b^M$, and hence for every $s'' \in s'R_b^N$ we have that $s'' \in s'R_b^M$, and that $(s'', s'') \in \mathcal{R}$.

Finally suppose that $(u, u') = (t, s)$. We must consider the cases where $b = a$ and where $b \neq a$. So suppose that $b = a$. By construction, $tR_a^N = \{t^\gamma \mid \gamma \in \Gamma\}$, and hence $v = t^\gamma$ for some $\gamma \in \Gamma$. Hence we can take $s^\gamma \in sR_a^M$, and note that as

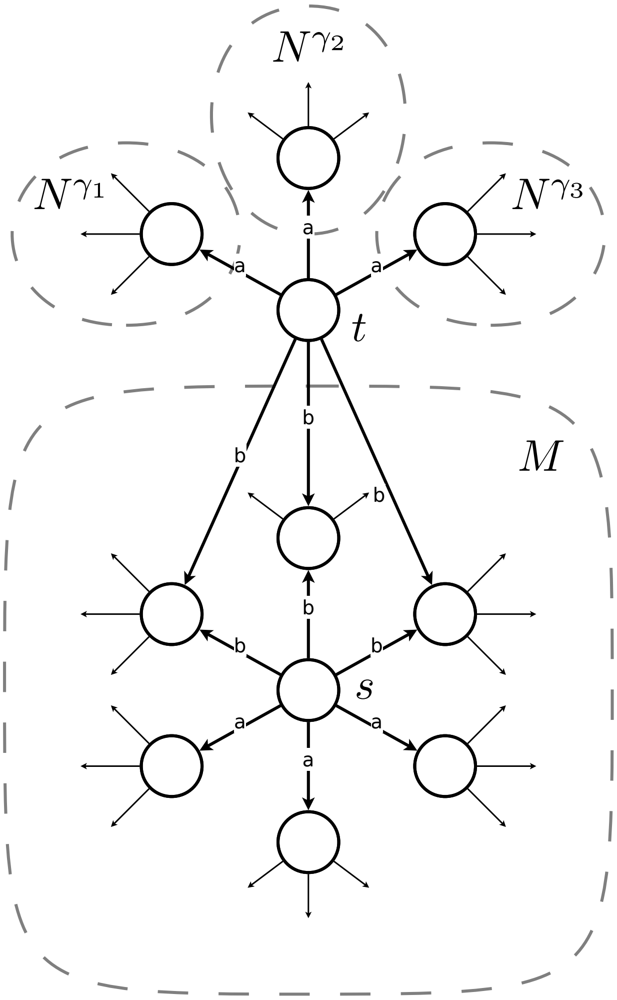

Figure 5.1: The model N is constructed by taking the model M and the models N^γ for every $\gamma \in \Gamma$, and connecting them with an extra node t. t is connected via an a-edge to t^γ from each of the N^γ, and is also connected via a b-edge to each b-successor of s in M. N_t is then the desired a-refinement of M_s.

$\mathcal{R}^\gamma$ is an a-simulation from M_{s^γ} to $N^\gamma_{t^\gamma}$, we know that $(t^\gamma, s^\gamma) \in \mathcal{R}^\gamma \subseteq \mathcal{R}$. Suppose that $b \neq a$. Then by construction, $tR^M_b = sR^M_b$, hence for every $t' \in tR^M_b$, we have that $t' \in sR^M_b$, and hence we know that $(t', t') \in \mathcal{R}$.

Therefore $\mathcal{R}$ satisfies **forth-**b for every $b \in A$.

back A similar argument to the above shows that $\mathcal{R}$ satisfies **back-**a for every $b \in A - \{a\}$.

Therefore $\mathcal{R}$ is an a-simulation, and $N_t \preceq_a M_s$.

Finally we show that $N_t \models \nabla_a \Gamma$. We must show that for each $\gamma \in \Gamma$ that $N_{t^\gamma} \models \gamma$. This follows from the fact that N_{t^γ} is bisimilar to $N^\gamma_{t^\gamma}$. This is obvious, as N contains a duplicate of N^γ, and N does not add any additional edges originating from states in S^{N^γ}. Hence from bisimulation invariance, $N_{t^\gamma} \models \gamma$ for every $\gamma \in \Gamma$, and hence $N_t \models \nabla_a \Gamma$.

As $N_t \preceq_a M_s$, and $N_t \models \nabla_a \Gamma$ we therefore have that $M_s \models \triangleright_a \nabla_a \Gamma$.

Conversely, suppose that $M_s \models \nabla_a \Gamma$. Then there exists a Kripke model $N_t \preceq_a M_s$, via some a-simulation $\mathcal{R}$, such that $N_t \models \nabla_a \Gamma$. From the definition of the cover operator, this implies that $N_t \models \Box_a \bigvee_{\gamma \in \Gamma} \gamma \wedge \bigwedge_{\gamma \in \Gamma} \Diamond_a \gamma$. In particular we note that for every $\gamma \in \Gamma$, $N_t \models \Diamond_a \gamma$, and so there exists some $t^\gamma \in tR^N_a$ such that $N_{t^\gamma} \models \gamma$. As $t^\gamma \in tR^N_a$, and $(t, s) \in \mathcal{R}$, by **forth-**a there exists some $s^\gamma \in sR^M_a$ such that $(t^\gamma, s^\gamma) \in \mathcal{R}$. Hence $\mathcal{R}$ is also an a-simulation from N_{t^γ} to M_{s^γ}, and so $M_{s^\gamma} \models \triangleright_a \gamma$. As for every $\gamma \in \Gamma$ we have that $s^\gamma \in sR^M_a$, we also have that $M_s \models \Diamond_a \triangleright_a \gamma$. Therefore we finally have that $M_s \models \bigwedge_{\gamma \in \Gamma} \Diamond_a \triangleright_a \gamma$.

Therefore **RK** is sound.

RComm Suppose that $M_s \in K$ is a Kripke model such that $M_s \models \nabla_b \{\triangleright_a \gamma \mid \gamma \in \Gamma\}$, where $a \neq b$.

We need to show that $M_s \models \triangleright_a \nabla_b \Gamma$. To do this we follow the same strategy as for proving **RK**: we construct an a-refinement $N_t \in K$, and show that $N_t \models \nabla_b \Gamma$.

We begin by constructing the model N_t. Consider $\gamma \in \Gamma$. From $M_s \models \nabla_b \{\triangleright_a \gamma \mid \gamma \in \Gamma\}$, there exists a state $s^\gamma \in sR^M_b$ such that $M_{s^\gamma} \models \triangleright_a \gamma$. Therefore there exists a Kripke model $N^\gamma_{t^\gamma} \preceq_a M_{s^\gamma}$, via some a-simulation $\mathcal{R}^\gamma$, such that $N^\gamma_{t^\gamma} \models \gamma$. Without loss of generality we assume that the N^γ are disjoint.

Let t be a state not in S^M or any of the S^{N^γ}. Then we construct a Kripke model $N = (S^N, R^N, V^N)$ where:

$$\begin{aligned}
S^N &= \{t\} \cup S^M \bigcup_{\gamma\in\Gamma} S^{N^\gamma} \\
R^N_b &= \{(t, t^\gamma) \mid \gamma \in \Gamma\} \cup R^M_b \cup \bigcup_{\gamma\in\Gamma} R^{N^\gamma}_b \\
R^N_c &= \{(t, t') \mid t' \in sR^M_c\} \cup R^M_c \cup \bigcup_{\gamma\in\Gamma} R^{N^\gamma}_c \text{ for } c \in A - \{b\} \\
V^N(p) &= \begin{cases} \{t\} \cup V^M(p) \cup \bigcup_{\gamma\in\Gamma} V^{N^\gamma}(p) & \text{if } s \in V^M(p) \\ V^M(p) \cup \bigcup_{\gamma\in\Gamma} V^{N^\gamma}(p) & \text{otherwise} \end{cases} \quad \text{for } p \in P
\end{aligned}$$

We construct an a-simulation $\mathcal{R}$ from N_t to M_s, where:

$$\mathcal{R} = \{(t, s)\} \cup \{(s', s') \mid s' \in S^M\} \cup \bigcup_{\gamma\in\Gamma} \mathcal{R}^\gamma$$

We note that $\mathcal{R}$ is an a-simulation, by similar arguments as used in the proof for **RK**. In particular, this means that $N_t \preceq_a M_s$.

We also note that for every $\gamma \in \Gamma$, that $N_{t^\gamma} \models \gamma$, by similar arguments as used in the proof for **RK**. In particular, this means that $N^\gamma_t \models \nabla_b \Gamma$.

Therefore $M_s \models \triangleright_a \nabla_b \Gamma$.

The converse, $\triangleright_a \nabla_b \Gamma \to \nabla_b \{\triangleright_a \gamma \mid \gamma \in \Gamma\}$ follows a similar proof to the relevant part in the proof for **RK**.

Therefore **RComm** is sound.

RDist Suppose that $M_s \in K$ is a Kripke model such that $M_s \models \bigwedge_{b\in B} \triangleright_a \nabla_b \Gamma_b$, where $B \subseteq A$.

We need to show that $M_s \models \triangleright_a \bigwedge_{b\in B} \nabla_b \Gamma_b$. To do this we follow the same strategy as for proving **RK**: we construct an a-refinement $N_t \in K$, and show that $N_t \models \triangleright_a \bigwedge_{b\in B} \nabla_b \Gamma_b$.

We begin by constructing the model N_t. Suppose that $a \in B$. Then we have $M_s \models \triangleright_a \nabla_a \Gamma_a$, and by **RK** this implies that $M_s \models \bigwedge_{\gamma\in\Gamma_a} \gamma$. We also have that for every $b \in B - \{a\}$ that $M_s \models \triangleright_a \nabla_a \Gamma_b$, and by **RComm** this implies that $M_s \models \nabla_b \{\triangleright_a \gamma \mid \gamma \in \Gamma_b\}$, and by the definition of the cover operator, this implies that $M_s \models \bigwedge_{\gamma\in\Gamma_b} \Diamond_b \triangleright_a \gamma$. Hence for every $b \in B$ and $\gamma \in \Gamma_b$, we have that $\Diamond_b \rhd_a \gamma$. This implies that for each $b \in B$ and each $\gamma \in \Gamma_b$ that there exists some $s^{b,\gamma} \in sR_b^M$ such that $M_{s^{a,\gamma}} \models \rhd_a \gamma$. Therefore there exists a Kripke model $N^{b,\gamma}_{t^{b,\gamma}} \preceq_a M_{s^{a,\gamma}}$ such that $N^{b,\gamma}_{t^{b,\gamma}} \models \gamma$. Without loss of generality we may assume that the $N^{b,\gamma}$ are disjoint.

Let t be a state not in S^M or any of the $S^{N^{b,\gamma}}$. Then we construct a Kripke model $N = (S^N, R^N, V^N)$ where:

$$\begin{aligned}
S^N &= \{t\} \cup S^M \cup \bigcup_{b \in A, \gamma \in \Gamma_b} S^{N^{b,\gamma}} \\
R_b^N &= \{(t, t^{b,\gamma}) \mid \gamma \in \Gamma_b\} \cup R_b^M \cup \bigcup_{c \in B, \gamma \in \Gamma_c} R_b^{N^{c,\gamma}} \text{ for } b \in B \\
R_b^N &= \{(t, t') \mid t' \in sR_b^M\} \cup R_b^M \cup \bigcup_{c \in B, \gamma \in \Gamma_c} R_b^{N^{c,\gamma}} \text{ for } b \in A \setminus B \\
V^N(p) &= \begin{cases} \{t\} \cup V^M(p) \cup \bigcup_{b \in B, \gamma \in \Gamma_b} V^{N^{b,\gamma}}(p) & \text{if } s \in V^M(p) \\ V^M(p) \cup \bigcup_{b \in B, \gamma \in \Gamma_b} V^{N^{b,\gamma}}(p) & \text{otherwise} \end{cases}
\end{aligned}$$

We construct an a-simulation $\mathcal{R}$ from N_t to M_s, where:

$$\mathcal{R} = \{(t, s)\} \cup \{(s', s') \mid s' \in S^M\} \bigcup_{b \in A, \gamma \in \Gamma_b} \mathcal{R}^\gamma$$

We note that this is an a-simulation, by similar arguments as used in the proof for **RK**. In particular, this means that $N_t \preceq_a M_s$.

We also note that for every $b \in A$, and $\gamma \in \Gamma_b$ that $N_{t^\gamma} \leftrightarroweq N^\gamma_{t^\gamma}$, by similar arguments as used in the proof for **RK**. In particular, this means that as $N^\gamma_{t^\gamma} \models \gamma$ that we also have $N_{t^\gamma} \models \gamma$, for every $b \in A$ and $\gamma \in \Gamma_b$. Therefore $N_t \models \nabla_b \Gamma_b$ for every $b \in A$, and therefore $N_t \models \bigwedge_{b \in A} \nabla_b \Gamma_b$.

Therefore $M_s \models \rhd_a \bigwedge_{b \in A} \nabla_b \Gamma_b$ and **RDist** is sound.

Therefore the axiomatisation $\mathbf{RML^K}$ is sound. □

We note that if the implication in **RDist** is strengthened to an equality, the resulting axiom is also sound. However this is easily derivable from the other axioms in $\mathbf{RML^K}$.

Lemma 5.2.2. *The following is derivable in* $\mathbf{RML^K}$*.*

$$\vdash \bigwedge_{b \in A} \rhd_a \nabla_b \Gamma_b \leftrightarrow \rhd_a \bigwedge_{b \in A} \nabla_b \Gamma_b$$

where Γ_b is a set of b-disjunctive normal formulae for every $b \in A$.

Proof (Sketch). The forward direction is the axiom **RDist**.

The converse can be derived in a more general form as $\triangleright_a(\phi\wedge\psi)\rightarrow\triangleright_a\phi\wedge\triangleright_a\psi$. The derivation is similar to the derivation for $\Box_a(\phi\wedge\psi)\rightarrow\Box_a\phi\wedge\Box_a\psi$ in the modal logic $\mathcal{L}^{\mathbf{K}}$, using the axiom **R** in place of **K**. $\Box$

We show the completeness of the axiomatisation $\mathbf{RML}^{\mathbf{K}}$ by a provably correct translation from $\mathcal{L}_{\triangleright}$ to $\mathcal{L}$. Completeness then follows from the completeness of $\mathcal{L}^{\mathbf{K}}$.

We introduce some equivalences that will be used by our translation.

Lemma 5.2.3. *The following are provable equivalences using* $\mathbf{RML}^{\mathbf{K}}$*:*

1. $\triangleright_a(\phi\vee\psi)\leftrightarrow\triangleright_a\phi\vee\triangleright_a\psi$
2. $\triangleright_a(\pi\wedge\bigwedge_{b\in B}\nabla_b\Gamma_b)\leftrightarrow\pi\wedge\bigwedge_{\gamma\in\Gamma_a}\Diamond_a\triangleright_a\gamma\wedge\bigwedge_{b\in B}\nabla_b\{\triangleright_a\gamma\mid\gamma\in\Gamma_b\}$ *where* π *is propositional,* $B\subseteq A$*, and* $a\in B$
3. $\triangleright_a(\pi\wedge\bigwedge_{b\in B}\nabla_b\Gamma_b)\leftrightarrow\pi\wedge\bigwedge_{b\in B}\nabla_b\{\triangleright_a\gamma\mid\gamma\in\Gamma_b\}$ *where* π *is propositional,* $B\subseteq A$*, and* $a\notin B$

Proof. (1) is derivable from **P** and **R** using the same strategy used to prove that $\Diamond_a(\phi\vee\psi)\leftrightarrow\Diamond_a\phi\vee\Diamond_a\psi$ is derivable from **P** and **K**.

(2) and (3) are derivable, by using **P**, **R** and **RP** to bring the propositional part π outside the $\triangleright_a$ operator, using **RDist** $\triangleright_a$ operator into the cover operators inside the conjunction, and then using **RK** or **RComm** as appropriate for each cover operator. $\Box$

Lemma 5.2.4. *Every formula of* $\mathcal{L}_{\triangleright}$ *is provably equivalent to a formula of* $\mathcal{L}$ *with the axiomatisation* $\mathbf{RML}^{\mathbf{K}}$*.*

Proof. Let $\alpha\in\mathcal{L}_{\triangleright}$. We assume without loss of generality that all $\blacktriangleright$ operators are expressed as $\triangleright$ operators, by the equivalence $\blacktriangleright_a\phi\leftrightarrow\neg\triangleright_a\neg\phi$. We prove by induction on the number of occurrences of $\triangleright$ in α that α is equivalent to a $\triangleright$-free formula, and therefore to a formula in $\mathcal{L}$. The base case where α contains no $\triangleright$ operators is trivial, as a $\triangleright$-free formula is a formula in $\mathcal{L}$. Suppose instead that α contains $n+1$ $\triangleright$ operators, and assume that any formula with n $\triangleright$ operators is provably equivalent to a formula in $\mathcal{L}$. We use the axioms of $\mathbf{RML}^{\mathbf{K}}$ to show that α is provably equivalent to a formula with n $\triangleright$ operators, and that therefore by the induction hypothesis it is provably equivalent to a formula in $\mathcal{L}$.

Given any subformula from α of type $\triangleright_a\beta$, such that β is $\triangleright$-free. Without loss of generality, by Lemma 5.1.2 we may assume that β is in disjunctive normal form. We prove by induction on the structure of β that $\triangleright_a\beta$ is provably equivalent to a formula $\chi \in \mathcal{L}$. The induction hypothesis is that for any proper subformula ϕ of β that $\triangleright_a\phi$ is equivalent to a formula in $\mathcal{L}$.

The base case is when β is a propositional formula. In this case, from **P** and **RP**, we have that $\triangleright_a\beta \leftrightarrow \beta$, and therefore $\triangleright_a\beta$ is equivalent to a formula in $\mathcal{L}$.

The inductive case is when $\beta = \phi \vee \psi$, or when $\beta = \pi \wedge \bigwedge_{b\in B} \nabla_b\Gamma_b$, where $B \subseteq A$. We note that we can use the equivalences from Lemma 5.2.3 to push the $\triangleright_a$ operator inside so that it is applied to subformulae of β. We can then use the induction hypothesis to replace each occurrence of the $\triangleright_a$ operator applied to a subformula of β with an equivalent formula in $\mathcal{L}$. The resulting formula is also in $\mathcal{L}$.

Therefore by the induction, $\triangleright_a\beta$ is equivalent to a formula $\chi \in \mathcal{L}$ for every $\beta \in \mathcal{L}$.

Hence replacing $\triangleright_a\beta$ in α with χ gives an equivalent formula that contains only n $\triangleright$ operators.

Therefore by the induction, α is equivalent to a formula in $\mathcal{L}$. □

The rest of the completeness proof is to show that, given the above translation into $\mathcal{L}$, we can show completeness by using these translations along with the completeness of $\mathcal{L}^{\mathbf{K}}$.

Corollary. *Let $\phi \in \mathcal{L}_\triangleright$ be given and $\psi \in \mathcal{L}$ be semantically equivalent to ϕ. If ψ is a theorem in $\mathcal{L}^{\mathbf{K}}$, then ϕ is a theorem in $\mathbf{RML}^{\mathbf{K}}$.*

Proof. Let $\phi \in \mathcal{L}_\triangleright$ and let $\psi \in \mathcal{L}$ be semantically equivalent to ϕ. By Lemma 5.2.4, we can obtain some $\phi' \in \mathcal{L}$ that is semantically equivalent to ϕ (and thus also to ψ) by following the given translation steps. We can extend a derivation of ψ to a derivation of ϕ' as the two are semantically equivalent in $\mathcal{L}^{\mathbf{K}}$, and by the completeness of $\mathcal{L}^{\mathbf{K}}$ this equivalence is derivable. As $\mathbf{RML}^{\mathbf{K}}$ is a conservative extension of $\mathcal{L}^{\mathbf{K}}$, this equivalence is therefore also derivable in $\mathbf{RML}^{\mathbf{K}}$. The derivation can be further extended to ϕ by observing that all of the reduction steps in Lemma 5.2.4 are provable equivalences in $\mathbf{RML}^{\mathbf{K}}$. Therefore ϕ is a theorem in $\mathbf{RML}^{\mathbf{K}}$. □

Lemma 5.2.5. *The axiom schema $\mathbf{RML}^{\mathbf{K}}$ is complete with respect to the semantic class K.*

Proof. Let $\phi \in \mathcal{L}_{\triangleright}$ such that $K \vDash_{\triangleright} \phi$. Then by Lemma 5.2.4, there exists a semantically equivalent formula $\psi \in \mathcal{L}$ which is $\triangleright$-free. As $K \vDash_{\triangleright} \phi$ and $\phi \leftrightarrow \psi$, then $K \vDash_{\triangleright} \psi$. As ψ is $\triangleright$-free, then it follows that $K \vDash \psi$, and by the completeness of $\mathbf{RML}^{\mathbf{K}}$ it follows that $\vdash_{\mathbf{K}} \psi$. Therefore by Corollary 5.2 we have that $\vdash_{\mathbf{RML}^{\mathbf{K}}} \phi$. $\square$

Theorem 5.2.6. *The axiomatisation* $\mathbf{RML}^{\mathbf{K}}$ *is sound and complete with respect to the semantic class K.*

Proof. The soundness proof is given in Lemma 5.2.1 and the completeness proof is given in Lemma 5.2.5. $\square$

We note that, as in the axiomatisation for the single-agent epistemic and doxastic logics, the completeness proofs above were performed with a provably correct translation from $\mathcal{L}_{\triangleright}$ to $\mathcal{L}$, under the semantics of $\mathcal{L}_{\triangleright}^{\mathbf{K}}$. This shows that $\mathcal{L}_{\triangleright}^{\mathbf{K}}$ is expressively equivalent to $\mathcal{L}^{\mathbf{K}}$, and allows us to show several results. In particular, $\mathcal{L}_{\triangleright}^{\mathbf{K}}$ is decidable.

Theorem 5.2.7. *The logic* $\mathcal{L}_{\triangleright}^{\mathbf{K}}$ *is decidable.*

This can be shown by following similar reasoning as used for the proof of Theorem 4.2.9.

CHAPTER 6

Multi-agent refinement quantified doxastic logic

In this chapter we provide an axiomatisation of the multi-agent refinement quantified doxastic logic. The axiomatisation for the doxastic variant differs from the modal variant due to the restriction to doxastic models, and due to the interpretation of the refinement quantifier, which is also restricted to considering only doxastic models.

6.1 Technical preliminaries

In Chapter 4 we gave an axiomatisation of the single-agent variant of the refinement quantified doxastic logic. This axiomatisation relied on a prenex normal form that restricted formulae by prohibiting nested modalities. The prenex normal form simplified the axiomatisation and proof of soundness by avoiding complications that would arise due to the transitive and Euclidean properties of *KD45* models when nested modalities are present.

For the axiomatisation of the multi-agent variant, we introduce what we call an alternating disjunctive normal form that serves the same purpose as the prenex normal form served for the single-agent variant. The alternating disjunctive normal form can be seen as a generalisation of the prenex normal form to multi-agent logics. It can alternatively be seen as a restriction of the previously defined disjunctive normal form. In multi-agent doxastic logic, it is not possible to write every formula without using nested modalities; for example, a formula involving nested modalities belonging to different agents, such as $\Box_a \Box_b p$ cannot be expressed without nested modalities. Thus the prenex normal form previously introduced is not appropriate in this setting. Rather, our alternating disjunctive normal form prohibits modalities of the same agent from being nested directly within one another, without a modality belonging to another agent appearing inbetween. For example, $\Box_a \Box_b \Diamond_a p$ is allowed, whereas $\Box_a \Diamond_a p$ is not. This allows us to avoid the same complications due to transitivity and Euclideaness that the prenex normal form allowed us to avoid.

Similar to the prenex normal form used in the single-agent case, we will define the alternating disjunctive normal form in terms of basic modal operators first, and then provide another form which is in terms of the cover operator. We will also show an important property of these formulae which we rely on for our soundness proof.

Definition 6.1.1 (Alternating disjunctive normal form)**.** A formula in a-alternating disjunctive normal form is defined by the following abstract syntax, where α is a formula in a-alternating disjunctive normal form:

$$\begin{array}{rcl} \alpha & ::= & \delta \mid \alpha \vee \alpha \\ \delta & ::= & \pi \mid \Box_b \gamma_b \mid \Diamond_b \gamma_b \mid \delta \wedge \delta \end{array}$$

where π stands for a propositional formula, $b \in A - \{a\}$, and γ_b stands for a formula in b-alternating disjunctive normal form.

A formula in alternating disjunctive normal form is defined by the following abstract syntax, where α is a formula in alternating disjunctive normal form:

$$\begin{array}{rcl} \alpha & ::= & \delta \mid \alpha \vee \alpha \\ \delta & ::= & \pi \mid \Box_a \gamma_a \mid \Diamond_a \gamma_a \mid \delta \wedge \delta \end{array}$$

where π stands for a propositional formula, $a \in A$, and γ_a stands for a formula in a-alternating disjunctive normal form.

We will now show that every formula of $\mathcal{L}$ is equivalent to a formula in alternating disjunctive normal form, under the semantics of doxastic logic.

Lemma 6.1.1. *We have the following equivalences in* $\mathcal{L}^{\mathbf{KD45}}$*:*

$$\begin{array}{rcl} \Box_a(\pi \vee (\alpha \wedge \Box_a \beta)) & \leftrightarrow & (\Box_a(\pi \vee \alpha) \wedge \Box_a \beta) \vee (\Box_a \pi \wedge \neg \Box_a \beta) \\ \Box_a(\pi \vee (\alpha \wedge \Diamond_a \beta)) & \leftrightarrow & (\Box_a(\pi \vee \alpha) \wedge \Diamond_a \beta) \vee (\Box_a \pi \wedge \neg \Diamond_a \beta) \end{array}$$

This is proven by Meyer and van der Hoek [19] for the single-agent epistemic logic, however the same proof also applies to $\mathcal{L}^{\mathbf{KD45}}$.

Meyer and van der Hoek remarked that the only use of the reflexivity axiom of $\mathcal{L}^{\mathbf{S5}}$, $\mathbf{T}$, in the proof, is in the form of the theorems $\vdash \Box\Box\phi \to \Box\phi$, and $\vdash \Box\neg\Box\phi \to \neg\Box\phi$. Therefore the proof holds for any logic which replaces $\mathbf{T}$ with axioms entailing both of these properties. Both of these properties are obviously valid in $\mathcal{L}^{\mathbf{KD45}}$, and therefore the proof by Meyer and van der Hoek [19] applies to this result.

Lemma 6.1.2. *Every formula of $\mathcal{L}$ is equivalent to a formula in alternating disjunctive normal form, under the semantics of $\mathcal{L}^{\mathbf{KD45}}$.*

Proof. We use a proof similar to the proof for prenex normal form, given by Meyer and van der Hoek [19].

Let $\alpha \in \mathcal{L}$. Without loss of generality, by Lemma 5.1.1, we may assume that α is in negation normal form. We prove by induction over the structure of α that α is equivalent to a formula in alternating disjunctive normal form. The induction hypothesis is that every strict subformula of α has an equivalent in alternating disjunctive normal form.

The proof is similar to the proof for Lemma 5.1.2, except for the cases where $\alpha = \Box_a\phi$, $\alpha = \Diamond_a\phi$ and $\alpha\phi \vee \psi$.

Suppose that $\alpha = \Box_a\phi$. By the induction hypothesis, there is a formula ϕ' in alternating disjunctive normal form that is equivalent to ϕ. Suppose that ϕ' is not in a-alternating disjunctive normal form (otherwise we are done). Then ϕ' contains some conjunct of the form $\Box_a\beta$ or $\Diamond_a\beta$. Suppose that we have a conjunct of the form $\Box_a\beta$. Then we can rewrite ϕ' as $\phi' = \pi \vee (\alpha \wedge \Box_a\beta)$. By Lemma 6.1.1, we get that $\phi \equiv (\Box_a(\pi \vee \alpha) \wedge \Box_a\beta) \vee (\Box_a\pi \wedge \neg\Box_a\beta)$. We can use the other equivalence from Lemma 6.1.1 in the case that ϕ' contains a conjunct of the form $\Diamond_a\beta$. Proceeding in this fashion, we can pull out each occurrence of $\Box_a\beta$ or $\Diamond_a\beta$ inside ϕ' until we have rewritten ϕ' in a-alternating disjunctive normal form, as ϕ''. Then $\Box_a\phi''$ is in alternating disjunctive normal form.

Suppose that $\alpha = \Diamond_a\phi$. By the induction hypothesis, there is a formula ϕ' in alternating disjunctive normal form that is equivalent to ϕ. Suppose that ϕ' is not in a-alternating disjunctive normal form (otherwise we are done). Then ϕ' contains some conjunct of the form $\Box_a\beta$ or $\Diamond_a\beta$. Suppose that we have a conjunct of the form $\Box_a\beta$. Then we can rewrite ϕ' as $\phi' = \pi \vee (\alpha \wedge \Box_a\beta)$. Then we get that $\phi \equiv \Diamond_a\pi \vee (\Diamond_a\alpha \wedge \Box_a\beta)$. We get similar if ϕ' contains a conjunct of the form $\Diamond_a\beta$. Proceeding in this fashion, we can pull out each occurrence of $\Box_a\beta$ or $\Diamond_a\beta$ inside ϕ' until we have rewritten ϕ' in a-alternating disjunctive normal form, as ϕ''. Then $\Box_a\phi''$ is in alternating disjunctive normal form.

Suppose that $\alpha = \phi \wedge \psi$. By the induction hypothesis, there are formulae ϕ' and ψ' in alternating disjunctive normal form that are equivalent to ϕ and ψ respectively. Then $\phi \wedge \psi \leftrightarrow \phi' \wedge \psi'$. As ϕ' and ψ' are in alternating disjunctive normal form, then $\phi' = \delta_1 \vee \cdots \vee \delta_m$ and $\psi' = \gamma_1 \vee \cdots \vee \gamma_n$ for some $m, n \geq 0$, where each of the δ_i and γ_i are conjunctions. Then we can rewrite α as a disjunction of conjunctions, by the following equivalence:

$$\phi' \wedge \psi' \leftrightarrow \bigvee_{i \leq m, j \leq n} \delta_i \wedge \gamma_j$$

The resulting formula is in alternating disjunctive normal form.

Therefore every formula of $\mathcal{L}$ is equivalent to a formula in alternating disjunctive normal form. $\square$

We note that as the alternating disjunctive normal form introduced here is a generalisation of the prenex normal form of Definition 4.1.1, conversion into this form can result in a formula that is exponentially larger than the original formula in the worst case.

As in the proof for refinement quantified modal logic, we formulate our axiomatisation in terms of the cover operator, ∇, so we use a cover logic version of this alternating disjunctive normal form.

Definition 6.1.2 (Cover logic disjunctive normal form)**.** A formula in a-cover logic alternating disjunctive normal form is defined by the following abstract syntax:

$$\alpha ::= \pi \wedge \bigwedge_{b \in B} \nabla_b \Gamma_b \mid \alpha \vee \alpha$$

where π stands for a propositional formula, $B \subseteq A - \{a\}$, and Γ_b stands for a finite, non-empty set of formulae in b-cover alternating disjunctive normal form.

A formula in cover logic alternating disjunctive normal form is defined by the following abstract syntax:

$$\alpha ::= \pi \wedge \bigwedge_{a \in B} \nabla_a \Gamma_a \mid \alpha \vee \alpha$$

where π stands for a propositional formula, $B \subseteq A$, and Γ_a stands for a finite, non-empty set of formulae in a-cover logic alternating disjunctive normal form.

Lemma 6.1.3. *Every formula of $\mathcal{L}$ is equivalent to a formula in cover logic alternating disjunctive normal form, under the semantics of $\mathcal{L}^{\mathbf{KD45}}$.*

Proof. Let $\alpha \in \mathcal{L}$. Without loss of generality, by Lemma 6.1.2, we can write α in alternating disjunctive normal form. By Lemma 5.1.2, we can rewrite α as an equivalent form using the modal version of disjunctive normal form. We note that if a subformula ϕ of α is in a-alternating disjunctive normal form, then ϕ does not contain any a-modalities at the top level, and therefore the result of the conversion on the subformula is a formula in a-cover logic alternating disjunctive normal form. Thus the result of the conversion is a formula in cover logic alternating disjunctive normal form. $\square$

The cover logic alternating disjunctive normal form will be used in our completeness proofs.

Finally we show an important property of the alternating disjunctive normal form that we rely on for our completeness proofs.

Lemma 6.1.4. *If ϕ is a formula in a-alternating disjunctive normal form, and $M_s \in KD45$ is a doxastic model such that $M_s \vDash \phi$, then there exists a doxastic model $N_t \in KD45$ such that $N_t \vDash \phi$ and $tR_a^N = R_a^N t = \{t\}$ and $R_b^N t = \emptyset$ for every $b \in A - \{a\}$.*

Proof. Suppose that ϕ is an a-alternating disjunctive normal formula, and that $M_s \in KD45$ is a doxastic model such that $M_s \vDash \phi$.

Let t be a state not in S^M. Then we construct a Kripke model $N = (S^N, R^N, V^N)$ where:

$$\begin{array}{rcl} S^N & = & \{t\} \cup S^M \\ R_a^N & = & \{(t,t)\} \cup R_a^M \\ R_b^N & = & \{(t,s') \mid s' \in sR_b^M\} \cup R_b^M \text{ for every } b \in A - \{a\} \\ V^N(p) & = & \begin{cases} \{t\} \cup V^M(p) & \text{if } s \in V^M(p) \\ V^M(p) & \text{otherwise} \end{cases} \end{array}$$

We note that $tR_a^N = R_a^N t = \{t\}$ and $R_b^N t = \emptyset$ for every $b \in A - \{a\}$.

We must first show that N is a doxastic model, and then that $N_t \vDash \phi$. The latter will be shown by proving that for every $s' \in tR_b^N$, the state $N_{s'}$ is bisimilar to $M_{s'}$.

First we show that N is a doxastic model. The relation R_a^N consists of the relation R_a^M, combined with the relationship (t,t). The latter addition ensures the serial property given the new element in S^N, and as there are no other relationships involving t, its addition preserves the transitivity and Euclideaness of R_a^M. Hence R_a^N is serial, transitive and Euclidean. As R_b^M is serial for S^M, and $tR_b^M = sR_b^M \neq \emptyset$, then R_b^N is serial for S^N. The relation R_b^N for $b \in A - \{a\}$ consists of the relation R_b^M, combined with relationships from t to all of the successors of $s \in S^M$. As R_b^M is transitive and Euclidean, the additional relationships, which are simply duplicates of relationships starting at s from R_b^M, also satisfy transitivity and Euclideaness, therefore R_b^N is also transitive and Euclidean. Therefore N is a doxastic model.

We next show that for every $b \in A - \{a\}$ and each successor $s' \in sR_b^M$ of s, the state $N_{s'}$ is bisimilar to $M_{s'}$. This is by the bisimulation relation $\mathcal{R} =$

$\{(s', s') \mid s' \in S^M\}$, mapping states in M to their corresponding states in N. This is clearly a bisimulation as the only state in S^N which is not in S^M is t, and there are no edges leading to t from any state in S^M.

As ϕ is in a-alternating disjunctive normal form, ϕ has the form $\phi = \delta_1 \vee \cdots \vee \delta_m$, where each δ_i has the form $\delta_i = \gamma_{i1} \wedge \cdots \wedge \gamma_{in_i}$, and each γ_{ij} is either a propositional formula, or has the form $\Box_b\psi$ or $\Diamond_b\psi$ for some $b \in A - \{a\}$. As $M_s \vDash \phi$, there exists some $i = 1, \ldots, m$ such that $M_s \vDash \delta_i$. Therefore, for every $j = 1, \ldots, n_i$, we have that $M_s \vDash \gamma_{ij}$. Suppose that γ_{ij} is a propositional formula. Then by construction N_t has the same valuation as M_s, and hence is equivalent under propositional formulae. Therefore $N_t \vDash \gamma_{ij}$. Suppose instead that $\gamma_{ij} = \Diamond_b\psi$ for some $b \in A - \{a\}$ and some formula ψ. Then there exists some $s' \in sR_b^M$ such that $M_{s'} \vDash \psi$. From above, we know that $N_{s'}$ is bisimilar to $M_{s'}$, and so by bisimulation invariance we have that $N_{s'} \vDash \psi$. By construction, $s' \in sR_b^M = tR_b^N$, and hence $N_t \vDash \Diamond_b\psi$. A similar argument can be used for the case where $\gamma_{ij} = \Box_b\psi$. Hence for every $j = 1, \ldots, n_i$, we have that $N_t \vDash \gamma_i j$.

Hence $N_t \vDash \delta_i$ and so $N_t \vDash \psi$. $\square$

6.2 Axiomatisation

We provide an axiomatisation of the multi-agent refinement quantified doxastic logic, $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$, and prove its soundness and completeness.

Definition 6.2.1 (Axiomatisation $\mathbf{RML}^{\mathbf{KD45}}$)**.** The axiomatisation $\mathbf{RML}^{\mathbf{KD45}}$ is a substitution schema consisting of the following axioms:

$$
\begin{array}{rl}
\mathbf{P} & \text{All propositional tautologies} \\
\mathbf{K} & \Box(\phi \rightarrow \psi) \rightarrow \Box\phi \rightarrow \Box\psi \\
\mathbf{D} & \Box\phi \rightarrow \Diamond\phi \\
\mathbf{4} & \Box\phi \rightarrow \Box\Box\phi \\
\mathbf{5} & \Diamond\phi \rightarrow \Box\Diamond\phi \\
\mathbf{R} & \blacktriangleright_a(\phi \rightarrow \psi) \rightarrow \blacktriangleright_a\phi \rightarrow \blacktriangleright_a\psi \\
\mathbf{RP} & \blacktriangleright_a\alpha \leftrightarrow \alpha \text{ where } \alpha \text{ is a propositional formula} \\
\mathbf{RComm} & \triangleright_a\nabla_b\Gamma_b \leftrightarrow \nabla_b\{\triangleright_a\gamma \mid \gamma \in \Gamma_b\} \text{ where } a \neq b \\
\mathbf{RDist} & \bigwedge_{b\in B} \triangleright_a\nabla_b\Gamma_b \rightarrow \triangleright_a \bigwedge_{b\in B} \nabla_b\Gamma_b \text{ where } B \subseteq A \\
\mathbf{RKD45} & \triangleright_a\nabla_a\Gamma_a \leftrightarrow \bigwedge_{\gamma\in\Gamma_a} \Diamond_a\triangleright_a\gamma
\end{array}
$$

where for every $a \in A$, the set Γ_a is a non-empty set of a-alternating disjunctive normal formulae.

Along with the rules:

$$
\begin{array}{rl}
\mathbf{MP} & \text{From } \vdash \phi \rightarrow \psi \text{ and } \vdash \phi, \text{ infer } \vdash \psi \\
\mathbf{NecK} & \text{From } \vdash \phi \text{ and } \vdash \Box_a\phi \\
\mathbf{NecR} & \text{From } \vdash \phi \text{ and } \vdash \blacktriangleright_a\phi
\end{array}
$$

The axiomatisation $\mathbf{RML}^{\mathbf{KD45}}$ is similar to the axiomatisation $\mathbf{RML}^{\mathbf{K}}$ in many respects. The notable differences are the inclusion of the **KD45** axioms, **D**, **4** and **5**, and the fact that the axioms **RComm**, **RDist** and **RKD45** are restricted to sets of a-alternating disjunctive normal formulae. This latter restriction is needed for our soundness proofs. Our soundness proofs follow the same style as those for $\mathbf{RML}^{\mathbf{K}}$: we assume that we have a model that satisfies a certain formula, and then manually construct a refinement of that model, using smaller refinements of the original model. The difference between the soundness proofs for $\mathbf{RML}^{\mathbf{K}}$ and the soundness proofs for $\mathbf{RML}^{\mathbf{KD45}}$ are that in the soundness proof for $\mathbf{RML}^{\mathbf{KD45}}$ the refinement that we construct must be a doxastic model, whereas in the soundness proof for $\mathbf{RML}^{\mathbf{K}}$, the refinement can be any Kripke model. The fact that the formulae are a-alternating disjunctive normal formulae means that we can invoke Lemma 6.1.4 in various places, which allows us to ensure that the refinement that we construct is a doxastic model.

Example 6.2.1. Once again we recall the coin-flipping example given in Example 3.3.1. We note that the derivation given in Example 5.2.1, of $\neg\Box_a p \wedge \neg\Box_a \neg p \wedge \neg\Box_b p \wedge \neg\Box_b \neg p \to \triangleright_a(\Box_a p \wedge \neg\Box_b p)$ can be adapted to a derivation in $\mathbf{RML^{KD45}}$, by replacing the instance of **RK** with **RKD45**, noting that the formulae inside the cover operators at each application of the axioms **RKD45**, **RComm** and **RDist** are all in an appropriate a-alternating disjunctive normal form or b-alternating disjunctive normal form.

Example 6.2.2. We give an example derivation using $\mathbf{RML^{KD45}}$, based on the coin-flipping example of Example 3.3.2. We suppose that neither Alice nor Bob believes that the coin has landed on heads or that it has landed on tails, and we show that it is possible for Alice to learn that the coin has landed on heads, whilst Bob continues to believe that Alice doesn't believe either way. Thus we give a derivation of $(\neg\Box_a p \wedge \neg\Box_a \neg p \wedge \Box_b \neg\Box_a p \wedge \Box_b \neg\Box_a \neg p) \to \triangleright_a(\Box_a p \wedge \Box_b \neg\Box_a p)$. Let $\phi = \neg\Box_a p \wedge \neg\Box_a \neg p \wedge \Box_b \neg\Box_a p \wedge \Box_b \neg\Box_a \neg p)$.

$\vdash \phi \to \neg\Box_a \neg p \wedge \Box_b \neg\Box_a p$	(**P**)
$\vdash \phi \to \Diamond_a p \wedge \Box_b \Diamond_a \neg p$	(Definition of $\Diamond$)
$\vdash \phi \to \Diamond_a p \wedge \Box_b(\Diamond_a \neg p \wedge \Diamond_a \top)$	(Consequence of **D**)
$\vdash \phi \to \Diamond_a \neg\neg p \wedge \Box_b(\Diamond_a \neg\neg\neg p \wedge \Diamond_a \neg\neg\top)$	(**P**)
$\vdash \phi \to \Diamond_a \neg \blacktriangleright_a \neg p \wedge \Box_b(\Diamond_a \neg \blacktriangleright_a \neg\neg p \wedge \Diamond_a \neg \blacktriangleright_a \neg\top)$	(**RP**)
$\vdash \phi \to \Diamond_a \triangleright_a p \wedge \Box_b(\Diamond_a \triangleright_a \neg p \wedge \Diamond_a \triangleright_a \top)$	(Definition of $\triangleright$)
$\vdash \phi \to \triangleright_a \nabla_a\{p\} \wedge \Box_b(\Diamond_a \triangleright_a \neg p \wedge \Diamond_a \triangleright_a \top)$	(**RKD45**)
$\vdash \phi \to \triangleright_a \nabla_a\{p\} \wedge \Box_b \triangleright_a \nabla_a\{\neg p, \top\}$	(**RKD45**)
$\vdash \phi \to \triangleright_a \nabla_a\{p\} \wedge \nabla_b\{\triangleright_a \nabla_a\{\neg p, \top\}\}$	(Definition of ∇)
$\vdash \phi \to \triangleright_a \nabla_a\{p\} \wedge \triangleright_a \nabla_b\{\nabla_a\{\neg p, \top\}\}$	(**RComm**)
$\vdash \phi \to \triangleright_a(\nabla_a\{p\} \wedge \nabla_b\{\nabla_a\{\neg p, \top\}\})$	(**RDist**)
$\vdash \phi \to \triangleright_a(\Box_a p \wedge \Box_b \Diamond_a \neg p)$	(Definition of ∇)
$\vdash \phi \to \triangleright_a(\Box_a p \wedge \Box_b \neg\Box_a p)$	(Definition of $\Diamond$)

We will now show that the axiomatisation is sound with respect to *KD45* models.

Lemma 6.2.1. *The axiomatisation* $\mathbf{RML^{KD45}}$ *is sound with respect to the semantic class KD45.*

Proof. The soundness of the axioms **P** and **K**, and the rules **MP** and **NecK** can be shown by the same reasoning used to show that they are sound in basic modal logic. As the axioms **RP** and **R**, and the rule **NecR** involve only a single agent, their soundness can be shown by the same reasoning used to show that they are sound in the single-agent refinement quantified modal logic [23].

All that remains to be shown is the soundness of **RKD45**, **RComm**, and **RDist**.

RKD45 Suppose that $M_s \in \mathit{KD45}$ is a doxastic model such that $M_s \models \bigwedge_{\gamma \in \Gamma_a} \Diamond_a \triangleright_a \gamma$, where Γ_a is a non-empty set of a-alternating disjunctive normal formulae.

We need to show that $M_s \models \triangleright_a \nabla_a \Gamma_a$. To do this we will construct a model $N_t \in \mathit{KD45}$, construct an a-simulation from N_t to M_s to show that $N_t \preceq_a M_s$, and finally show that $N_t \models \nabla_a \Gamma_a$. Compared to the proof of **RK** from Lemma 5.2.1, we must construct N_t such that $N_t \in \mathit{KD45}$, and we must prove that this is so.

We begin by constructing the model N_t. Consider $\gamma \in \Gamma_a$. From $M_s \models \Diamond_a \triangleright_a \gamma$, there exists a state $s^\gamma \in sR^M$ such that $M_{s^\gamma} \models \triangleright_a \gamma$. Therefore there exists a doxastic model $N^\gamma_{t^\gamma} \preceq_a M_{s^\gamma}$, via some a-simulation $\mathcal{R}^\gamma$, such that $N^\gamma_{t^\gamma} \models \gamma$. Without loss of generality we assume that the N^γ are disjoint. We may also assume, by Lemma 6.1.4, that for every $\gamma \in \Gamma$, that $t^\gamma R_a^{N^\gamma} = R_a^{N^\gamma} t^\gamma = \{t^\gamma\}$ and that $R_b^{N^\gamma} t^\gamma = \emptyset$ for every $b \in A - \{a\}$.

Let t be a state not in S^M or any of the S^{N^γ}. Then we construct a Kripke model $N = (S^N, R^N, V^N)$ where:

$$
\begin{aligned}
S^N &= \{t\} \cup S^M \cup \bigcup_{\gamma \in \Gamma_a} S^{N^\gamma} \\
R_a^N &= \{(t, t^\gamma) \mid \gamma \in \Gamma_a\} \cup \{(t^\gamma, t^{\gamma'}) \mid \gamma, \gamma' \in \Gamma_a\} \cup R_a^M \cup \bigcup_{\gamma \in \Gamma_a} R_a^{N^\gamma} \\
R_b^N &= \{(t, t') \mid t' \in sR_b^M\} \cup R_b^M \cup \bigcup_{\gamma \in \Gamma_a} R_b^{N^\gamma} \text{ for } b \in A - \{a\} \\
V^N(p) &= \begin{cases} \{t\} \cup V^M(p) \cup \bigcup_{\gamma \in \Gamma_a} V^{N^\gamma}(p) & \text{if } s \in V^M(p) \\ V^M(p) \cup \bigcup_{\gamma \in \Gamma_a} V^{N^\gamma}(p) & \text{otherwise} \end{cases} \quad \text{for } p \in P
\end{aligned}
$$

A representation of this model is pictured in Figure 6.1.

We show that N is a doxastic model. We must show two cases: that R_a^N is serial, transitive and Euclidean, and that R_b^N is also serial, transitive and Euclidean, for every $b \in A - \{a\}$. As M and each N^γ are doxastic models, the relations R_c^M and $R_c^{N^\gamma}$ are serial, transitive and Euclidean, for every $c \in A$. We observe that S^N is a union of $\{t\}$, S^M and each S^{N^γ}. As Γ_a is non-empty there is at least one relationship (t, t^γ) in R_a^N, therefore as R_a^N also contains R_a^M and each $R_a^{N^\gamma}$, each of which are serial, then R_a^N is also serial. Similarly, as M is serial, sR_b^M is non-empty, and so there is at least one relationship (t, t') where $t' \in sR_b^M$ in R_b^N; therefore as R_b^N also contains R_b^M and $R_b^{N^\gamma}$, then R_b^N is also serial. To show the transitivity and Euclideaness of R_a^N, we first observe that the relation can be considered in three independent parts; the relationships between t and t^γ, and between each t^γ; the relation R_a^M; and the relations $R_a^{N^\gamma}$. The

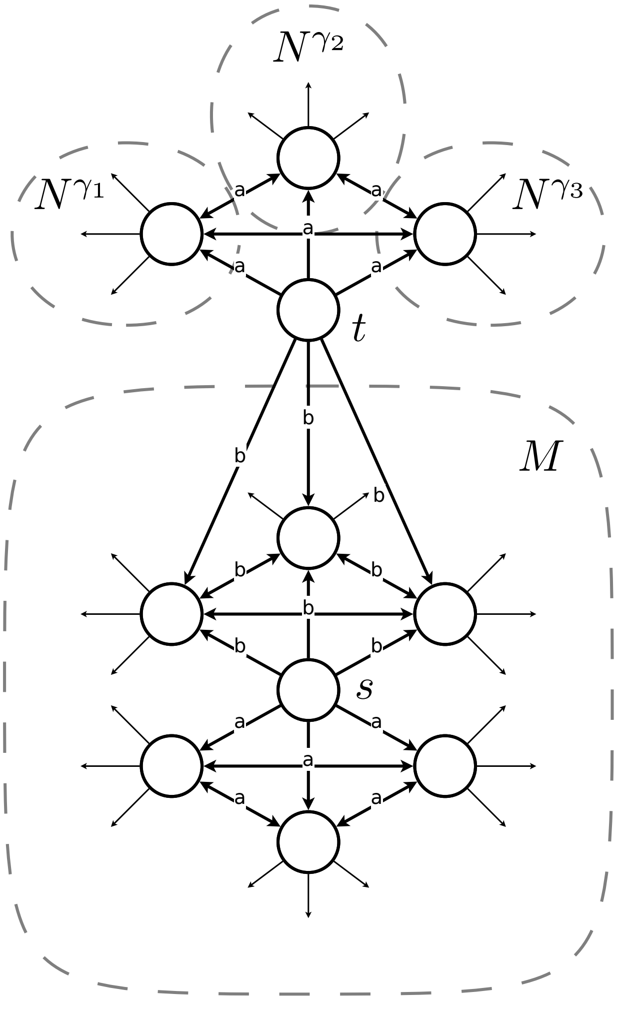

Figure 6.1: The model N is constructed by taking the model M and the models N^γ for every $\gamma \in \Gamma_a$, and connecting them with an extra node t. t is connected via an a-edge to t^γ from each of the N^γ, and is also connected via a b-edge to each b-successor of s in M. We must also have a-edges between each of the t^γ to ensure that the resulting model is transitive and Euclidean. N_t is then the desired a-refinement of M_s.

relation R_a^M is disjoint from the other relations, and so can be considered in isolation. By hypothesis we have assumed that $t^\gamma R_a^{N^\gamma} = \{t^\gamma\}$, so as there are no relationships from any t^γ to a different state in S^{N^γ}, we can also consider the each $R_a^{N^\gamma}$ separately from the rest of R_a^N. We note that as R_a^M and each $R_a^{N^\gamma}$ are transitive and Euclidean, and as the remainder of the relation is essentially an equivalence relation without the relationship (t, t), it also has transitivity and Euclideaness. Therefore R_a^N is transitive and Euclidean. We can show that R_b^N is also transitive and Euclidean from the fact that R_b^M and each $R_b^{N^\gamma}$ is transitive and Euclidean, and otherwise disjoint from each other, and the remainder of R_b^N is simply a duplicate of relationships from s to its successors, then R_b^N is also transitive and Euclidean. Therefore N is a doxastic model.

We construct an a-simulation $\mathcal{R}$ from N_t to M_s, where:

$$\mathcal{R} = \{(t, s)\} \cup \{(s', s') \mid s' \in S^M\} \cup \bigcup_{\gamma \in \Gamma_a} \mathcal{R}^\gamma$$

We must show that $\mathcal{R}$ satisfies **atoms**, **forth-**b for every $b \in A$, and **back-**b for every $b \in A - \{a\}$.

atoms We note that, by construction, the valuation of N matches the valuation of its corresponding states in M and each N^γ, and the valuation of N_t matches that of M_s. Therefore $\mathcal{R}$ satisfies **atoms**.

forth We next show that $\mathcal{R}$ satisfies **forth-**b for every $b \in A$. Let $b \in A$ and let $(u, v) \in \mathcal{R}$.

Suppose that $(u, v) \in \mathcal{R}^\gamma$ for some $\gamma \in \Gamma_a$. Suppose further that $b = a$ and $u = t^\gamma$ for some $\gamma \in \Gamma_a$. Then $t^\gamma R_a^N = tR_a^N = \{t^{\gamma'} \mid \gamma' \in \Gamma_a\}$. As M is a doxastic model, we have that $s^\gamma R_a^M = sR_a^M = \{s^{\gamma'} \mid \gamma' \in \Gamma_a\}$, and as each $\mathcal{R}^{\gamma'}$ is an a-simulation between $N^{\gamma'}_{t^{\gamma'}}$ and $M_{s^{\gamma'}}$, we have that for every $\gamma' \in \Gamma_a$, $(t^{\gamma'}, s^{\gamma'}) \in \mathcal{R}^{\gamma'} \subseteq \mathcal{R}$. Otherwise consider any other $(u, v) \in \mathcal{R}^\gamma$. Then as $\mathcal{R}^\gamma$ is an a-simulation, it satisfies **forth-**b for every $b \in A$. Hence for every $u' \in uR_b^{N^\gamma}$, there exists some $v' \in vR_b^M$ such that $(u', v') \in \mathcal{R}^\gamma \subseteq \mathcal{R}$. Suppose that $b = a$ and $u = t^\gamma$ for some $\gamma \in \Gamma_a$.

Suppose instead that $(u, v) = (s', s')$ for some $s' \in S^M$. Then we note that $s'R_b^N = s'R_b^M$, and hence for every $s'' \in s'R_b^N$ we have that $s'' \in s'R_b^M$, and that $(s'', s'') \in \mathcal{R}$.

Finally suppose that $(u, u') = (t, s)$. Then suppose that $b = a$. By construction, $tR_a^N = \{t^\gamma \mid \gamma \in \Gamma_a\}$, and hence $v = t^\gamma$ for some $\gamma \in \Gamma_a$. Hence we can take $s^\gamma \in sR_a^M$, and note that as $\mathcal{R}^\gamma$ is an a-simulation from M_{s^γ} to $N^\gamma_{t^\gamma}$, we know that $(t^\gamma, s^\gamma) \in \mathcal{R}^\gamma \subseteq \mathcal{R}$. Suppose that $b \neq a$. Then by construction, $tR_b^M = sR_b^M$, hence for every $t' \in tR_b^M$, we have that $t' \in sR_b^M$, and hence we know that $(t', t') \in \mathcal{R}$. Hence $\mathcal{R}$ satisfies **back-**b for every $b \in A$. matches that of M_s. Therefore $\mathcal{R}$ satisfies **atoms**.

back A similar argument to above shows that $\mathcal{R}$ satisfies **back-**b for every $b \in A - \{a\}$.

Therefore $\mathcal{R}$ is an a-simulation, and $N_t \preceq_a M_s$.

Finally we show that $N_t \models \nabla_a \Gamma_a$. We must show that for each $\gamma \in \Gamma_a$ that $N_{t^\gamma} \models \gamma$. As each γ is an a-alternating disjunctive normal formula, a similar argument to that used in Lemma 6.1.4 can be used to show that each of the successors of N_{t^γ} are bisimilar to the corresponding successors of $N^\gamma_{t^\gamma}$. This is obvious, as N contains a duplicate of N^γ, and N does not contain any additional edges originating from states in S^{N^γ}, except for a-edges from t^γ. Hence we have that $N_{t^\gamma} \models \gamma$ for every $\gamma \in \Gamma_a$, and hence $N_t \models \nabla_a \Gamma$.

As $N_t \preceq_a M_s$, and $N_t \models \nabla_a \Gamma_a$ we therefore have that $M_s \models \triangleright_a \nabla_a \Gamma_a$.

Conversely, suppose that $M_s \models \nabla_a \Gamma_a$. Then there exists a doxastic model $N_t \preceq_a M_s$, via some a-simulation $\mathcal{R}$, such that $N_t \models \nabla_a \Gamma_a$. From the definition of the cover operator, this implies that $N_t \models \Box_a \bigvee_{\gamma \in \Gamma_a} \gamma \wedge \bigwedge_{\gamma \in \Gamma_a} \Diamond_a \gamma$. In particular we note that for every $\gamma \in \Gamma_a$, $N_t \models \Diamond_a \gamma$, and so there exists some $t^\gamma \in tR_a^N$ such that $N_{t^\gamma} \models \gamma$. As $t^\gamma \in tR_a^N$, and $(t, s) \in \mathcal{R}$, by **forth-**a there exists some $s^\gamma \in sR_a^M$ such that $(t^\gamma, s^\gamma) \in \mathcal{R}$. Hence $\mathcal{R}$ is also an a-simulation from N_{t^γ} to M_{s^γ}, and so $M_{s^\gamma} \models \triangleright_a \gamma$. As for every $\gamma \in \Gamma_a$ we have that $s^\gamma \in sR_a^M$, we also have that $M_s \models \Diamond_a \triangleright_a \gamma$. Therefore we finally have that $M_s \models \bigwedge_{\gamma \in \Gamma_a} \Diamond_a \triangleright_a \gamma$.

Therefore **RKD45** is sound.

RComm Suppose that $M_s \in KD45$ is a doxastic model such that $M_s \models \nabla_b \{\triangleright_a \gamma \mid \gamma \in \Gamma_b\}$, where $a \neq b$ and Γ_b is a non-empty set of b-alternating disjunctive normal formulae.

We need to show that $M_s \models \triangleright_a \nabla_b \Gamma_b$. To do this we follow the same strategy as for proving **RKD45**: we construct an a-refinement $N_t \in KD45$, and show that $N_t \models \nabla_b \Gamma$.

We begin by constructing the model N_t. Consider $\gamma \in \Gamma$. From $M_s \models \nabla_b \{\triangleright_a \gamma \mid \gamma \in \Gamma_b\}$, there exists a state $s^\gamma \in sR_b^M$ such that $M_{s^\gamma} \models \triangleright_a \gamma$. Therefore there exists a doxastic model $N^\gamma_{t^\gamma} \preceq_a M_{s^\gamma}$, via some a-simulation $\mathcal{R}^\gamma$, such that $N^\gamma_{t^\gamma} \models \gamma$. Without loss of generality we assume that the N^γ are disjoint. We may also assume, by Lemma 6.1.4, that for every $\gamma \in \Gamma$, $t^\gamma R_b^{N^\gamma} = R_b^{N^\gamma} t^\gamma = \{t^\gamma\}$ and that $R_c^{N^\gamma} t^\gamma = \emptyset$ for every $c \in A - \{b\}$.

Let t be a state not in S^M or any of the S^{N^γ}. Then we construct a Kripke model $N = (S^N, R^N, V^N)$ where:

$$\begin{aligned}
S^N &= \{t\} \cup S^M \bigcup_{\gamma \in \Gamma_b} S^{N^\gamma} \\
R_b^N &= \{(t, t^\gamma) \mid \gamma \in \Gamma_b\} \cup \{(t^\gamma, t^{\gamma'}) \mid \gamma, \gamma' \in \Gamma_b\} \cup R_b^M \cup \bigcup_{\gamma \in \Gamma_b} R_b^{N^\gamma} \\
R_c^N &= \{(t, t') \mid t' \in sR_c^M\} \cup R_c^M \cup \bigcup_{\gamma \in \Gamma_b} R_c^{N^\gamma} \text{ for } c \in A - \{b\} \\
V^N(p) &= \begin{cases} \{t\} \cup V^M(p) \cup \bigcup_{\gamma \in \Gamma_b} V^{N^\gamma}(p) & \text{if } s \in V^M(p) \\ V^M(p) \cup \bigcup_{\gamma \in \Gamma_b} V^{N^\gamma}(p) & \text{otherwise} \end{cases} \quad \text{for } p \in P
\end{aligned}$$

We note that N is a doxastic model, by similar arguments as used in the proof for **RKD45**.

We construct an a-simulation $\mathcal{R}$ from N_t to M_s, where:

$$\mathcal{R} = \{(t, s)\} \cup \{(s', s') \mid s' \in S^M\} \cup \bigcup_{\gamma \in \Gamma_b} \mathcal{R}^\gamma$$

We note that $\mathcal{R}$ is an a-simulation, by similar arguments as used in the proof for **RKD45**. In particular, this means that $N_t \preceq_a M_s$.

We also note that for every $\gamma \in \Gamma_b$, that $N_{t^\gamma} \vDash \gamma$, by similar arguments as used in the proof for **RKD45**. In particular, this means that $N_t^\gamma \vDash \nabla_b \Gamma_b$.

Therefore $M_s \vDash \triangleright_a \nabla_b \Gamma_b$.

The converse, $\triangleright_a \nabla_b \Gamma_b \to \nabla_b \{\triangleright_a \gamma \mid \gamma \in \Gamma_b\}$ follows a similar proof to the relevant part in the proof for **RKD45**.

Therefore **RComm** is sound.

RDist Suppose that M_s is a doxastic model such that $M_s \vDash \bigwedge_{b \in B} \triangleright_a \nabla_b \Gamma_b$, where $B \subseteq A$, and Γ_b is a set of b-disjunctive formulae for each $b \in B$.

We need to show that $M_s \vDash \triangleright_a \bigwedge_{b \in B} \nabla_b \Gamma_b$. To do this we follow the same strategy as for proving **RKD45**: we construct an a-refinement $N_t \in KD45$, and show that $N_t \vDash \triangleright_a \bigwedge_{b \in B} \nabla_b \Gamma_b$.

We begin by constructing the model N_t. Suppose that $a \in B$. Then we have $M_s \models \rhd_a \nabla_a \Gamma_a$, and by **RKD45** this implies that $M_s \models \bigwedge_{\gamma \in \Gamma_a} \gamma$. We also have that for every $b \in B - \{a\}$ that $M_s \models \rhd_a \nabla_a \Gamma_b$, and by **RComm** this implies that $M_s \models \nabla_b \{\rhd_a \gamma \mid \gamma \in \Gamma_b\}$, and by the definition of the cover operator, this implies that $M_s \models \bigwedge_{\gamma \in \Gamma_b} \Diamond_b \rhd_a \gamma$. Hence for every $b \in B$ and $\gamma \in \Gamma_b$, we have that $\Diamond_b \rhd_a \gamma$. This implies that for each $b \in B$ and each $\gamma \in \Gamma_b$ that there exists some $s^{b,\gamma} \in sR_b^M$ such that $M_{s^{a,\gamma}} \models \rhd_a \gamma$. Therefore there exists a doxastic model $N^{b,\gamma}_{t^{b,\gamma}} \preceq_a M_{s^{a,\gamma}}$ such that $N^{b,\gamma}_{t^{b,\gamma}} \models \gamma$. Without loss of generality we may assume that the $N^{b,\gamma}$ are disjoint. We may also assume, by Lemma 6.1.4, that for every $b \in B$ and $\gamma \in \Gamma$, that $t^\gamma R_b^{N^\gamma} = R_b^{N^\gamma} t^\gamma = \{t^\gamma\}$ and that $R_c^{N^\gamma} t^\gamma = \emptyset$ for every $c \in A - \{b\}$.

Let t be a state not in S^M or any of the $S^{N^{b,\gamma}}$. Then we construct a Kripke model $N = (S^N, R^N, V^N)$ where:

$$
\begin{aligned}
S^N &= \{t\} \cup \bigcup_{b \in A, \gamma \in \Gamma_b} S^{N^{b,\gamma}} \\
R_b^N &= \{(t, t^{b,\gamma}) \mid b \in A, \gamma \in \Gamma_b\} \cup \{(t^{b,\gamma}, t^{b,\gamma'}) \mid \gamma, \gamma' \in \Gamma_b\} \\
&\quad \cup \bigcup_{c \in A, \gamma \in \Gamma_b} R_b^{N^{c,\gamma}} \text{ for } b \in A \\
R_b^N &= \{(t, t') \mid t' \in sR_b^M\} \cup R_b^M \cup \bigcup_{c \in B, \gamma \in \Gamma_c} R_b^{N^{c,\gamma}} \text{ for } b \in A \setminus B \\
V^N(p) &= \begin{cases} \{t\} \cup \bigcup_{b \in A, \gamma \in \Gamma_b} V^{N^{b,\gamma}}(p) & \text{if } s \in V^M(p) \\ \bigcup_{b \in A, \gamma \in \Gamma_b} V^{N^{b,\gamma}}(p) & \text{otherwise} \end{cases}
\end{aligned}
$$

We note that N is a doxastic model, by similar arguments as used in the proof for **RKD45**.

We construct an a-simulation $\mathcal{R}$ from N_t to M_s, where:

$$
\mathcal{R} = \{(t, s)\} \cup \bigcup_{b \in A, \gamma \in \Gamma_b} \mathcal{R}^\gamma
$$

We note that this is an a-simulation, by similar arguments as used in the proof for **RKD45**. In particular, this means that $N_t \preceq_a M_s$.

We also note that for every $\gamma \in \Gamma_b$, that $N_{t^\gamma} \models \gamma$, by similar arguments as used in the proof for **RKD45**. In particular, this means that $N_t^\gamma \models \nabla_b \Gamma_b$.

Therefore $M_s \models \rhd_a \bigwedge_{b \in A} \nabla_b \Gamma_b$ and **RDist** is sound.

Therefore the axiomatisation $\mathbf{RML^{KD45}}$ is sound. □

We note that, similar to $\mathbf{RML^K}$, the converse of **RDist** is derivable within $\mathbf{RML^{KD45}}$.

Lemma 6.2.2. *The following is derivable in* $\mathbf{RML^{KD45}}$*.*

$$\vdash \bigwedge_{b\in A} \rhd_a \nabla_b \Gamma_b \leftrightarrow \rhd_a \bigwedge_{b\in A} \nabla_b \Gamma_b$$

where Γ_b *is a set of* b*-alternating disjunctive normal formulae for every* $b \in A$*.*

We show the completeness of the axiomatisation $\mathbf{RML^{KD45}}$ in a similar fashion to the completeness proof of $\mathbf{RML^K}$, by a provably correct translation from $\mathcal{L}_\rhd$ to $\mathcal{L}$. Completeness then follows from the completeness of $\mathcal{L}^{\mathbf{KD45}}$.

As for the completeness proof for $\mathbf{RML^K}$, we introduce some similar equivalences that will be used by our translation.

Lemma 6.2.3. *The following are provable equivalences using* $\mathbf{RML^{KD45}}$*:*

1. $\rhd_a(\phi \vee \psi) \leftrightarrow \rhd_a\phi \vee \rhd_a\psi$
2. $\rhd_a(\pi \wedge \bigwedge_{b\in B} \nabla_b\Gamma_b) \leftrightarrow \pi \wedge \bigwedge_{\gamma\in\Gamma_a} \Diamond_a\rhd_a\gamma \wedge \bigwedge_{b\in B} \nabla_b\{\rhd_a\gamma \mid \gamma \in \Gamma_b\}$ *where* π *is propositional,* $B \subseteq A$*, and* $a \in B$
3. $\rhd_a(\pi \wedge \bigwedge_{b\in B} \nabla_b\Gamma_b) \leftrightarrow \pi \wedge \bigwedge_{b\in B} \nabla_b\{\rhd_a\gamma \mid \gamma \in \Gamma_b\}$ *where* π *is propositional,* $B \subseteq A$*, and* $a \notin B$

where for every $a \in A$*, the set* Γ_a *is a non-empty set of* a*-alternating disjunctive normal formulae.*

This can be shown by following similar reasoning as used for the proof of Lemma 5.2.3, but by substituting $\mathbf{RML^K}$ axioms for $\mathbf{RML^{KD45}}$ axioms.

Lemma 6.2.4. *Every formula of* $\mathcal{L}_\rhd^{\mathbf{KD45}}$ *is provably equivalent to a formula of* $\mathcal{L}^{\mathbf{KD45}}$*.*

This can be shown following similar reasoning as used for the proof of Lemma 5.2.4. We note that, of course, we use the appropriate axioms of $\mathbf{RML^{KD45}}$ in place of those from $\mathbf{RML^K}$, and we also use the equivalences from Lemma 6.2.3 in place of those from Lemma 5.2.3. We also use the assumption that formulae are in the cover logic alternating disjunctive normal form, rather than the disjunctive normal form.

The rest of the completeness proof follows the same reasoning as used for the completeness proof of $\mathbf{RML^K}$.

Corollary. *Let $\phi \in \mathcal{L}_{\triangleright}^{\mathbf{KD45}}$ be given and $\psi \in \mathcal{L}^{\mathbf{KD45}}$ be semantically equivalent to ϕ. If ψ is a theorem in $\mathcal{L}^{\mathbf{KD45}}$, then ϕ is a theorem in $\mathbf{RML}^{\mathbf{KD45}}$.*

Lemma 6.2.5. *The axiom schema $\mathbf{RML}^{\mathbf{KD45}}$ is complete with respect to the semantic class KD45.*

We note that Corollary 6.2 and Lemma 6.2.5 can be shown by following similar reasoning as used for the proofs of Corollary 5.2 and Lemma 5.2.5 respectively.

Theorem 6.2.6. *The axiomatisation $\mathbf{RML}^{\mathbf{KD45}}$ is sound and complete with respect to the semantic class KD45.*

Proof. The soundness proof is given in Lemma 6.2.1 and the completeness proof is given in Lemma 6.2.5. $\square$

We note that, as in the axiomatisation for the single-agent epistemic and doxastic logics, the completeness proofs above were performed with a provably correct translation from $\mathcal{L}_{\triangleright}$ to $\mathcal{L}$, under the semantics of $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$. This shows that $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$ is expressively equivalent to $\mathcal{L}^{\mathbf{KD45}}$, and allows us to show several results. In particular, that $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$ is decidable.

Theorem 6.2.7. *The logic $\mathcal{L}_{\triangleright}^{\mathbf{KD45}}$ is decidable.*

This can be shown by following similar reasoning as used for the proof of Theorem 4.2.9.

CHAPTER 7

Conclusions

In this thesis we have provided sound and complete axiomatisations for the single-agent refinement quantified epistemic and doxastic logics, and for the multi-agent refinement quantified modal and doxastic logics. The completeness proofs for each of these axiomatisations is performed via a provably correct translation from the refinement quantified modal logics to the basic modal logics, thus showing that each of the refinement quantified versions of the logics we have considered are expressively equivalent to the basic modal versions. This allows us to prove a number of properties about these refinement quantified modal logics from the properties of the corresponding basic modal logics, including decidability and the finite model property, and allows us to derive model-checking and decision procedures for these logics via the translations.

The logics presented in this thesis do not lend themselves directly to practical applications, however results in logics based on those we have considered may conceivably see practical uses in formal verification of software in the future. The logics that we have considered quantify over all possible informative updates, whilst practical applications often call for restrictions on the informative updates that are possible. For example, in the setting of a game, the moves that can and cannot be made are dictated by the rules of the game, or in the setting of a security protocol, the communication that occurs is dictated by the protocol, and thus the information that may be gleaned from these moves or communication is accordingly restricted.

Furthermore, the notion of informative updates captured by refinements is better suited to epistemic settings than to general modal or doxastic settings, and so the omission of a result in the multi-agent epistemic setting is significant. Although refinements can correspond to notions other than informative updates, the correspondence to informative updates is our primary motivation. Extensions of these logics to other settings, such as to epistemic logic, may borrow from the techniques used in the results we have presented in this thesis, particularly the cover logic normal forms and the technique of a provably correct translation to basic modal logics. However whether these techniques are applicable or not depend on the setting that we are considering; for example van Ditmarsch, Pinchinat and French [23] have already demonstrated that adding refinement quantifiers to the modal logic **K4** results in a logic that is strictly more expressive, and thus the technique of a provably correct translation to the corresponding basic modal logic is not feasible.

There are several immediate avenues for future work based on the results presented in this thesis. We have not considered efficient decision or model-checking procedures, succinctness results, or the effects of adding common knowledge operators to the logics. Also we have already mentioned that results in the multi-agent refinement quantified epistemic logic have yet to be given.

The provably correct translation used in the completeness proofs can be used to derive a decision procedure for the logics we have considered. However, we note that such a decision procedure would have a non-elementary time complexity. This complexity is due to the translation to the various normal forms that we have considered. Conversion to prenex normal form, disjunctive normal form, or alternating disjunctive normal form results in a formula that is exponentially larger than the original in the worst-case. As this conversion must be performed for each refinement quantifier in the given formula, a formula with nested refinement quantifiers results in an exponential increase in formula size for each nested quantifier, thus the overall size of the resulting formula is non-elementary in the size of the original formula. A more efficient decision procedure would be desirable.

A decision procedure based on a tableau method was provided by van Ditmarsch, French and Pinchinat [23] for the single-agent refinement quantified modal logic, that runs in 2EXP time in the worst case. A modification of this tableau method may be applicable to the multi-agent refinement quantified epistemic and doxastic logics, and is likely to perform better than the non-elementary translation that we already have. A decision procedure for the single-agent refinement quantified epistemic and doxastics logics that runs in 2EXP time has already been presented by Hales, French and Davies [12].

Efficient model-checking procedures for these logics have yet to be considered. Model-checking procedures can easily be derived from any decision procedure that we may have, however it is often the case that there are model-checking procedures for a logic that are more efficient than the corresponding decision procedures. It is unknown whether there are model-checking procedures for any of the logics considered so far, that are more efficient than the decision procedures.

The provably correct translation used in the completeness proofs can also be used to show that the refinement quantified modal logics we have considered are expressively equivalent to their corresponding basic modal logics, so we raise the question of whether the refinement quantified modal logics are more succinct than the basic modal logics. It was shown by van Ditmarsch, French and Pinchinat [23] that the multi-agent refinement quantified modal logic is exponentially more succinct than the basic modal logic. Succinctness results for the refinement quantified doxastic and epistemic logics have yet to be considered.

We may also wish to consider the addition of common knowledge operators to the refinement quantified modal logics. In an epsitemic setting, common knowledge refers to a situation where a group of agents knows some fact ϕ, and every agent in that group knows that the group knows that ϕ, and so on ad infinitum. A common question in dynamic epistemic logic is whether certain knowledge can become common knowledge for a group of agents; common knowledge tends to be considerably more difficult for agents to obtain than normal knowledge. Common knowledge operators are used to signify that a group of agents holds certain knowledge as common knowledge. Combined with refinement quantifiers, this would allow us to pose questions such as whether or not certain common knowledge is attainable.

We finally remark on the challenges involved in an axiomatisation for the multi-agent refinement quantified epistemic logic. We can recognise the axiomatisations for the multi-agent refinement quantified modal and doxastic logics as generalisations of the axiomatisations for the single-agent variants of these logics. In each case, we generalise the **RK** or **RKD45** axiom, and introduce additional axioms, **RComm** and **RDist** to deal with the interactions between different agents in the logics. These axiomatisations have been designed so that the proof strategy of a provably correct translation is feasible in order to show that the axiomatisations are complete. Consequently the soundness proofs for these axioms have been more challenging than the completeness proofs.

We have briefly considered an axiomatisation for the multi-agent refinement quantified epistemic logic. The obvious generalisations however do not appear to be sound. In the cases for K and $KD45$, the soundness proofs for our new axioms involve taking refinements in successor states, each of which satisfy a particular formula, and combining these refinements into a larger model where the corresponding successor states continue to satisfy the corresponding formulae. The main obstacle in the epistemic setting is that we do not yet have a result that allows us to safely combine such refinements into a larger model whilst the successor states continue to satisfy the formulae we require them to. In the case of K, this was trivial, as we combine the refinements using inward edges only, thus the successor states are bisimilar to the refinements they were constructed from, and thus satisfy the same formulae. In the case of $KD45$, we relied on the properties of alternating disjunctive normal formulae, in particular

Lemma 6.1.4, which allowed us to safely combine the refinements in the presence of the transitive and Euclidean edges. Whilst every $\mathcal{L}^{\mathbf{S5}}$ formula is equivalent to a formula in alternating disjunctive normal form, Lemma 6.1.4 does not hold in $\mathcal{L}^{\mathbf{S5}}$. This is due to a combination of the reflexivity and symmetric properties of *S5* models. It is because of these properties, particularly the symmetric property, that we cannot expect models to continue to satisfy the formulae we require them to satisfy when we combine them into larger models.

A modification of the alternating disjunctive normal form may lead to a method to safely combine the refinements in the manner we require, and thus give us a provably correct translation. On the other hand, it may be the case that the refinement quantified epistemic logic is more expressive than epistemic logic, in which case the technique of a provably correct translation would be infeasible. Previously Baltag, Solecki and Moss [5] gave a provably correct translation from the action model logic to epistemic logic; we note that if we had a provably correct translation from the refinement quantified epistemic logic to epistemic logic, then we could combine these two translations to give a translation from the arbitrary action model logic to epistemic logic.

APPENDIX A

Modal, doxastic and epistemic logics

In this appendix we give a brief introduction to modal, doxastic and epistemic logics, the required background material for this paper.

To introduce modal logic, we will begin with the definition of a Kripke model.

Let A be a non-empty, finite set of agents, and let P be a non-empty, countable set of propositional atoms.

Definition A.0.2 (Kripke model)**.** A *Kripke model* $M = (S, R, V)$ consists of a *domain* S, which is a set of states (or worlds), an *accessibility relation* $R : A \to \mathcal{P}(S \times S)$, and a *valuation* $V : P \to \mathcal{P}(S)$.

The class of all Kripke models is called K. We write $M \in K$ to denote that M is a Kripke model.

For $R(a)$, we write R_a. We write sR_a for $\{t \mid (s,t) \in R_a\}$ and we write R_at for $\{s \mid (s,t) \in R_a\}$. As we will be required to discuss several models at once, we will use the convention that $M = (S^M, R^M, V^M)$, $N = (S^N, R^N, V^N)$, and so on. For $s \in S^M$ we will let M_s refer to the pair (M, s), also known as the pointed Kripke model of M at state s.

Modal logics can be said to model notions such as necessity, knowledge or belief. To help with the introduction however, we will for now adopt an analogy using knowledge. Under the *possible worlds* interpretation of modal logic, we are concerned with a collection of agents, and each agent can be said to live in a world where there are a certain set of facts that are true or false. The agent may not have complete knowledge about which of those facts are true or false, and so it may consider several possibilities. We say that an agent considers several possible worlds, where the truth of each fact may be different in each hypothetical world. Thus each agent considers a set of worlds to be possible, and we say that the agent *knows* a statement about the world to be true if that statement is true in all of the worlds that it considers possible. For example, if an agent knew everything that there was to know about the world, then it would consider only one world possible: the actual world; whereas if an agent was uncertain about a particular statement ϕ about the world, then the agent would consider at least two worlds possible: a world where ϕ is true and a world where ϕ is false.

Kripke models are an encoding of this possible worlds interpretation. Each state of the Kripke model represents a world, and the facts that are true at that world are represented by the propositional atoms that are true at the state, according to the valuation of the Kripke model. At each state, the worlds that an agent considers possible are those worlds that are accessible from that state according to that agent's accessibility relation.

It should be emphasised that the worlds that an agent considers possible are not a global property of the Kripke model, but rather the possible worlds can vary from state to state. This is reflected in the fact that the possible worlds are represented by the states accessible from a particular state, rather than by a fixed subset of states of the Kripke model. Furthermore it should be noted that the statements about the world that an agent can know about is not limited just to statements directly about the facts of the world; agents may be aware of the knowledge of other agents.

In the basic modal logic, we usually refer to the notion of necessity. Thus if a statement is true in all of the worlds that an agent considers possible, then the agent considers that statement to be necessary. We will explain the difference between necessity, knowledge and belief later; for the moment the difference is not important.

We will now give a precise definition of the language, $\mathcal{L}$, of basic modal logic, and its semantics over a general class of Kripke models.

Definition A.0.3 (Language of $\mathcal{L}$)**.** Given a non-empty, finite set of agents A and a non-empty, countable set of propositional atoms P, the language of $\mathcal{L}$ is defined by the following abstract syntax:

$$\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \Box_a \alpha$$

where $p \in P$, $a \in A$ and $\alpha \in \mathcal{L}$.

Standard abbreviations include: $\top ::= p \vee \neg p$ for some $p \in P$; $\bot ::= \neg\top$; $\phi \vee \psi ::= \neg(\neg\phi \wedge \neg\psi)$; $\phi \rightarrow \psi ::= \neg\phi \vee \psi$; and $\Diamond_a \phi ::= \neg\Box_a\neg\phi$.

The symbol $\Box_a$ is said to represent necessity in modal logics. Thus we may read $\Box_a\phi$ as "a considers ϕ necessary". In epistemic logic, a variant of modal logic, $\Box_a$ is said to represent knowledge, and in doxastic logic, $\Box_a$ is said to represent belief; in these settings we may read $\Box_a\phi$ as "a knows ϕ" or "a believes ϕ". The dual operator, $\Diamond_a$ represents possibility; hopefully it is intuitive that if a statement is not considered to be necessarily false, then it must be considered possible. Similarly, if a statement is not known to be false, or if it is not believed to be false, then it must be considered possible. Thus we may read $\Diamond_a\phi$ as "a considers ϕ possible".

Definition A.0.4 (Semantics of $\mathcal{L}^{\mathbf{C}}$)**.** Let C be a class of Kripke models, and let $M = (S, R, V) \in C$ be a Kripke model taken from C. The interpretation of ϕ in a pointed Kripke model M_s is defined inductively:

$$\begin{array}{lll} M_s & \vDash & p \text{ iff } s \in V_p \\ M_s & \vDash & \neg\phi \text{ iff } M_s \nvDash \phi \\ M_s & \vDash & \phi \wedge \psi \text{ iff } M_s \vDash \phi \text{ and } M_s \vDash \psi \\ M_s & \vDash & \Box_a\phi \text{ if for all } t \in S : (s,t) \in R_a \text{ implies } M_t \vDash \phi \end{array}$$

We say that a formula ϕ is *satisfied* by a pointed Kripke model $M_s \in C$ if and only if $M_s \vDash \phi$. We say that ϕ is satisfied by a Kripke model $M \in C$ if and only if $M_s \vDash \phi$ for some $s \in S^M$. We say that ϕ is *satisfied* by a class of Kripke models C if and only if it is satisfied by every Kripke model $M \in C$. We say that ϕ is *valid* in a Kripke model $M \in C$ if and only if $M_s \vDash \phi$ for every $s \in S^M$. We write $M \vDash \phi$. We say that ϕ is *valid* in a class of Kripke models C if and only if $M \vDash \phi$ for every $M \in C$. We write $\vDash_C \phi$.

We take a moment to remark on the interpretation of the modal logic. We note that modal formulae are interpreted in pointed Kripke models. In our possible worlds interpretation, a pointed Kripke model M_s can be said to be a Kripke model where we have designated a particular state, s, as the actual or current world. Thus we may consider all other worlds to be purely hypothetical. A propositional variable is interpreted with respect to the valuation at the current state of the Kripke model. Negation and conjunction are interpreted in the intuitive manner, both at the current state. Probably the most unfamiliar aspect of the interpretation of modal logic is the interpretation of the $\Box_a$ operator. The $\Box_a$ operator effectively moves consideration from the current state of the pointed Kripke model to each of the states accessible from the current state, under the Kripke model's accessibility relation. This captures the notion that we described earlier, that an agent considers a statement to be necessary if that statement is true in all of the worlds that it considers possible. It should be clear that under this interpretation, the dual operator, $\Diamond_a$, also captures the notion that an agent considers a statement to be possible if that statement is true in at least one of the worlds that it considers possible.

Modal logics are interpreted over classes of Kripke models. The simplest normal modal logic, $\mathcal{L}^{\mathbf{K}}$, is interpreted over the class of all Kripke models, K. Variants of modal logic are interpreted over subclasses of K, which impose constraints on the accessibility relations of the Kripke models. Doxastic logic, $\mathcal{L}^{\mathbf{KD45}}$, and epistemic logic, $\mathcal{L}^{\mathbf{S5}}$, are two variants of modal logic that are interpreted over different classes of epistemic models: the class of doxastic models, *KD45*, and the class of epistemic models, *S5*, respectively. As we will see later, the constraints that are imposed on these classes of models gives us properties that make doxastic and epistemic logic represent simple notions that we have about belief and knowledge.

We will now define doxastic and epistemic models, first defining the relational properties that we will use to describe them.

Definition A.0.5 (Relational properties)**.** Let S be a set and let $R \subseteq S \times S$. Then we define the following properties of R.

1. R is *serial* if and only if for every $s \in S$ there exists some $t \in S$ such that $(s,t) \in R$.
2. R is *reflexive* if and only if for every $s \in S$ we have that $(s,s) \in R$.
3. R is *transitive* if and only if for every $(s,t),(t,r) \in R$ we also have that $(s,r) \in R$.
4. R is *Euclidean* if and only if for every $(s,t),(s,r) \in R$ we also have that $(t,r) \in R$.
5. R is *symmetric* if and only if for every $(s,t) \in R$ we also have that $(t,s) \in R$.

We note that if R is reflexive then it is also serial. We also note that R is reflexive and Euclidean, if and only if it is reflexive, transitive and symmetric.

Definition A.0.6 (Doxastic model)**.** A *doxastic model* is a Kripke model $M = (S,R,V)$ such that the relation R_a is serial, transitive, and Euclidean for all $a \in A$. The class of all doxastic models is called *KD45*. We write $M \in$ *KD45* to denote that M is a doxastic model.

Definition A.0.7 (Epistemic model)**.** An *epistemic model* is a Kripke model $M = (S,R,V)$ such that the relation R_a is reflexive, transitive and symmetric for all $a \in A$. The class of all epistemic models is called *S5*. We write $M \in$ *S5* to denote that M is an epistemic model.

We note that every epistemic model is also a doxastic model, as if a relation R_a is reflexive, transitive and symmetric, then it is also serial and Euclidean, and therefore satisfies the constraints of a doxastic model.

The logics $\mathcal{L}^{\mathbf{K}}$, $\mathcal{L}^{\mathbf{KD45}}$ and $\mathcal{L}^{\mathbf{S5}}$ are instances of $\mathcal{L}^{\mathbf{C}}$ with classes K, *KD45* and *S5* respectively. It should be noted that $\mathcal{L}^{\mathbf{KD45}}$ is a conservative extension of $\mathcal{L}^{\mathbf{K}}$, and $\mathcal{L}^{\mathbf{S5}}$ is a conservative extension of $\mathcal{L}^{\mathbf{KD45}}$ (and also of $\mathcal{L}^{\mathbf{K}}$). This means that every valid formula in $\mathcal{L}^{\mathbf{K}}$ is also valid in $\mathcal{L}^{\mathbf{KD45}}$, and likewise for $\mathcal{L}^{\mathbf{KD45}}$ and $\mathcal{L}^{\mathbf{S5}}$. This is because any formula that is valid with respect to a particular class of Kripke models is also valid for any subclass of those Kripke models.

Proposition A.0.8 (Properties of $\mathcal{L}^{\mathbf{K}}$, $\mathcal{L}^{\mathbf{KD45}}$ and $\mathcal{L}^{\mathbf{S5}}$)**.** *Let $\phi, \psi \in \mathcal{L}$. Then the following are valid:*

1. $\models_K \Box_a(\phi \to \psi) \to \Box_a\phi \to \Box_a\psi$
2. $\models_K \Box_a\phi \land \Box_a\psi \leftrightarrow \Box_a(\phi \land \psi)$
3. $\models_K \Diamond_a\phi \lor \Diamond_a\psi \leftrightarrow \Diamond_a(\phi \lor \psi)$
4. $\models_{KD45} \Box_a\phi \to \Diamond_a\phi$
5. $\models_{KD45} \Box_a\phi \to \Box_a\Box_a\phi$
6. $\models_{KD45} \neg\Box_a\phi \to \Box_a\neg\Box_a\phi$
7. $\models_{S5} \Box_a\phi \to \phi$

Proof.

1. Let $M_s \in K$ such that $M_s \models \Box_a(\phi \to \psi)$ and $M_s \models \Box_a\phi$. Then for every $t \in sR_a$, we have $M_t \models \phi \to \psi$ and $M_t \models \phi$. Therefore for every $t \in sR_a$, we have $M_t \models \psi$, and therefore we have that $M_s \models \Box_a\psi$.
2. Exercise.
3. Exercise.
4. Let $M_s \in KD45$ such that $M_s \models \Box_a\phi$. Then for every $t \in sR_a$, we have $M_t \models \phi$. As $M_s \in KD45$, we have that R_a is serial. Therefore $sR_a \neq \emptyset$, and so there exists some $t \in sR_a$ such that $M_t \models \phi$. Therefore $\Diamond_a\phi$.
5. Let $M_s \in KD45$ such that $M_s \models \Box_a\phi$. Then consider $t \in sR_a$ and $t' \in tR_a$. As $M_s \in KD45$, we have that R_a is transitive. Therefore $t' \in sR_a$, and so $M_{t'} \models \phi$. Therefore $M_t \models \Box_a\phi$, and so $M_s \models \Box_a\Box_a\phi$.
6. Let $M_s \in KD45$ such that $M_s \models \neg\Box_a\phi$. Then there exists some $t \in sR_a$ such that $M_t \models \neg\phi$. Then consider $t' \in sR_a$. As $M_s \in KD45$, we have that R_a is Euclidean. Therefore $t \in t'R_a$, and so $M_t \models \neg\Box_a\phi$. Therefore $M_s \models \Box_a\neg\Box_a\phi$.

7. Let $M_s \in S5$ such that $M_s \models \Box_a \phi$. Then for every $t \in sR_a$, we have $M_t \models \phi$. As $M_s \in S5$, we have that R_a is reflexive. Therefore $s \in sR_a$, and so $M_s \models \phi$.

□

We note that the properties described in Proposition A.0.8 capture intuitive properties of belief and knowledge that lead to doxastic and epistemic logics being referred to as the logics of belief and knowledge. The property (1) is commonly known as the *distribution axiom*, and ensures that an agent knows all logical consequences of its own knowledge. This property is also sometimes referred to as logical omniscience. The property (4) is commonly known as the *consistency axiom*, and ensures that an agent cannot have inconsistent beliefs, i.e. beliefs that contradict one another. The related, but stronger property, (7) is commonly known as the *truth axiom*, and ensures that an agent can only know statements that are actually true in the current world. The properties (5), and (6) are commonly known as the *positive introspection axiom* and the *negative introspection axiom* respectively, and ensures that if an agent knows (or believes) a statement, then they know that they know it (or believe that they believe it), and likewise if they do not know or believe a statement, then they know or believe that this is the case.

We note that the proofs of properties (4), (5) and (6) rely only on the serial, transitive and Euclidean properties of *KD45* models, respectively. We also note that the proof of property (7) relies only on the reflexive property of *S5* models. In fact, the class of Kripke models that satisfy each of these properties are precisely the class of Kripke models that satisfy the corresponding relational properties we have described; for example, the class of Kripke models that satisfy property (4) is precisely the class of models that have the serial property, and the class of Kripke models that satisfy property (5) is precisely the class of models that have the transitive property. This demonstrates a connection between the structural constraints imposed upon doxastic and epistemic models, and the properties we have described above, and we will refer back to this connection shortly.

We will now provide a Hilbert-style axiomatisation for $\mathcal{L}^{\mathbf{K}}$, $\mathcal{L}^{\mathbf{KD45}}$ and $\mathcal{L}^{\mathbf{S5}}$. A Hilbert-style axiomatisation gives a method for deriving valid formulae in a logic. The type of axiomatisation that we give is called a substitution schema, which comprises of a set of axioms and a set of rules. The axioms are statements that contain variables, and substituting the variables for well-formed formulae gives a valid statement in the logic. The rules provide a method for deriving a new valid statement from already known valid statements. The main results of this paper are axiomatisations of this style for the variants of refinement quantified modal logics.

Definition A.0.8 (Axiomatisation **K**)**.** The axiomatisation **K** is a substitution schema consisting of the following axioms:

P	All propositional tautologies
K	$\Box_a(\phi \to \psi) \to \Box_a\phi \to \Box_a\psi$

Along with the rules:

MP	From $\vdash \phi \to \psi$ and $\vdash \phi$, infer $\vdash \psi$
NecK	From $\vdash \phi$ infer $\vdash \Box_a\phi$

We say that a formula ϕ is *provable* or *derivable* under an axiomatisation if and only if it can be derived using some finite sequence of axioms and rules from that axiomatisation, and we write $\vdash \phi$. When we are discussing multiple axiomatisations at once, we may add a subscript to the turnstile symbol, e.g. $\vdash_{\mathbf{K}} \phi$, to make it more explicit which logic we are referring to.

An axiomatisation of a logic must have two important properties: *soundness* and *completeness*. An axiomatisation is sound with respect to a logic if and only if every formula that can be derived from the axiomatisation is also valid in that logic. An axiomatisation is complete if and only if every formula that is valid in the logic is also derivable.

The axiomatisation **K**, given above, is shown to be sound and complete by Hughes and Creswell [13].

The axiomatisations for $\mathcal{L}^{\mathbf{KD45}}$ and $\mathcal{L}^{\mathbf{S5}}$ are extensions of the axiomatisation **K**, which we will now provide.

Definition A.0.9 (Axiomatisation **KD45**)**.** The axiomatisation **KD45** is a substitution schema consisting of the axioms and rules of **K**, along with the additional axioms:

D	$\Box_a\phi \to \Diamond_a\phi$
4	$\Box_a\phi \to \Box_a\Box_a\phi$
5	$\Diamond_a\phi \to \Box_a\Diamond_a\phi$

This axiomatisation is shown to be sound and complete by Gabbay [10].

We note that the additional axioms for **KD45** correspond to the properties of *KD45* models described in Proposition A.0.8. The axiom of **5** is equivalent to property (6) from Proposition A.0.8. The axioms **D**, **4** and **5** serve to restrict the logic to considering only the Kripke models that are serial, transitive and Euclidean.

Definition A.0.10 (Axiomatisation **S5**)**.** The axiomatisation **S5** is a substitution schema consisting of the axioms and rules of **K**, along with the additional axioms:

$$\begin{array}{ll} \mathbf{T} & \Box_a\phi \rightarrow \phi \\ \mathbf{5} & \Diamond_a\phi \rightarrow \Box_a\Diamond_a\phi \end{array}$$

This axiomatisation is shown to be sound and complete by Hughes and Creswell [13].

Once again, we note that the additional axioms for **S5** correspond to the properties of *S5* models described in Proposition A.0.8. The axioms **T** and **5** serve to restrict the logic to considering only the Kripke models that are reflexive and Euclidean. We note that the reflexive and Euclidean properties together imply transitivity and symmetry. Although it is not immediately obvious, the axioms **D** and **4** from doxastic logic are both derivable using the axiomatisation **S5**.

Finally we remark on the differences between the notions of belief and knowledge that are modelled by doxastic and epistemic logics. As we have seen, doxastic and epistemic logics both have a notion of positive introspection and negative introspection. This is reflected by the transitive and Euclidean properties of *KD45* and *S5* models. These properties are shown to be defining of these logics by their representations in the axiomatisations **KD45** and **S5**, as the axioms **4** and **5**. The differences between the logics are due to the differences between the consistency axiom and the truth axiom. This is reflected by the serial property of *KD45* models, and by the reflexive property of *S5* models, and these properties are shown to be defining of these logics by the axioms **D** and **T** respectively. Thus we can simply say that knowledge is true belief, and that belief is like knowledge, with the exception that beliefs do not necessarily have to be true, but they do have to be consistent with one another.

There are of course many other variants of modal logics which correspond to slightly different notions of knowledge or belief. For example, one may suppose that it is reasonable to be ignorant of one's own ignorance – that we don't have negative introspection in knowledge – and thus we have the modal logic **S4**, which has the axioms and rules of **K** along with the axioms **T** and **4**, but not **5**. Another example, one may suppose that beliefs do not necessarily have to be consistent with one another – they may be incosistent – thus we have the modal logic **K45**, which has the axioms and rules of **K** along with the axioms **4** and **5**, but not **D**. For this thesis we only consider the logics of **K**, **KD45** and **S5**.

Bibliography

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

[2] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of symbolic logic*, pages 510–530, 1985.

[3] C. Areces, H. van Ditmarsch, T. French, and R. Goldblatt. Undecidability for arbitrary public announcement logic. *Advances in Modal Logic*, 7:23–42, 2008.

[4] P. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. de Lima. What can we achieve by arbitrary announcements?: A dynamic take on fitch's knowability. In *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge*, pages 42–51. ACM, 2007.

[5] A. Baltag and L.S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004.

[6] M. Bilkova, A. Palmigiano, and Y. Venema. Proof systems for the coalgebraic cover modality. *Advances in Modal Logic*, 7:1–21, 2008.

[7] P. Blackburn, M. de Rijke, and Y. Venema. *Modal logic*, volume 53. Cambridge Univ Pr, 2002.

[8] M.J. Fischer and R.E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.

[9] T. French. Bisimulation quantifiers for modal logics. *PhD thesis, University of Western Australia*, 2006.

[10] D.M. Gabbay. *Many-dimensional modal logics: theory and applications*, volume 148. North Holland, 2003.

[11] J. Gerbrandy and W. Groeneveld. Reasoning about information change. *Journal of Logic, Language and Information*, 6(2):147–169, 1997.

[12] J. Hales, T. French, and R. Davies. Refinement quantified logics of knowledge. *Electronic Notes in Theoretical Computer Science*, 278:85–98, 2011.

[13] G.E. Hughes and M.J. Cresswell. *A new introduction to modal logic*. Burns & Oates, 1996.

[14] J.S. Miller and L.S. Moss. The undecidability of iterated modal relativization. *Studia Logica*, 79(3):373–407, 2005.

[15] J. Plaza. Logics of public communications. *Synthese*, 158(2):165–179, 2007.

[16] V.R. Pratt. A near-optimal method for reasoning about action* 1. *Journal of Computer and System Sciences*, 20(2):231–254, 1980.

[17] K. Segerberg. The basic dynamic doxastic logic of agm. *Frontiers in belief revision*, 22:57–84, 2001.

[18] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and computation*, 204(11):1620–1662, 2006.

[19] W. van der Hoek and J.J.C. Meyer. *Epistemic logic for AI and computer science*. Cambridge Univ Pr, 2004.

[20] H. van Ditmarsch. The logic of knowledge games: showing a card. In *Proceedings of the BNAIC*, volume 99, pages 35–42. Maastricht University, 1999.

[21] H. van Ditmarsch. Knowledge games. *Bulletin of Economic Research*, 53(4):249–273, 2001.

[22] H. van Ditmarsch and T. French. Simulation and information: Quantifying over epistemic events. *Knowledge Representation for Agents and Multi-Agent Systems*, pages 51–65, 2009.

[23] H. van Ditmarsch, T. French, and S. Pinchinat. Future event logic: axioms and complexity. In *Advances in Modal Logic*, pages 24–27, 2010.

[24] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic epistemic logic*. Springer Verlag, 2007.

[25] F. Wolter and M. Zakharyaschev. Dynamic description logics. *Advances in Modal Logic*, 1998:431–446, 1998.