

Refinement Quantified Logics of Knowledge and Belief for Multiple Agents

James Hales¹ Tim French Rowan Davies

*Computer Science and Software Engineering
The University of Western Australia
Perth, Australia*

Abstract

Given the “possible worlds” interpretation of modal logic, a refinement of a Kripke model is another Kripke model in which an agent has ruled out some possible worlds to be consistent with some new information. The refinements of a finite Kripke model have been shown to correspond to the results of applying arbitrary action models to the Kripke model [10]. Refinement modal logics add quantifiers over such refinements to existing modal logics. Work by van Ditmarsch, French and Pinchinat [11] gave an axiomatisation for the refinement modal logic over the class of unrestricted Kripke models, for a single agent. Recent work by Hales, French and Davies [13] extended these results, restricting the quantification to the class of doxastic and epistemic models for a single agent. Here we extend these results further, to the classes of doxastic and epistemic models for multiple agents. The generalisation to multiple agents for doxastic and epistemic models is not straightforward and requires novel techniques, particularly for the epistemic case. We provide sound and complete axiomatisations for the considered logics, and a provably correct translations to their underlying modal logics, corollaries of which are expressivity and decidability results.

Keywords: Modal logic, Epistemic logic, Doxastic logic, Bisimulation quantifier, Refinement quantifier, Temporal epistemic logic, Multi-agent system, Action models

1 Introduction

This paper examines the extension of multi-agent doxastic and epistemic logics by refinement quantifiers. Refinement quantifiers were introduced by van Ditmarsch and French [10] to capture a general notion of informative updates in the context of epistemic logic. Informative updates, such as public announcements correspond to an agent receiving new information and incorporating this into their knowledge state. These are discussed in great detail by van Ditmarsch, van der Hoek and Kooi [12].

When we move from explicit updates to arbitrary updates we move from the question “Is ϕ true after the agent learns ψ ?”, to the question, “Is it possible

¹ Acknowledges the support of the Prescott Postgraduate Scholarship.

that the agent can learn *something* in such a way that ϕ is true?”. When the new information is constrained to be expressible as an epistemic formula we have Arbitrary Public Announcement Logic which was shown to be undecidable by van Ditmarsch and French [9]. A refinement quantifier is a weaker operator than an arbitrary announcement. Refinement quantified logics have been shown to be decidable, and axiomatized for the logics K and the modal μ -calculus by van Ditmarsch, French and Pinchinat [11], and for single-agent $KD45$ and $S5$ by Hales, French and Davies [13]. This paper goes to the original motivation for refinement quantifiers as informative updates in multi-agent systems. We present sound and complete axiomatizations for refinement quantified multi-agent $KD45$ and $S5$, and derive expressivity and decidability results for each logic.

Refinements may be thought of as the result of an iterated process of duplicating and removing successor worlds in a Kripke model. This process will preserve an agent’s positive knowledge (things they know) but may not preserve their general knowledge (things they merely suspected to be true are not guaranteed to be true in a refinement). We define the refinements of a Kripke model using a relationship between Kripke models, also called a refinement, which is simply the reverse direction of a simulation, which are a generalisation of bisimulations [5]. As we are in a multi-agent setting, we specify refinement relations with respect to sets of agents, so that the knowledge of all other agents is preserved, except possibly what they may know of the refined agents knowledge state.

Refinements have an important role in multi-agent epistemic logic. Refinements preserve the positive knowledge of all agents [10], so they naturally generalise anything that may be considered an informative update. Examples include public announcements [3], group announcements [1] and action models [12]. A refinement quantifier over some epistemic property, ϕ , corresponds to the question of whether we can provide information to the set of agents in such a way that ϕ will be true. Dynamics systems of knowledge have applications in reasoning about games, autonomous agent negotiation, and communication systems, and in these contexts the refinement operation corresponds to whether one knowledge state is reachable from another.

Our strategy for proving completeness follows the approaches of D’Agostino and Lenzi [6], (and subsequently [11,13]) of giving a provably correct translation into a sublanguage with a known completeness result.

2 Technical Preliminaries

We recall the definitions given by van Ditmarsch, French, and Pinchinat [11] in describing the refinement modal logic, and adapt those definitions to be based on doxastic logic, $KD45$, and epistemic logic, $S5$. Specifically, we restrict the Kripke models under discussion to those in the class of $KD45$ models when we are discussing the extension of the refinement modal logic to $KD45$, which we call the refinement doxastic logic, or $KD45_{\forall}$, and to those in the class of $S5$ models when we are discussing the extension to $S5$, which we call the refinement

epistemic logic, or $S5_{\forall}$.

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

Definition 2.1 (Kripke model) *A Kripke model $M = (S, R, V)$ consists of a domain S , which is a set of states (or worlds), accessibility $R : A \rightarrow \mathcal{P}(S \times S)$, and a valuation $V : P \rightarrow \mathcal{P}(S)$. The class of all Kripke models is called \mathcal{K} . We write $M \in \mathcal{K}$ to denote that M is a Kripke model.*

For $R(a)$ we write R_a . Given two states $s, s' \in S$, we write $R_a(s, s')$ to denote that $(s, s') \in R_a$. We write sR_a for $\{t \mid (s, t) \in R_a\}$. As we will be required to discuss several models at once, we will use the convention that $M = (S, R, V)$, $M' = (S', R', V')$, $M^\gamma = (S^\gamma, R^\gamma, V^\gamma)$, etc. For $s \in S$ we will let M_s refer to the pair (M, s) , or the pointed Kripke model M at state s .

Definition 2.2 (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 $\mathcal{KD45}$. We write $M \in \mathcal{KD45}$ to denote that M is a doxastic model.*

Definition 2.3 (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 $\mathcal{S5}$. We write $M \in \mathcal{S5}$ to denote that M is an epistemic model.*

This paper covers results in both $\mathcal{KD45}$ and $\mathcal{S5}$; as such we will assume that all models are doxastic models when discussing $\mathcal{KD45}$ or $\mathcal{KD45}_{\forall}$, and that all models are epistemic models when discussing $\mathcal{S5}$ or $\mathcal{S5}_{\forall}$.

Definition 2.4 (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, and write $M_s \leftrightarrow M'_{s'}$ to denote that there is a bisimulation between M_s and $M'_{s'}$.

Definition 2.5 (Simulation and refinement) *Let M and M' be Kripke models and let $B \subseteq A$ be a set of agents. A non-empty relation $\mathcal{R} \subseteq S \times S'$ is a B -simulation if and only if it satisfies **atoms**, **forth- a** for every $a \in A$ and **back- a** for every $a \in A \setminus B$.*

If $s \in S$ and $s' \in S'$ such that $(s, s') \in \mathcal{R}$, we call $M'_{s'}$ a B -simulation of M_s and call M_s a B -refinement of $M'_{s'}$. We write $M'_{s'} \succeq_B M_s$, or equivalently, $M_s \preceq_B M'_{s'}$, to denote this.

In the case where $B = A$ we use the terms *simulation* and *refinement* in place of *A-simulation* and *A-refinement*, and we write $M'_s \succeq M_s$, or equivalently, $M_s \preceq M'_s$. In the case where $B = \{a\}$ for some $a \in A$ we simply use the terms *a-simulation* and *a-refinement*, and we write $M'_s \succeq_a M_s$ or $M_s \preceq_a M'_s$.

The refinements of a Kripke model correspond to the results of executing arbitrary action models on the Kripke model [10]. In an epistemic setting, a refinement can therefore be considered as the result of an informative update in which each agent's positive knowledge is preserved, and other knowledge may potentially vary [10]. A B -refinement corresponds to a more restricted informative update in which only the agents in the set B are directly provided with information, and the knowledge of agents not in B is preserved, except for where it concerns non-positive knowledge of agents in B . An example of this is given in the following section, in Example 3.5. How the notion of B -refinements relates to action models is a question left for future work.

3 Syntax and semantics

Here we define the syntax and semantics of the logics $KD45_\forall$ and $S5_\forall$, which restrict the logic K_\forall , defined by van Ditmarsch, French, and Pinchinat to models and refinements of models that are in $KD45$ or $S5$ respectively.

The same syntax used for K_\forall is used for $KD45_\forall$ and $S5_\forall$, and so we will define it only once, as \mathcal{L}_\forall .

Definition 3.1 (Language of \mathcal{L}_\forall) *Given a finite set of agents A and a set of propositional atoms P , the language of \mathcal{L}_\forall is inductively defined as*

$$\phi ::= p \mid \neg\phi \mid (\phi \wedge \phi) \mid \Box_a\phi \mid \forall_B\phi$$

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

We use all of the standard abbreviations for modal logics, in addition to the abbreviation $\exists_B\phi ::= \neg\forall_B\neg\phi$. We abbreviate $\exists_{\{a\}}$ and $\forall_{\{a\}}$ as \exists_a and \forall_a respectively. Similarly we abbreviate \exists_A and \forall_A as \exists and \forall respectively.

We refer to the language \mathcal{L} , of modal formulae, which is simply \mathcal{L}_\forall without the \forall_B operator, and the language \mathcal{L}_0 , of propositional formulae, which is \mathcal{L} without the \Box_a operator. We use the notation $\phi \leq \psi$ to mean that ϕ is a (non-strict) subformula of ψ .

We also use the cover operator, following the definitions given by Bílková, Palmigiano, and Venema [4]. The cover operator, $\nabla_a\Gamma$ is an abbreviation defined by $\nabla_a\Gamma ::= \Box_a \bigvee_{\gamma \in \Gamma} \gamma \wedge \bigwedge_{\gamma \in \Gamma} \Diamond_a\gamma$, where Γ is a finite set of formulae. We note that the modal operators \Box_a , \Diamond_a and ∇_a are interdefineable, as $\Box_a\phi \leftrightarrow \nabla_a\{\phi\} \vee \nabla_a\emptyset$ and $\Diamond_a\phi \leftrightarrow \nabla_a\{\phi, \top\}$. This is the basis of our axiomatisations, as it is for the axiomatisation of the single-agent logic $K_\forall^{(1)}$, presented by van Ditmarsch, French and Pinchinat [11], and the axiomatisations for the single-agent logics $KD45_\forall^{(1)}$ and $S5_\forall^{(1)}$ presented by Hales, French and Davies [13]. The cover operator allows us to define normal forms for modal logics that allow us to only consider conjunctions of modalities in specific situations for our axiomatisations and provably correct translations.

The semantics for K_{\forall} , $KD45_{\forall}$ and $S5_{\forall}$ are very similar, and so we will introduce a generalised semantics that can be applied to all three.

Definition 3.2 (Semantics of C_{\forall}) *Let \mathcal{C} be a class of Kripke models, and let $M = (S, R, V)$ be a Kripke model taken from the class \mathcal{C} . The interpretation of $\phi \in \mathcal{L}_{\forall}$ is defined by induction.*

$$\begin{aligned} M_s \models p & \quad \text{iff } s \in V_p \\ M_s \models \neg\phi & \quad \text{iff } M_s \not\models \phi \\ M_s \models \phi \wedge \psi & \quad \text{iff } M_s \models \phi \text{ and } M_s \models \psi \\ M_s \models \Box_a \phi & \quad \text{iff for all } t \in S : (s, t) \in R_a \text{ implies } M_t \models \phi \\ M_s \models \forall_B \phi & \quad \text{iff for all } M'_{s'} \in \mathcal{C} : M_s \succeq_B M'_{s'} \text{ implies } M'_{s'} \models \phi \end{aligned}$$

The logics K_{\forall} , $KD45_{\forall}$ and $S5_{\forall}$ are instances of C_{\forall} with the classes \mathcal{K} , $\mathcal{KD45}$ and $\mathcal{S5}$ respectively. The difference between these logics are the class of models that formulae are interpreted over, and the class of models that the refinement quantifier, \forall_B quantifies over. It should be emphasised that the interpretation of the refinement operator, \forall_B , varies for each logic, as the refinements considered in the interpretation of each logic must be taken from the appropriate class of models. It is for this reason that $KD45_{\forall}$ and $S5_{\forall}$ are not conservative extensions of K_{\forall} . For example, $\exists_a \Box_a \perp$ is valid in K_{\forall} , but not in $KD45_{\forall}$ or $S5_{\forall}$. This is because given any pointed model in \mathcal{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, $\Box_a \perp$ is satisfied, and hence $\exists_a \Box_a \perp$ is satisfied in the original model. However because of the seriality of $\mathcal{KD45}$ and $\mathcal{S5}$ models, $\Box_a \perp$ is not even satisfiable in $KD45_{\forall}$ or $S5_{\forall}$, and as \exists_a quantifies over $\mathcal{KD45}$ or $\mathcal{S5}$ models in these cases, therefore $\exists_a \Box_a \perp$ is not satisfiable either.

Lemma 3.3 *The logics K_{\forall} , $KD45_{\forall}$ and $S5_{\forall}$ are bisimulation invariant.*

The proof for bisimulation invariance in K_{\forall} , given by van Ditmarsch, French and Pinchinat [11] applies to $KD45_{\forall}$ and $S5_{\forall}$.

Example 3.4 Imagine a scenario where an agent is presented with three cards face down, and asked to identify which is the ace of spades (let's suppose it is the *left* card). As an agent's knowledge is only ever based on reliable evidence, it follows that given any informative update, there is always a further informative update after which the agent knows the location of the ace:

$$left \rightarrow \forall \exists \Box left. \quad (1)$$

This scenario is represented in Figure 1. We can also imagine a corresponding scenario in terms of the agent's *belief* rather than *knowledge*. Here an agent may believe that the ace is in fact the *centre* card, despite this not being the case. In this setting the formula (1) does not hold. We also note that once an agent holds a belief, no informative update will cause the agent to revise that belief.

$$\Box(right \vee centre) \rightarrow \forall \Box(right \vee centre). \quad (2)$$

That is, we do not consider belief *revision* in the sense of [2] but rather belief *refinement*. This situation is depicted in Figure 2. This allows incorrect information, but requires that the information provided is consistent because of the requirement that $\mathcal{KD45}$ models are serial.




Fig. 1. The initial state of an agent's uncertainty from Example 3.4, with an example refinement, both in $\mathcal{S5}$.




Fig. 2. The initial state of an agent's uncertainty from Example 3.4, with an example refinement, this time in $\mathcal{KD45}$.

Example 3.5 Now consider a situation where we have three agents: James, Rowan and Tim, and three coloured cards: Red, Green Blue. The cards are dealt one to each player, and it becomes universally known that Tim does not have the blue card. If James has a red card, then James is able to deduce that Tim must have the green card and hence Rowan has the blue card. However, if James has the blue card, then he remains uncertain about which card Tim has. This situation is represented in Figure 3. Here we use the triple $[b, g, r]$ to indicate that in the corresponding world, James has the blue card, Rowan has the green card, and Tim has the red card.

Suppose now that Tim is able to request information from an oracle. He asks whether Rowan has the red card, and the oracle answers. James hears Tim ask the question, but he does not hear the answer. Rowan hears Tim ask the oracle a question, but was not sure if Tim asked whether Rowan has the red card or the green card. The new knowledge state is represented in Figure 4

Note that while only Tim queried the oracle (so the informative update was a refinement over the agent set $\{Tim\}$), Rowan and James were able to learn about *how* Tim's knowledge changed.

These examples show how refinements captured a very general notion of informative update. Quantifying over refinements therefore allows us to determine whether certain knowledge states are achievable among a set of agents, and has applications in designing and verifying security protocols [7] and reasoning about bidding strategies in games [8].

4 Refinement doxastic logic

In this section we consider the refinement doxastic logic, $KD45_{\forall}$. We provide a sound and complete axiomatisation of the multi-agent refinement doxastic logic, and provide expressivity and decidability results.




Fig. 3. The initial knowledge state in Example 3.5, in $S5$. Reflexive, transitive and Euclidean edges are omitted.




Fig. 4. The final knowledge state in Example 3.5. It can be seen to be a $\{Tim\}$ -refinement of the scenario in Figure 3 by relating worlds with equal assignments.

As with previous work by van Ditmarsch, French and Pinchinat [11], and Hales, French and Davies [13], completeness, expressivity and decidability results are shown by a provably correct translation from the language of refinement modal formulae to the language of modal formulae. The translation and axiomatisation rely on a special normal form for doxastic logic, in terms of the cover operator. We first introduce this normal form in the technical preliminaries of this section, and in the following subsection we introduce the axiomatisation, prove its soundness, and provide the provably correct translation from which our other results follow.

4.1 Technical preliminaries

The axiomatisation of the single-agent refinement modal logic by van Ditmarsch, French and Pinchinat [11] relied on a cover logic disjunctive normal form, formulated in terms of the cover operator. The following axiomatisation of the single-agent refinement doxastic logic by Hales, French and Davies [13] relied on a restricted version of this normal form, the cover logic prenex normal form. The cover logic prenex normal form restricts the formulae inside cover operators to be only propositional formulae. In the single-agent doxastic logic, all formulae may be expressed in this prenex normal form, however this is not true in the multi-agent setting. We introduce a generalisation of the prenex normal form to the multi-agent setting, which we call the cover logic alternating disjunctive normal form.

We first introduce the (non-cover logic) alternating disjunctive normal form, and then introduce its cover logic version. These are analogous to the prenex normal form and the corresponding cover logic version used by Hales, French and Davies [13].

Definition 4.1 (Alternating disjunctive normal form (ADNF)) *A formula in a-alternating disjunctive normal form (abbreviated as a-ADNF) is defined by the following abstract syntax, where α is a formula in a-ADNF:*

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

where $\pi \in \mathcal{L}_0$, $b \in A \setminus \{a\}$, and γ_b stands for a formula in b -ADNF.

A formula in alternating disjunctive normal form (abbreviated as ADNF) is defined by the following abstract syntax, where α is a formula in ADNF:

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

where $\pi \in \mathcal{L}_0$, $a \in A$, and γ_a stands for a formula in a -ADNF.

The alternating disjunctive normal form essentially prohibits direct nestings of modal operators of a particular agent inside modal operators of the same agent. For example, the formula $\Box_a \Diamond_a p$ is *not* in ADNF, because the \Diamond_a operator is nested directly within the \Box_a operator, but the formula $\Box_a \Box_b \Diamond_a p$ is in ADNF because although \Diamond_a is nested within the \Box_a operator, there is a \Box_b operator inbetween. In the case where there is only one agent in the language, this is the same as the prenex normal form of Hales, French and Davies [13], where modal operators may only contain propositional formulae. We will now show that every formula of \mathcal{L} is equivalent to a formula in ADNF, under the semantics of $KD45$.

Lemma 4.2 *We have the following equivalences in $KD45$:*

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

This is proven by Meyer and van der Hoek [14] for the single-agent epistemic logic, $S5^{(1)}$, however the same proof also applies to $KD45$.

Meyer and van der Hoek remarked that the only use of the reflexivity axiom of $S5$, \mathbf{T} , in the proof, is in the form of the theorems $\vdash \Box \Box \phi \rightarrow \Box \phi$, and $\vdash \Box \neg \Box \phi \rightarrow \neg \Box \phi$. Therefore the proof holds for any logic which replaces \mathbf{T} with axioms entailing both of these properties. Both of these properties are valid in $KD45$, and therefore the proof by Meyer and van der Hoek [14] applies to this result.

Lemma 4.3 *Every formula of \mathcal{L} is equivalent to a formula in ADNF, under the semantics of $KD45$.*

Proof. We use a similar reasoning to the proof for prenex normal form, given by Meyer and van der Hoek [14]. If we proceed by induction, assuming that all formulae within a -modal operators are already in the appropriate a -ADNF, then we can iteratively apply the equivalences from Lemma 4.2 in order to replace subformulae where modal operators belonging to a particular agent appear directly within a modal operator of the same agent. \square

The axiomatisation of $KD45_{\nabla}$ that we present in the next section is described in terms of the cover operator, ∇ , which we introduced previously. As a cover logic version of the prenex normal form was used in the axiomatisation

of the single-agent logic $KD45_{\forall}^{(1)}$, we use a cover logic version of the ADNF for the multi-agent $KD45_{\forall}$ logic.

We first introduce the (non-alternating) cover logic disjunctive normal form.

Definition 4.4 (Cover logic disjunctive normal form (CDNF)) *A formula in cover logic disjunctive normal form (abbreviated as CDNF) is defined by the following abstract syntax:*

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

where $\pi \in \mathcal{L}_0$, $B \subseteq A$ and Γ_b stands for a finite set of formulae in CDNF.

Lemma 4.5 *Every formula of \mathcal{L} is equivalent to a formula in cover logic disjunctive normal form, under the semantics of K .*

This is shown by Janin and Walukiewicz [15] in the context of the modal μ -calculus, where the CDNF also contains μ and ν operators, but we note that the result and proof holds if the μ and ν operators are removed. We also note that as this result is in K , it also holds for $KD45$ and $S5$.

The CADNF is essentially a cover logic version of the ADNF. It prohibits cover operators of a particular agent from appearing directly within the scope of cover operators of the same agent. In the case where there is only one agent in the language, this is the same as the cover logic prenex normal form of Hales, French and Davies [13].

Definition 4.6 (Cover logic alternating disj. normal form (CADNF)) *A formula in a -cover logic alternating disjunctive normal form (abbreviated as a -CADNF) is defined by the following abstract syntax:*

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

where $\pi \in \mathcal{L}_0$, $B \subseteq A \setminus \{a\}$, and Γ_b stands for a finite, non-empty set of formulae in b -CADNF.

A formula in cover logic alternating disjunctive normal form (abbreviated as CADNF) is defined by the following abstract syntax:

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

where $\pi \in \mathcal{L}_0$, $B \subseteq A$, and Γ_a stands for a finite, non-empty set of formulae in a -CADNF.

Lemma 4.7 *Every formula of \mathcal{L} is equivalent to a formula in cover logic alternating disjunctive normal form, under the semantics of $KD45$.*

Proof. By Lemma 4.3 we can write any formula in ADNF. By Lemma 4.5, we can rewrite this formula in ADNF into CDNF. We note that as the results for Lemma 4.5 are for K -equivalence, the algorithm for this conversion preserves

the property that modal operators of a particular agent are not nested inside modal operators of the same agent (as otherwise the resulting formula would not be K -equivalent). Therefore the result is in CADNF. \square

The CADNF is used in the formulation of our axiomatisation, and is relied upon for our soundness proofs, and as the basis of our provably correct translation used for the completeness, expressivity and decidability results.

4.2 Axiomatisation

We provide an axiomatisation of the multi-agent refinement quantified doxastic logic, $KD45_{\forall}$, and prove its soundness and completeness. Expressivity and decidability results are given as corollaries.

Definition 4.8 (Axiomatisation \mathbf{RML}_{KD45}) *The axiomatisation \mathbf{RML}_{KD45} is a substitution schema consisting of the following axioms:*

$$\begin{array}{l}
\mathbf{P} \text{ All propositional tautologies} \\
\mathbf{K} \ \Box_a(\phi \rightarrow \psi) \rightarrow (\Box_a\phi \rightarrow \Box_a\psi) \\
\mathbf{D} \ \Box_a\phi \rightarrow \Diamond_a\phi \\
\mathbf{4} \ \Box_a\phi \rightarrow \Box_a\Box_a\phi \\
\mathbf{5} \ \Diamond_a\phi \rightarrow \Box_a\Diamond_a\phi \\
\mathbf{R} \ \forall_B(\phi \rightarrow \psi) \rightarrow (\forall_B\phi \rightarrow \forall_B\psi) \\
\mathbf{RP} \ \forall_B\pi \leftrightarrow \pi \text{ where } \pi \text{ is a propositional formula} \\
\mathbf{RKD45} \ \exists_B\nabla_a\Gamma_a \leftrightarrow \nabla_a(\{\exists_B\gamma \mid \gamma \in \Gamma_a\} \cup \{\top\}) \text{ where } a \in B \\
\mathbf{RComm} \ \exists_B\nabla_a\Gamma_a \leftrightarrow \nabla_a\{\exists_B\gamma \mid \gamma \in \Gamma_a\} \text{ where } a \notin B \\
\mathbf{RDist} \ \exists_B \bigwedge_{c \in C} \nabla_c\Gamma_c \leftrightarrow \bigwedge_{c \in C} \exists_B\nabla_c\Gamma_c \text{ where } C \subseteq A
\end{array}$$

where for every $a \in A$, the set Γ_a is a non-empty set of formulae in a -ADNF. Along with the rules:

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

The axiomatisation \mathbf{RML}_{KD45} is similar to the axiomatisation for the single-agent case considered previously[13] in many respects. Notable differences are that \mathbf{RML}_{KD45} relies on the formulae in cover operators being in a -ADNF instead of being propositional formulae, and that \mathbf{RML}_{KD45} also introduces the \mathbf{RComm} and \mathbf{RDist} axioms, which are required to handle the interactions between different agents.

Each of the $\mathbf{RKD45}$, \mathbf{RComm} axioms serve to push \exists_B operators inside a modality, specifically a cover operator, so that it is applied to each formula inside the cover operator. The axiom \mathbf{RDist} allows us to push the \exists_B operator inside conjunctions in a very specific case. The axiom \mathbf{R} allows us to push \exists_B operators inside disjunctions, and the \mathbf{RP} axiom allows us to eliminate \exists_B operators applied to propositional atoms. This forms the basis of our provably correct translation from \mathcal{L}_{\forall} to \mathcal{L} .

The restriction to a -ADNF for the **RKD45**, **RComm** and **RDist** axioms is necessary to resolve inconsistencies that may be caused by positive and negative introspection of belief that is present in $KD45$. For example, we will look at a counter-example to the **RKD45** axiom if we relax the restriction. Let $M = (S, R, V)$ where $S = \{1, 2, 3\}$, $R_a = \{(1, 2), (1, 3), (2, 3), (3, 2)\}$ and $V(p) = \{2\}$. Then $M_1 \models \nabla_a \{\exists_a \Box_a p, \exists_a \Box_a \neg p, \top\}$. However the formula $\nabla_a \{\Box_a p, \Box_a \neg p\}$ is not satisfiable in $KD45$ and so $M_1 \not\models \exists_a \nabla_a \{\Box_a p, \Box_a \neg p\}$. This problem is not present in K_\forall , but exists in $KD45_\forall$ because of positive and negative introspection. The same problem is present in $KD45_\forall^{(1)}$ (the counter-example used here involves only one agent), and is avoided by use of the cover logic prenex normal form.

Lemma 4.9 *The axiomatisation \mathbf{RML}_{KD45} is sound with respect to the semantic class $KD45$.*

Proof. The soundness of the axioms **P** and **K**, **D**, **4**, **5** and the rules **MP** and **NecK** can be shown by the same reasoning used to show that they are sound in $KD45$. The soundness of the axioms **RP** and **R**, and the rule **NecR** can be shown by the same reasoning used to show that they are sound in the single-agent refinement quantified modal logic, as shown by van Ditmarsch, French and Pinchinat [11].

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

RKD45 (\implies) Trivial. The refinements at successors that we require to satisfy the right-hand side of the equivalence are the successors of the refinement that is entailed by the left-hand side of the equivalence.


(\Leftarrow) Let $M_s \in KD45$ be a doxastic model such that $M_s \models \nabla_a (\{\exists_B \gamma \mid \gamma \in \Gamma_a\} \cup \{\top\})$, where $a \in B$ and γ is an a -ADNF for every $\gamma \in \Gamma_a$. Then for every $\gamma \in \Gamma_a$ there exists some $t^\gamma \in sR_a$ and $M_{t^\gamma}^\gamma \in KD45$ such that $M_{t^\gamma}^\gamma \preceq_B M_{t^\gamma}$ (via a B -simulation \mathfrak{R}^γ) and $M_{t^\gamma}^\gamma \models \gamma$. Without loss of generality, we assume that each of the M^γ are disjoint.

We need to show that $M_s \models \exists_B \nabla_a \Gamma_a$. To do this we will construct a model $M'_{s'} \in KD45$, such that $M'_{s'} \preceq_B M_s$ and show that $M'_{s'} \models \nabla_a \Gamma_a$.

We begin by constructing the model $M' = (S', R', V')$ where:

$$\begin{aligned} S' &= \{s'\} \cup \{t^{\gamma'} \mid \gamma \in \Gamma_a\} \cup S \cup \bigcup_{\gamma \in \Gamma_a} S^\gamma \\ R'_a &= \{(s', t^{\gamma'}) \mid \gamma \in \Gamma_a\} \cup \{t^{\gamma'} \mid \gamma \in \Gamma_a\}^2 \cup R_a \cup \bigcup_{\gamma \in \Gamma_a} R_a^\gamma \\ R'_b &= \{(s', t) \mid t \in sR_b\} \cup \{(t^{\gamma'}, u) \mid u \in t^\gamma R_b^\gamma\} \cup R_b \cup \bigcup_{\gamma \in \Gamma_a} R_b^\gamma \\ V'(p) &= \{s' \mid s \in V(p)\} \cup \{t^{\gamma'} \mid \gamma \in \Gamma_a, t^\gamma \in V^\gamma(p)\} \cup V(p) \cup \bigcup_{\gamma \in \Gamma_a} V^\gamma(p) \end{aligned}$$

where $b \in A \setminus \{a\}$, $p \in P$, and s' and each of the $t^{\gamma'}$ are new states that do not appear in S or any of the S^γ . This construction is shown in Figure 5.

Fig. 5. Construction of refinement $M'_{s'} \preceq_B M_s$.

We note that by construction $M' \in \mathcal{KD45}$ and $M'_{s'} \preceq_B M_s$, via the B -simulation $\mathfrak{R} = \{s', s\} \cup \{(t^{\gamma'}, t^\gamma) \mid \gamma \in \Gamma_a\} \cup \{(t, t) \mid t \in S\} \cup \bigcup_{\gamma \in \Gamma_a} \mathfrak{R}^\gamma$.

We must show that $M'_{s'} \models \nabla_a \Gamma_a$. To do this we show for every $\gamma \in \Gamma_a$ that $M'_{t^{\gamma'}} \models \gamma$. Let $\gamma \in \Gamma_a$. As γ is an a -ADNF, it is a disjunction of conjunctions of propositional formulae and formulae of the form $\Box_b \phi$ and $\Diamond_b \phi$, where $b \neq a$. By construction, the valuation of $M'_{t^{\gamma'}}$ is identical to the valuation of $M_{t^{\gamma'}}$, therefore for any $\pi \in \mathcal{L}_0$ we have that $M'_{t^{\gamma'}} \models \pi$ if and only if $M_{t^{\gamma'}} \models \pi$. Furthermore we note for every $u \in S^\gamma$ that $M'_u \leftrightarrow M_u^\gamma$, and so for any $\psi \in \mathcal{L}$ we have that $M'_{t^{\gamma'}} \models \psi$ if and only if $M_{t^{\gamma'}}^\gamma \models \psi$; in particular this is the case when $\psi = \Box_b \phi$ or $\psi = \Diamond_b \phi$. By construction, $t^{\gamma'} R'_b = t^\gamma R_b$, and so $M'_{t^{\gamma'}} \models \Box_b \phi$ if and only if $M_{t^{\gamma'}}^\gamma \models \Box_b \phi$ if and only if $M_{t^\gamma}^\gamma \models \Box_b \phi$. As by hypothesis we have $M_{t^\gamma}^\gamma \models \gamma$ we therefore have that $M'_{t^{\gamma'}} \models \gamma$. As this holds for every $\gamma \in \Gamma_a$, it follows that $M'_{s'} \models \nabla_a \Gamma_a$.

As $M'_{s'} \preceq_B M_s$, and $M'_{s'} \models \nabla_a \Gamma_a$ we therefore have that $M_s \models \exists_B \nabla_a \Gamma_a$. Therefore **RKD45** is sound.

RComm (\implies) Trivial.

(\impliedby) The proof is similar to the proof for **RKD45**. The construction for **RComm**, and the B -simulation used to show that it is a B -refinement are identical. Slightly different reasoning must be used to show that this construction is a B -refinement, reflecting the fact that in this case $a \notin B$, but this is

straightforward. The reasoning used to show that in the resulting B -refinement that $M'_{s'} \models \nabla_a \Gamma_a$ is identical to the reasoning used in **RKD45**.

RDist (\implies) Trivial.

(\impliedby) Let $M_s \in \mathcal{KD45}$ be a doxastic model such that $M_s \models \bigwedge_{c \in C} \exists_B \nabla_c \Gamma_c$, where $B, C \subseteq A$ and Γ_c is a set of c -ADNF for every $c \in C$. Then for every $c \in C$ there exists some $M_{s^c}^c \in \mathcal{KD45}$ such that $M_{s^c}^c \preceq_B M_s$ (via a B -simulation \mathfrak{R}^c) and $M_{s^c}^c \models \nabla_c \Gamma_c$. Without loss of generality, we assume that each of the M^c are disjoint.

We need to show that $M_s \models \exists_B \bigwedge_{c \in C} \nabla_c \Gamma_c$. To do this we will construct a model $M'_{s'} \in \mathcal{KD45}$ such that $M'_{s'} \preceq_B M_s$ and show that $M'_{s'} \models \bigwedge_{c \in C} \nabla_c \Gamma_c$.

We begin by constructing the model $M' = (S', R', V')$ where:

$$\begin{aligned} S' &= \{s'\} \cup S \cup \bigcup_{c \in C} S^c \\ R'_c &= \{(s', t) \mid t \in s^c R_c^c\} \cup R_c \cup \bigcup_{d \in C} R_c^d \text{ for } c \in C \\ R'_c &= \{(s', t) \mid t \in s R_c\} \cup R_c \cup \bigcup_{d \in C} R_c^d \text{ for } c \notin C \\ V'(p) &= \{s' \mid s \in V(p)\} \cup V(p) \cup \bigcup_{c \in C} V^c(p) \end{aligned}$$

where $p \in P$ and s' is a new state that does not appear in S or any of the S^c .

We note that by construction $M' \in \mathcal{KD45}$ and $M'_{s'} \preceq_B M_s$, via the B -simulation $\mathfrak{R} = \{(s', s)\} \cup \{(t, t) \mid t \in S\} \cup \bigcup_{c \in C} \mathfrak{R}^c$.

We use a similar argument to that used in **RKD45** to show that $M'_{s'} \models \bigwedge_{c \in C} \nabla_c \Gamma_c$. We note for every $c \in C$ and $u \in S^c$ that $M'_u \leftrightarrow M_u^c$. By construction we have that $s' R'_c = s^c R_c^c$, and therefore from $M_{s^c}^c \models \nabla_c \Gamma_c$ we have that $M'_{s'} \models \nabla_c \Gamma_c$.

As $M'_{s'} \preceq_B M_s$, and $M'_{s'} \models \nabla_c \Gamma_c$ we therefore have that $M_s \models \exists_B \bigwedge_{c \in C} \nabla_c \Gamma_c$. Therefore **RDist** is sound.

Therefore the axiomatisation **RML_{KD45}** is sound. \square

The construction we use to show the soundness of **RKD45** is similar to the constructions used to show the soundness of the axioms **GK** for the axiomatisation of $K_{\nabla}^{(1)}$ [11] and **GKD45** for the axiomatisation of $KD45_{\nabla}^{(1)}$ [13]. In each construction we begin by assuming for each $\gamma \in \Gamma$ the existence of initial refinements $M_{t^\gamma}^\gamma$ at some successor of M_s , where $M_{t^\gamma}^\gamma \models \gamma$. The construction for **GK**, shown in Figure 6 simply combines each initial refinement M^γ with a new state s' into a combined refinement $M'_{s'}$. As the only new edges on any of the initial refinements are in-bound edges, the interpretation of formulae in these initial refinements are preserved by this construction. In particular this means for each $\gamma \in \Gamma$ that $M_{t^\gamma}^\gamma \models \gamma$, and so it is then a simple matter to show that $M'_{s'} \models \nabla \Gamma$. The construction used for **GKD45** on the other hand, shown in Figure 7, does not use the whole of each initial refinement, but rather only uses a duplicate $t^{\gamma'}$ of the root state t^γ that shares the same valuation, but not the same successors. This works because the axiom **GKD45** assumes that each




Fig. 6. Construction of refinement $M'_{s'} \preceq M_s$ to show the soundness of **GK** in $K_{\forall}^{(1)}$




Fig. 7. Construction of refinement $M'_{s'} \preceq M_s$ to show the soundness of **GKD45** in $KD_{45}^{(1)}$

$\gamma \in \Gamma$ is a propositional formula, so only duplicating the valuation is required to preserve the interpretation of these formulae. Moreover, using only the root state is necessary because using the whole of each initial refinement, as in the construction for **GK**, would be problematic in the setting of KD_{45} ; the edges from s' to each state t^γ may necessitate the addition of additional edges from s' to other states in each M^γ , in order to satisfy the transitive property of KD_{45} , and this would mean that we may not have $M'_{s'} \models \nabla\Gamma$.

The construction used for **RKD45**, shown in Figure 5, uses elements from both the **GK** and **GKD45** constructions. As in **GK**, we include the initial refinements into the combined refinement, and we only add in-bound edges to each initial refinement, so that we get the same bisimilarity property relied upon for the **GK** construction. However rather than adding edges from the new root state s' to each state t^γ , we instead introduce *proxy states* $t^{\gamma'}$, corresponding to each t^γ . The proxy states duplicate the propositional valuations and b -successors, where $b \neq a$, of each root state t^γ . In this resulting combined refinement, each proxy state $t^{\gamma'}$ preserves the truth of every a -ADNF formula from the corresponding state t^γ . The proxy states avoid difficulties due to the additional transitive and Euclidean edges that would otherwise be required in order to satisfy the properties of KD_{45} , similar to the difficulties avoided by the construction used for **GKD45**. In the case where we have only a single agent, the construction for **RKD45** is essentially the same as the construction used for **GKD45**.

We now show the completeness of the axiomatisation **RML_{KD45}** by a provably correct translation from \mathcal{L}_\forall to \mathcal{L} . Completeness then follows from the completeness of KD_{45} .

We first introduce some equivalences used by our translation.

Lemma 4.10 *The following are provable equivalences using **RML_{KD45}**:*

- (i) $\exists_B(\phi \vee \psi) \leftrightarrow \exists_B\phi \vee \exists_B\psi$
- (ii) $\exists_B(\pi \wedge \bigwedge_{c \in C} \nabla_c \Gamma_c) \leftrightarrow \pi \wedge \bigwedge_{c \in C \cap B} \nabla_c(\{\exists_B\gamma \mid \gamma \in \Gamma_c\} \cup \{\top\}) \wedge \bigwedge_{c \in C \setminus B} \nabla_c\{\exists_B\gamma \mid$

$$\gamma \in \Gamma_c\}$$

where $\pi \in \mathcal{L}_0$, $B, C \subseteq A$, and for every $c \in C$ the set Γ_c is a non-empty set of c -ADNF formulae.

Proof. (1) is derivable using **P** and **R**. (2) is derivable using **P**, **R**, **RP**, **RKD45**, **RComm** and **RDist**. (3) is derivable using **P**, **R**, **RP**, **RComm** and **RDist**. \square

Lemma 4.11 *Every formula of \mathcal{L}_\forall is provably equivalent to a formula of \mathcal{L} via the axiomatisation of **RML_{KD45}**.*

Proof. We proceed by iteratively removing \exists_B -operators from our formula via provable equivalences. Take any subformula of the form $\exists_B\phi$, where $\phi \in \mathcal{L}$, and rewrite ϕ in CADNF. This is a provable equivalence in **RML_{KD45}** as the axioms of $KD45$ appear in **RML_{KD45}**. We then iteratively apply the equivalences from Lemma 4.10, pushing the \exists_B inside disjunctions and cover operators, until the only \exists_B operators are applied to propositional formulae. We can then use the axiom **RP** to remove these \exists_B operators. \square

Theorem 4.12 *The axiomatisation **RML_{KD45}** is sound and complete with respect to the semantic class $\mathcal{KD45}$.*

Proof. Soundness is proven in Lemma 4.9. As in [11,13], completeness is via the provably correct translation from Lemma 4.11 to a sublanguage \mathcal{L} of \mathcal{L}_\forall , which has completeness for the semantic class $\mathcal{KD45}$ via the axioms of $KD45$ that appear in **RML_{KD45}**. \square

Corollary 4.13 *The logic $KD45_\forall$ is expressively equivalent to $KD45$.*

Corollary 4.14 *The logic $KD45_\forall$ is decidable.*

We note that the decision procedure for $KD45_\forall$ is via the translation to $KD45$, and that this translation has a non-elementary complexity. Better decision procedures are left to future work.

5 Refinement epistemic logic

In this section we consider the refinement epistemic logic, $S5_\forall$. We provide an axiomatisation of the multi-agent refinement epistemic logic, and provide expressivity and decidability results. Our axiomatisation, soundness and completeness results follow the same general format of those we have seen previously, but using a different technique in place of a disjunctive normal form.

5.1 Technical preliminaries

Although the alternating disjunctive normal form is a valid normal form for epistemic logic, it is not sufficiently restricted to give a sound axiomatisation of a similar form to the axiomatisation **RML_{KD45}**. We instead introduce the notion of an *explicit* formula, and formulate our axiomatisation in terms of formulae in this form.

Definition 5.1 (Explicit formulae) *Let $\pi \in \mathcal{L}_0$ be a propositional formula, $B \subseteq A$ be a finite set of agents and for every $b \in B$ let $\Gamma_b \subseteq \mathcal{L}$ be a finite*

set of formulae. Let $\gamma^0 \in \mathcal{L}$ be a formula such that for every $b \in B$ we have $\gamma^0 \in \Gamma_b$. Let $\Phi = \{\phi \mid b \in B, \gamma \in \Gamma_b, \phi \leq \gamma\}$ be a set of subformulae of all of the formulae in each set Γ_b . Finally, let α be a formula of the form

$$\alpha = \pi \wedge \gamma^0 \wedge \bigwedge_{b \in B} \nabla_b \Gamma_b$$

Then α is an explicit formula if the following conditions hold:

- (i) For every $b \in B$, $\gamma \in \Gamma_b$, $\phi \in \Phi$: either $\vdash_{S5} \gamma \rightarrow \phi$ or $\vdash_{S5} \gamma \rightarrow \neg\phi$.
- (ii) For every $b \in B$, $\gamma \in \Gamma_b$, $\Box_b \phi \in \Phi$: $\vdash_{S5} \gamma \rightarrow \Box_b \phi$ if and only if for every $\gamma' \in \Gamma_b$ we have $\vdash_{S5} \gamma' \rightarrow \phi$.

Explicit formulae essentially remove a number of elements of choice from the interpretation of a formula: formulae appearing within the cover operators at the top level must explicitly specify which of their subformulae (and subformulae of other formulae appearing in cover operators) are true or false; the formula must explicitly specify which formula in the cover operator is true at the current state; and each cover operator $\nabla_b \Gamma_b$ in the formula must also explicitly agree on which $\Box_b \phi$ subformulae are true.

We will now show that every formula of \mathcal{L} is equivalent to a disjunction of explicit formulae while preserving its *S5* interpretation.

Lemma 5.2 *Every formula of \mathcal{L} is equivalent to a disjunction of explicit formulae, under the semantics of *S5*.*

Proof. Lemma 4.5 gives equivalence to a disjunction of formulae of the form $\pi \wedge \bigwedge_{b \in B} \nabla_b \Gamma_b$, so we only need show the lemma for such a formula, say δ .

Let $\Phi = \{\phi \mid b \in B, \gamma \in \Gamma_b, \phi \leq \gamma\}$. Then replace each γ in δ with a disjunction over the truth of the subformulae ϕ , giving an equivalent δ' :

$$\delta' = \pi \wedge \bigwedge_{b \in B} \nabla_b \left\{ \bigvee_{\Psi \subseteq \Phi} \left(\gamma \wedge \bigwedge_{\phi \in \Psi} \phi \wedge \bigwedge_{\phi \in \Phi \setminus \Psi} \neg\phi \right) \mid \gamma \in \Gamma_b \right\}$$

We move these disjunctions outwards by iteratively applying the equivalence

$$\nabla_b (\{\rho \vee \sigma\} \cup \Gamma) \equiv_{S5} \nabla_b (\{\rho\} \cup \Gamma) \vee \nabla_b (\{\sigma\} \cup \Gamma) \vee \nabla_b (\{\rho, \sigma\} \cup \Gamma)$$

Then distributivity of conjunction over disjunction gives a disjunction of formulae of the form

$$\pi \wedge \bigwedge_{b \in B} \nabla_b \Gamma'_b$$

where each Γ'_b has only elements of the form $\gamma \wedge \bigwedge_{\phi \in \Psi} \phi \wedge \bigwedge_{\phi \in \Phi \setminus \Psi} \neg\phi$. Then, by reflexivity this is equivalent to

$$\pi \wedge \bigwedge_{b \in B} \left(\bigvee_{\gamma' \in \Gamma'_b} \gamma' \wedge \nabla_b \Gamma'_b \right)$$

and again distributivity gives an equivalent disjunction of formulae of the form

$$\pi \wedge \bigwedge_{b \in B} (\gamma'_b \wedge \nabla_b \Gamma'_b)$$

where we omit inconsistent disjuncts, and the Γ'_b are as previously.

Then for every $b, c \in B$ we must have that $\gamma'_b \rightarrow \gamma'_c$ or $\gamma'_b \rightarrow \neg\gamma'_c$, since γ'_c is a conjunction of elements of $\phi \in \Phi$ and their negations, and γ'_b contains either ϕ or $\neg\phi$ for each. But, $\gamma'_b \rightarrow \neg\gamma'_c$ leads to the disjunct being inconsistent, so $\gamma'_b \rightarrow \gamma'_c$, and by the same reasoning $\gamma'_c \rightarrow \gamma'_b$, thus all γ'_b are equivalent.

So for any $b \in B$ we can let $\gamma'_0 = \gamma'_b$, and rewrite our disjunct as

$$\pi \wedge \gamma'_0 \wedge \bigwedge_{b \in B} (\nabla_b \Gamma'_b)$$

This is of the appropriate form for an explicit formula, and it satisfies the two conditions, as follows. The relevant set of subformulae is

$$\Phi' = \{\phi' \mid b \in B, \gamma' \in \Gamma'_b, \phi' \leq \gamma'\}$$

and each γ' still has the form

$$\gamma \wedge \bigwedge_{\phi \in \Psi} \phi \wedge \bigwedge_{\phi \in \Phi \setminus \Psi} \neg\phi$$

- (i) Always either $\gamma' \rightarrow \phi'$ or $\gamma' \rightarrow \neg\phi'$ because ϕ' is a conjunction of formulae $\phi \in \Phi$ and their negations, and for each $\gamma' \rightarrow \phi$ or $\gamma' \rightarrow \neg\phi$.
- (ii) Let $b \in B$, $\gamma' \in \Gamma'_b$ and $\Box_b \phi' \in \Phi'$. Suppose that $\vdash_{S5} \gamma' \rightarrow \Box_b \phi'$ and there exists some $\gamma'' \in \Gamma'_b$ such that $\not\vdash_{S5} \gamma'' \rightarrow \phi'$. Then from positive and negative introspection we have that $\vdash_{S5} \nabla_b \Gamma'_b \rightarrow \Box_b \phi'$. Furthermore, from the first property of an explicit formula we have that $\vdash_{S5} \gamma'' \rightarrow \neg\phi'$, and so $\vdash_{S5} \nabla_b \Gamma'_b \rightarrow \Diamond_b \neg\phi' \rightarrow \neg\Box_b \phi'$. But this is a contradiction.

Suppose instead that $\vdash_{S5} \gamma' \rightarrow \neg\Box_b \phi'$ and for every $\gamma'' \in \Gamma'_b$ we have that $\vdash_{S5} \gamma'' \rightarrow \phi'$. Then from the first hypothesis we have that $\vdash_{S5} \Diamond_b \neg\Box_b \phi' \rightarrow \Diamond_b \neg\phi'$, but from the second hypothesis we have that $\vdash_{S5} \nabla_b \Gamma'_b \rightarrow \Box_b \phi'$, a contradiction.

Therefore every formula is equivalent to a disjunction of explicit formulae. \square

We also note that if we take an explicit formula and remove some of the cover operators from the conjunction, that the result is still an explicit formula.

Lemma 5.3 *Let $\alpha = \pi \wedge \gamma^0 \wedge \bigwedge_{b \in B} \nabla_b \Gamma_b$ be an explicit formula, let $C \subseteq B$ and let $\beta = \pi \wedge \gamma^0 \wedge \bigwedge_{c \in C} \nabla_c \Gamma_c$. Then β is also an explicit formula.*

This result follows directly from the definition of explicit formulae once we realise that the set Φ of subformulae is the same for α and β .

Explicit formulae will be used to formulate our axiomatisation, and the equivalence of epistemic formulae to disjunctions of explicit formulae will be used as part of the provably correct translation that we use in the completeness proof. The requirement for explicit formulae will be discussed briefly after introducing the axiomatisation in the next section.

5.2 Axiomatisation

We provide an axiomatisation of multi-agent refinement epistemic logic, $S5_{\forall}$, and prove its soundness and completeness. Expressivity and decidability results are given as corollaries.

Definition 5.4 (Axiomatisation \mathbf{RML}_{S5}) *The axiomatisation \mathbf{RML}_{S5} is a substitution schema consisting of the following axioms:*

$$\begin{array}{l}
\mathbf{P} \text{ All propositional tautologies} \\
\mathbf{K} \ \Box_a(\phi \rightarrow \psi) \rightarrow (\Box_a\phi \rightarrow \Box_a\psi) \\
\mathbf{T} \ \Box_a\phi \rightarrow \phi \\
\mathbf{5} \ \Diamond_a\phi \rightarrow \Box_a\Diamond_a\phi \\
\mathbf{R} \ \forall_B(\phi \rightarrow \psi) \rightarrow (\forall_B\phi \rightarrow \forall_B\psi) \\
\mathbf{RP} \ \forall_B\pi \leftrightarrow \pi \text{ where } \pi \text{ is a propositional formula} \\
\mathbf{RS5} \ \exists_B(\gamma_0 \wedge \nabla_a\Gamma_a) \leftrightarrow \exists_B\gamma_0 \wedge \nabla_a(\{\exists_B\gamma \mid \gamma \in \Gamma_a\} \cup \{\top\}) \text{ where } a \in B \\
\mathbf{RComm} \ \exists_B(\gamma_0 \wedge \nabla_a\Gamma_a) \leftrightarrow \exists_B\gamma_0 \wedge \nabla_a\{\exists_B\gamma \mid \gamma \in \Gamma_a\} \text{ where } a \notin B \\
\mathbf{RDist} \ \exists_B \bigwedge_{c \in C} (\gamma_0 \wedge \nabla_c\Gamma_c) \leftrightarrow \bigwedge_{c \in C} \exists_B(\gamma_0 \wedge \nabla_c\Gamma_c) \text{ where } C \subseteq A
\end{array}$$

where $\gamma_0 \wedge \nabla_a\Gamma_a$ and $\gamma_0 \wedge \bigwedge_{c \in C} \nabla_c\Gamma_c$ are explicit formulae.

Along with the rules:

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

The axiomatisation \mathbf{RML}_{S5} is similar to the axiomatisation \mathbf{RML}_{KD45} , except that it contains the axioms for $S5$, and that the axioms $\mathbf{RS5}$, \mathbf{RComm} and \mathbf{RDist} rely on the notion of explicit formulae, and make specific reference to which formula from within the cover operator is true at the current state. As in \mathbf{RML}_{KD45} , the $\mathbf{RS5}$, \mathbf{RComm} , \mathbf{RDist} , \mathbf{R} and \mathbf{RP} axioms together form the basis of our provably correct translation to epistemic logic.

The restriction to cover logic alternating disjunctive normal form that \mathbf{RML}_{KD45} uses is not sufficient for the axiomatisation \mathbf{RML}_{S5} . For example, we consider a counter-example to $\mathbf{RS5}$ if we replace the restriction to explicit formulae with a restriction to alternating disjunctive normal formulae. Let $M = (S, R, V)$ where $S = \{1, 2\}$, $R_a = \{(1, 1), (1, 2), (2, 1), (2, 2)\}$, $R_b = \{(1, 1), (2, 2)\}$ and $V(p) = \{1\}$. Then $M_1 \models \exists_a\Box_b\Box_a p \wedge \Diamond_a\exists_a\Box_b\Box_a p \wedge \Diamond_a\exists_a\Box_b\Box_a \neg p$. However the formula $\nabla_a\{\Box_b\Box_a p, \Box_b\Box_a \neg p\}$ is not satisfiable in $S5$, and so $M_1 \not\models \exists_a(\Box_b\Box_a p \wedge \nabla_a\{\Box_b\Box_a p, \Box_b\Box_a \neg p\})$. In $KD45_{\forall}$ this is not a problem, but the problem arises in $S5$ because of the addition of reflexivity; in $S5_{\forall}$ we have that $\vdash_{S5} \Box_b\Box_a p \rightarrow \Box_a p$ and $\vdash_{S5} \Box_b\Box_a \neg p \rightarrow \Box_a \neg p$, and so we get the same contradiction as in our previous counter-example in $KD45_{\forall}$.

The construction that we use to show the soundness of $\mathbf{RS5}$ is similar to the construction that we used to show the soundness of $\mathbf{RKD45}$ in Lemma 4.9. The main differences are that in the construction for $\mathbf{RS5}$ all relations must be

reflexive, transitive and Euclidean, and that the root state of the combined refinement is just another $t^{\gamma'}$ state, and is not treated differently as the root state was for the **RKD45** construction. In the construction for **RKD45**, the fact that the interpretation of a -ADNF formulae is preserved comes from the fact that the states of the initial refinements that are included in the combined refinement are bisimilar to the corresponding states from the uncombined initial refinements. This bisimilarity is because the only new edges added in the construction on states in the initial refinements are in-bound edges. In $\mathcal{S5}$ we cannot *only* have in-bound edges into the states from the initial refinements, as in $\mathcal{S5}$ every edge must be symmetric. We must instead use a different strategy to ensure that combining the initial refinements does not vary the interpretation of each formula $\gamma \in \Gamma_a$ that we are interested in. The use of explicit formulae provides us with additional restrictions on the properties of each initial refinement M^γ , which are sufficient to show that combining the initial refinements does not vary the interpretation of the formulae $\gamma \in \Gamma_a$.

Lemma 5.5 *The axiomatisation **RML** $\mathcal{S5}$ is sound with respect to the semantic class $\mathcal{S5}$.*

Proof. The soundness of the axioms **P** and **K**, **T**, **5** and the rules **MP** and **NecK** can be shown by the same reasoning used to show that they are sound in $\mathcal{S5}$. 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, as shown by van Ditmarsch, French and Pinchinat [11].

All that remains to be shown is the soundness of **RS5**, **RComm** and **RDist**.

RS5 (\implies) Trivial.

(\impliedby) Let $M_s \in \mathcal{S5}$ be an epistemic model such that $M_s \models \exists_B \gamma_0 \wedge \nabla_a (\{\exists_B \gamma \mid \gamma \in \Gamma_a\} \cup \{\top\})$, where $a \in B$ and $\gamma_0 \wedge \nabla_a \Gamma_a$ is an explicit formula. Then for every $\gamma \in \Gamma_a$ there exists some $t^\gamma \in sR_a$ and $M_{t^\gamma}^\gamma \preceq_B M_{t^\gamma}$ (via a B -simulation \mathfrak{R}^γ) such that $M_{t^\gamma}^\gamma \models \gamma$. We also have that $M_s \models \exists_B \gamma_0$, where $\gamma_0 \in \Gamma_a$, and so we may assume that $M_{t^{\gamma_0}}^{\gamma_0} \preceq_B M_s$. Without loss of generality, assume that each of the M^γ are disjoint.

We need to show that $M_s \models \exists_B (\gamma_0 \wedge \nabla_a \Gamma_a)$. To do this we will construct a model $M'_{t^{\gamma_0}} \in \mathcal{S5}$ such that $M'_{t^{\gamma_0}} \preceq_B M_s$ and show that $M'_{t^{\gamma_0}} \models \gamma_0 \wedge \nabla_a \Gamma_a$.

We begin by constructing the model $M' = (S', R', V')$ where:

$$\begin{aligned} S' &= \{t^{\gamma'} \mid \gamma \in \Gamma_a\} \cup \bigcup_{\gamma \in \Gamma_a} S^\gamma \\ R'_a &= (\{t^{\gamma'} \mid \gamma \in \Gamma_a\})^2 \cup \bigcup_{\gamma \in \Gamma_a} R_a^\gamma \\ R'_b &= \bigcup_{\gamma \in \Gamma_a} \left((\{t^{\gamma'}\} \cup t^\gamma R_b^\gamma)^2 \cup R_b^\gamma \right) \\ V'(p) &= \{t^{\gamma'} \mid \gamma \in \Gamma_a, t^\gamma \in V(p)\} \cup \bigcup_{\gamma \in \Gamma_a} V^\gamma(p) \end{aligned}$$

where $b \in A \setminus \{a\}$, $p \in P$ and each $t^{\gamma'}$ is a new state that does not appear in S or any of the S^γ .

We note that by construction $M' \in \mathcal{S5}$ and $M'_{t^{\gamma_0}} \preceq_B M_s$, via the B -simulation $\mathfrak{R} = \{(t^{\gamma'}, t^\gamma) \mid \gamma \in \Gamma_a\} \cup \bigcup_{\gamma \in \Gamma_a} \mathfrak{R}^\gamma$.

For every $\gamma \in \Gamma_a$, we can view $t^{\gamma'}$ as initially being a bisimilar copy of t^γ in M^γ , but with its a -successors pruned. The result would be an B -refinement of $M_{t^\gamma}^\gamma$. M' is formed by joining each $t^{\gamma'}$ state with a -edges, the result being an B -refinement of our original model M .

We must show that $M'_{t^{\gamma_0}} \models \gamma_0 \wedge \nabla_a \Gamma_a$. To do this we show by induction on the structure of formulae in Φ that for every $\phi \in \Phi$, $\gamma \in \Gamma_a$ that: $M'_{t^{\gamma'}} \models \phi$ if and only if $M_{t^\gamma}^\gamma \models \phi$, and for every $u \in S^\gamma$ that $M'_u \models \phi$ if and only if $M_u^\gamma \models \phi$.

The base case, where $\phi = p$ for some $p \in P$ follows by construction. The case where $\phi = \neg\alpha$ or $\phi = \alpha \wedge \beta$ follows directly from the induction hypothesis.

Let $b \in A$, $\phi = \Box_b \alpha$, $\gamma \in \Gamma_b$ and $u \in S^\gamma$. Then $M'_u \models \Box_b \alpha$ if and only if for every $v \in uR'_b$ we have that $M'_v \models \alpha$. By construction either $uR'_b = uR_b^\gamma$ or $uR'_b = uR_b^\gamma \cup \{t^{\gamma'}\}$. Suppose that $uR'_b = uR_b^\gamma$. Then by the induction hypothesis, $M'_v \models \alpha$ for every $v \in uR'_b$ if and only if $M_v^\gamma \models \alpha$ for every $v \in uR_b^\gamma = uR_b^\gamma$ if and only if $M_u^\gamma \models \Box_b \alpha$. Suppose instead that $uR'_b = uR_b^\gamma \cup \{t^{\gamma'}\}$. By the induction hypothesis, $M'_v \models \alpha$ for every $v \in uR'_b$ if and only if $M_v^\gamma \models \alpha$ for every $v \in uR_b^\gamma$. We note that $t^{\gamma'} \in uR'_b$ if and only if $t^\gamma \in uR_b^\gamma$, and that by the induction hypothesis $M'_{t^{\gamma'}} \models \alpha$ if and only if $M_{t^\gamma}^\gamma \models \alpha$. Therefore $M'_v \models \alpha$ for every $v \in uR'_b$ if and only if $M_v^\gamma \models \alpha$ for every $v \in uR_b^\gamma$ if and only if $M_u^\gamma \models \Box_b \alpha$. Similar reasoning shows that $M'_{t^{\gamma'}} \models \Box_b \alpha$ if and only if $M_{t^\gamma}^\gamma \models \Box_b \alpha$; we note that to show this in the case where $b = a$ that we must rely on property (ii) in the definition of explicit formulae, and the fact that the a -successors of each state $t^{\gamma'}$ for $\gamma' \in \Gamma$ are only the other states $t^{\gamma''}$ for $\gamma'' \in \Gamma$.

Therefore by induction we have for every $\gamma \in \Gamma_a$ that $M'_{t^{\gamma'}} \models \gamma$ if and only if $M_{t^\gamma}^\gamma \models \gamma$. As for every $\gamma \in \Gamma_a$ we have that $M_{t^\gamma}^\gamma \models \gamma$, this gives us $M'_{t^{\gamma'}} \models \gamma$. This also gives us $M'_{t^{\gamma_0}} \models \gamma_0 \wedge \nabla_a \Gamma_a$. As we have shown that $M'_{t^{\gamma_0}} \preceq_B M_s$ this gives us $M_s \models \exists_B (\gamma_0 \wedge \nabla_a \Gamma_a)$. Therefore **RS5** is sound.

RComm (\implies) Trivial.

(\impliedby) As in the case for **RComm** in **RML_{KD45}**, the proof is similar to the proof for **RS5**. The inductive proof to show that the constructed refinement satisfies $\gamma_b \wedge \nabla_b \Gamma_b$ is similar to the proof for **RS5**, except that where we treat the agent a specially, we instead treat b specially.

RDist (\implies) Trivial.

(\impliedby) Let $M_s \in \mathcal{S5}$ be an epistemic model such that $M_s \models \bigwedge_{c \in C} \exists_B (\gamma_0 \wedge \nabla_c \Gamma_c)$. Then for every $c \in C$ there exists some $M_{s^c}^c \preceq_B M_s$ (via a B -simulation \mathfrak{R}^c) such that $M_{s^c}^c \models \gamma_0 \wedge \nabla_c \Gamma_c$. Without loss of generality, we assume that each of the M^c are disjoint.

We need to show that $M_s \models \exists_B \bigwedge_{c \in C} (\gamma_0 \wedge \nabla_c \Gamma_c)$. To do this we will construct a model $M'_s \in \mathcal{S5}$, show that $M'_s \preceq_B M_s$ and show that $M'_s \models \bigwedge_{c \in C} (\gamma_0 \wedge \nabla_c \Gamma_c)$.

We begin by constructing the model $M' = (S', R', V')$ where:

$$\begin{aligned}
S' &= \{s'\} \cup \bigcup_{c \in C} S^c \\
R'_c &= (\{s'\} \cup s^c R_c^c)^2 \cup R_c \cup \bigcup_{d \in C} R_c^d \text{ for } c \in C \\
R'_c &= (\{s'\} \cup s R_c)^2 \cup R_c \cup \bigcup_{d \in C} R_c^d \text{ for } c \notin C \\
V'(p) &= \{s' \mid s \in V(p)\} \cup V(p) \cup \bigcup_{c \in C} V^c(p)
\end{aligned}$$

where $p \in P$ and s' is a new state that does not appear in S or any of the S^c .

We note that by construction $M' \in \mathcal{S5}$ and $M'_{s'} \preceq_a M_s$, via the a -simulation \mathfrak{R} , where $\mathfrak{R} = (s', s) \cup \bigcup_{b \in A} \mathfrak{R}^b$.

We must show that $M'_{s'} \models \bigwedge_{c \in C} (\gamma_0 \cup \nabla_c \Gamma_c)$. To do this we show by induction on the structure of formulae in Φ that for every $\phi \in \Phi$, $c \in C$, that: $M'_{s'} \models \phi$ if and only if $M_{s^c} \models \phi$, and for every $u \in S^c$ that $M'_u \models \phi$ if and only if $M_u^c \models \phi$. We use similar reasoning as used in the proof for **RS5**.

Therefore we have that $M'_{s'} \preceq_B M_s$ and $M'_{s'} \models \bigwedge_{c \in C} (\gamma_0 \cup \nabla_c \Gamma_c)$. Therefore **RDist** is sound.

Therefore **RMLS5** is sound with respect to the semantic class $\mathcal{S5}$. \square

We show the completeness of the axiomatisation **RMLS5** in a similar fashion to the completeness proof of **RMLKD45**, by a provably correct translation from \mathcal{L}_\forall to \mathcal{L} . Completeness then follows from the completeness of $\mathcal{S5}$.

As for the completeness proof for **RMLKD45**, we introduce some similar equivalences that will be used by our translation.

Lemma 5.6 *The following are provable equivalences using **RMLS5**:*

- (i) $\exists_a(\phi \vee \psi) \leftrightarrow \exists_a \phi \vee \exists_a \psi$
- (ii) $\exists_a(\pi \wedge \gamma_0 \wedge \bigwedge_{b \in A} \nabla_b \Gamma_b) \leftrightarrow \pi \wedge \exists_a \gamma_0 \wedge \exists_a \bigwedge_{\gamma \in \Gamma_a} \diamond_a \exists_a \gamma \wedge \bigwedge_{b \in A} \nabla_b \{\exists_a \gamma \mid \gamma \in \Gamma_b\}$

where $\pi \wedge \gamma_0 \wedge \bigwedge_{b \in A} \nabla_b \Gamma_b$ is an explicit formula.

This can be shown by following similar reasoning as used for the proof of Lemma 4.10, but by substituting **RMLS5** axioms for **RMLKD45** axioms. We also rely on Lemma 5.3 to ensure that the result of applying **RDist** is a conjunction of \exists_a operators applied to explicit formulae.

Lemma 5.7 *Every formula of $\mathcal{S5}_\forall$ is provably equivalent to a formula of $\mathcal{S5}$.*

This can be shown using similar reasoning to Lemma 4.11, but instead of rewriting modal subformulae into cover logic alternating disjunctive normal form, we use Lemma 5.2 to rewrite subformulae as explicit formulae, and instead of using the equivalences from Lemma 4.10, we use the equivalences from Lemma 5.6.

Theorem 5.8 *The axiomatisation **RMLS5** is sound and complete with respect to the semantic class $\mathcal{S5}$.*

Corollary 5.9 *The logic $S5_{\forall}$ is expressively equivalent to $S5$.*

Corollary 5.10 *The logic $S5_{\forall}$ is decidable.*

As in the case for $KD45_{\forall}$, the decision procedure for $S5_{\forall}$ is via the translation to $S5$, and this translation has a non-elementary complexity. Better decision procedures are left to future work.

References

- [1] Ågotnes, T., P. Balbiani, H. van Ditmarsch and P. Seban, *Group announcement logic*, Journal of Applied Logic **8** (2010), pp. 62–81.
- [2] Alchourrón, C., P. Gärdenfors and D. Makinson, *On the logic of theory change: Partial meet contraction and revision functions*, Journal of symbolic logic (1985), pp. 510–530.
- [3] Balbiani, P., 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*, ACM, 2007, pp. 42–51.
- [4] Bílková, M., A. Palmigiano and Y. Venema, *Proof systems for the coalgebraic cover modality*, Advances in Modal Logic **7** (2008), pp. 1–21.
- [5] Blackburn, P., M. de Rijke and Y. Venema, “Modal logic,” Cambridge Univ Pr, 2002, 341–343 pp.
- [6] D’Agostino, G. and G. Lenzi, *An axiomatization of bisimulation quantifiers via the mu-calculus*, Theor. Comput. Sci. **338** (2005), pp. 64–95.
- [7] Dechesne, F. and Y. Wang, *To know or not to know: epistemic approaches to security protocol verification*, Synthese **177** (2010), pp. 51–76.
- [8] van Ditmarsch, H., *The logic of pit*, Synthese **149** (2006), pp. 343–374.
- [9] van Ditmarsch, H. and T. French, *Undecidability for arbitrary public announcement logic*, Advances in Modal Logic **7** (2008), pp. 23–42.
- [10] van Ditmarsch, H. and T. French, *Simulation and information: Quantifying over epistemic events*, Knowledge Representation for Agents and Multi-Agent Systems (2009), pp. 51–65.
- [11] van Ditmarsch, H., T. French and S. Pinchinat, *Future event logic: axioms and complexity*, in: *Advances in Modal Logic*, 2010, pp. 24–27.
- [12] van Ditmarsch, H., W. van der Hoek and B. Kooi, “Dynamic epistemic logic,” Springer Verlag, 2007.
- [13] Hales, J., T. French and R. Davies, *Refinement quantified logics of knowledge*, Electronic Notes in Theoretical Computer Science **278** (2011), pp. 85–98.
- [14] van der Hoek, W. and J. Meyer, “Epistemic logic for AI and computer science,” Cambridge Univ Pr, 2004, 35–38 pp.
- [15] Janin, D. and I. Walukiewicz, *Automata for the modal μ -calculus and related results*, Mathematical Foundations of Computer Science 1995 (1995), pp. 552–562.