Learning and Knowledge in Social Networks

MSc Thesis *(Afstudeerscriptie)*

written by

Robert M. Carrington
(born February 15th, 1989 in Oakland, California, USA)

under the supervision of **Dr Alexandru Baltag**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
August 20, 2013

Members of the Thesis Committee:
Dr Alexandru Baltag
Dr Nina Gierasimczuk
Prof Dr Dick de Jong
Dr Maria Aloni

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

For most purposes, the information an agent can readily access is just as important as the agent's knowledge. This thesis explores several approaches to reasoning about the information agents in a network can access. The first section introduces a modality for information from immediate connections, and axiomatizes the resulting epistemic logic (EAL). I also introduce and axiomatize iterated version of the logic (IAL), which considers information along multiple edges of access. I prove both of these logics complete for epistemic models equipped with an edge relation. I define and axiomatize two additional logics for access without completeness – a version with restricted access (RAL) and one extending the existing framework of epistemic friendship logic with iterated access modalities (IFL).

Key Words: dynamic epistemic logic, social network

To my grandparents. The problems you overcame were a lot tougher than any contained here.

Contents

Acknowledgements

I would like to thank Alexandru Baltag for indulging lots of meandering discussions on network-related topics. One of those winding roads lead to this thesis, so I am very grateful. Thanks also for your guidance though countless technical conundrums.

I would additionally like to thank my fellow Msc Logic students, especially those who worked through the hot summer alongside me. It was humbling to be surrounded by such precise and tenacious intellects. I look forward to seeing where the this eccentric little group ends up down the line. Stay in touch.

I would also like to thank Tanja Kassenar for countless administrative acts of benevolence. With any luck, the e-mail sending you this thesis will be the last of the pestering you get from me.

A special thanks to Jessica Olsen for telling me about the MSc Logic program over a game of pool three years ago. I never would have ended up here if we hadn't crossed paths in Tacoma that summer.

Finally, thank you to my parents and sister for being supportive of this endeavour. The idea of me moving to Amsterdam for two years could not have been easy to swallow, let alone for the purpose of studying logic. I count myself lucky.

Chapter 1

Introduction

> "Here lies a man who was able to surround himself with men far cleverer than himself." - grave of Andrew Carnegie, billionaire industrialist

For most purposes, the set of information an agent has ready access to is more important that the subset she knows independently. If we think about the daily tasks we perform, the knowledge we are operating on extends far beyond our own heads. We rely on various sources – friends, computers, books – to provide us with the knowledge we need. Relying solely on what we know as individuals, it would be impossible for us to function as we normally do. The reason this does not matter is that information we have ready access to functions in virtually the same way as what we already know. It is ready to be used whenever we need it. If we take this point to heart, then the most significant fact about an agent's situation is the information an agent has ready access to rather than the subset the agent already knows.

Furthermore, the *essential* facts about a situation might be about this kind of access to information. Consider a motivating example: Two children sit down for a test knowing the same things individually, but one child has the good fortune to be seated next to a clever friend. This child is in a position to score very well on the test, being able to peek at the answers of his friends. The other child is not as lucky, and stands to do poorly. The facts about individual knowers do not capture the important aspect of the situation, since the two children know the same things. The fact that *does* explain one child's superior position is a fact about access to information, namely that there is a clever friend sitting close by. Note also that this access cannot be recast into some kind of personal knowledge. The first child is cheating off of another student, answering his test based on someone else's information. It is not his knowledge set, nor his neighbors knowledge set, which explains the access. Rather it is the connection between them. Access to information is an important and moreover a distinct concept in representing social situations.

This thesis aims to incorporate notions of access into epistemic logic, to be able to better understand epistemic scenarios like the one described. The thesis focuses first on information access in a coordinated setting, where communication takes the form of network neighbors pooling their information. The goal is to formalize the reasoning underlying statements such as: "Aaron doesn't know

φ, but he is friends with Brad so he could find out" or "Brad is the only friend who knew φ, so Aaron must have learned φ from him."

The structure of the thesis is as follows:

1. In the remainder of this first chapter, I place this thesis within the context of recent philosophical and technical research. I also introduce epistemic logic and dynamic operators, which will be used throughout.

2. The second chapter introduces the central object of study: epistemic network models. These models are epistemic models augmented with an edge relation between agents. I define a logic called Epistemic Access Logic (EAL) for describing both what agents know and have access to via their network neighbors. I provide an axiomatization for EAL and show completeness for epistemic network models.

3. In the third chapter, I consider a language with modalities for the information of agents beyond immediate neighbors. This version is referred to as Iterated Access Logic (IAL). This logic can also be extended with a dynamic modality $[READ]$, which corresponds to the action of all agents reading the information available from other agents. I show that IAL and DIAL are also complete for epistemic network models.

4. In the fourth chapter I provide a logic for restricted information access (RAL). Access to information can, in this framework, be described on the level of individual issues. I present an axiomatization of RAL for which completeness is an open question.

5. The fifth chapter introduces an expanded version of Epistemic Friendship Logic (EFL), studied in studied in [3]. This version IFL expands the language by incorporating iterated modalities for knowledge and friendship.

6. The sixth chapter explores a different aspect of access – the power of individual agents to control information flow in a network. This logic, LCGC builds on a quite different framework from [1], but could be integrated with the previous logics in the future.

7. The final chapter is dedicated to conclusions and promising directions of further research.

This work falls primarily within the tradition of epistemic logic. This field aims to use logical methods in representing and reasoning about the knowledge of agents. The recognized pioneers of this field are G. H. von Wright and Jaakko Hintikka, whose work in the 1950s and 1960s laid the foundation for the formal systems researched today. Since that time, epistemic logic has been studied in such diverse areas as economics, philosophy and computer science. Two developments are particularly relevant for this work. The first is multi-agent epistemic models, which allow for multiple perpectives to be represented within the same model. Secondly, the introduction of dynamic modalities into epistemic logic. A dynamic modality creates formulas that are checked against a modified version of the underlying model, rather than the model itself. Since static models can only capture an agent's knowledge at a particular moment, dynamic modalities are essential for representing changes in what agents know. Both of these advances are highly relevant in the setting social networks, where we will be dealing with any number of agents all communicating with one another.

This thesis can also be seen as falling within the field of social epistemology. Instead of analyzing conditions within the mind of a single knower, social epistemologists look beyond the individual to examine how an agent's knowledge depends upon the larger community. Although the beginning of this study is difficult to isolate, major reference points for this field include Philip Kitcher's investigation into the scientific community in *The Advancement of Science* and Alvin Goldman's examination of social procedures more generally in *Knowledge in a Social World*.

Finally, this work is connected to the field of social network analysis. This area of research takes individuals and binary social ties as atomic, representing them as nodes and edges respectively. Researchers attempt to find mathematical properties of the resulting graphs that can explain features or behaviors within the original social group. This work will also make use of the graph-representation of social structure, and study how this structure affects interaction.

Chapter 2

Background

2.1 The Notion of Access in Context

Work in several areas has advanced the idea that the information we have ready access to is at least as important as the knowledge in our minds. Although this thesis does not assume the truth of any particular theory, it has certainly taken inspiration from and been influenced by these ideas.

The strongest interpretation of this sentiment comes from philosophy of mind, where many hold that the knowledge we have ready access to *qualifies as being* part of our mind. This is referred to as the extended mind thesis, after a 1998 paper published by Andy Clark and David Chalmers [9]. Supporters of this view hold that if the mind is understood as a set of functions – the capacity for thinking, remembering, calculating – then we have to include all the objects involved in those functions and not just the ones that happen to be inside our skulls. Memory is taken as a prime example. Phone number were, in years past, memorized by most people since they had to be dialled by hand. Nowadays with the advancement of technology, most people have those numbers stored and dialled automatically through their phones. If our cellphones are performing the same job that our neurons did previously, they should count as part of our minds just as much as our neurons do. For contributing to a mental function, Clark and Chalmers require that the object is (1) constantly and immediately available and (2) automatically endorsed. That is, trusted or relied upon just as readily as our own brain. The present work does not assume the same requirements, or assume that the external mind hypothesis is true at all. The arguments are mentioned here for a much more modest purpose than the original. This thesis is simply relying on the fact that information outside one's head can be just as indispensable as the information inside of it.

A less controversial body of work emphasizing the importance of information access is found in distributed computing. A distributed system is any set of computers that share a common task, and which interact in the process of completing that task. Familiar examples of distributed computing include e-mail services and the internet. Within distributed systems, it is usually not the case that each computer starts with all the information it will need to perform its portion of the shared task. However, as long as each computer has a procedure for acquiring the information it needs, this fact does not pose a problem. When faced with a question it cannot answer alone, the computer will access the information it needs from the memory of some other network member. Thus, in designing an algorithm for a distributed system, the more important question for successful computation is what information each computer will be able to access rather than what information is stored on that particular computer.

2.2 Epistemic Logic

We arrive at epistemic logic by interpreting the semantics of modal logic to represent the space of possibilities from one or more agents' perspectives. The syntax is similar to standard modal logic, except that the $\Diamond$ and $\Box$ modalities are replaced by a single K_a modality. $K_a\varphi$ is meant to be read as "Agent a knows φ."

Definition (Syntax of $\mathcal{L}_{\text{EL}}$)

$\varphi \ := \ p \mid \neg\varphi \mid \varphi \ \vee \ \varphi \mid K_a\varphi$

where $p \in \text{Prop}$, and $a \in \mathcal{A}$.

Now that we have the formulas of the language $\mathcal{L}_{\text{EL}}$, we provide the model they will be interpreted upon.

Definition An *epistemic model* $\mathcal{M}$ based on a set of agents $\mathcal{A}$ is a triple:

$$(S, (\sim_a)_{a\in\mathcal{A}}, V)$$

where $S \neq \emptyset$ is a set of states, for each $i \in \mathcal{A}$, $\sim_i$ is a binary equivalence relation on S, and $V : \text{Prop} \to \mathcal{P}(S)$ is a valuation.

Each accessibility relation $\sim_a$ connects worlds that the agent is unable to distinguish from the actual world. If φ is true in all the worlds accessible via $\sim_a$, then agent a knows that φ holds in the actual world.

Definition (Semantics of $\mathcal{L}_{\text{EL}}$):

$\mathcal{M}, w \models p$ iff $w \in V(p)$

$\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$

$\mathcal{M}, w \models \varphi \vee \psi$ iff $\mathcal{M}, w \models \varphi$ or $\mathcal{M}, w \models \psi$

$\mathcal{M}, w \models K_a\varphi$ iff for all v such that $w \sim_a v$ we have that $\mathcal{M}, v \models \varphi$

There are numerous proposals for the axioms for knowledge, each with their advantages and disadvantages. The set of axioms known as S5 however has become the default in epistemic logic, and it is what we shall use here.

S5 Axiomatization

(MP) From φ and $\varphi \to \psi$, derive ψ.

(N) From φ derive $K_a\varphi$.

(Prop) All validities of propositional logic

(K) $K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$

$(T)\ K_a\varphi \to \varphi$

$(4)\ K_a\varphi \to K_aK_a\varphi$

$(5)\ \neg K_a\varphi \to K_a\neg K_a\varphi$

Common Knowledge We can expand the language to include a new modality, C, meant to capture the notion of common knowledge, i.e. formulas that are taken as mutually assumed by all agents. Common knowledge can be described in natural language as that knowledge which "everyone knows, and everyone knows that everyone knows, and everyone knows that everyone knows that everyone knows, ... " and so on ad infinitum. Where φ is any formula of $\mathcal{L}_{EL}$, we add:

$$C\varphi$$

Truth for formulas of this form is defined by:

$$\mathcal{M}, w \models C\varphi \text{ iff } v \models \varphi \text{ for every v and every finite chain of the form}$$
$$w = w_o \sim_{a_1} w_1 \sim_{a_2} w_2 ... \sim_{a_n} w_n = t, \text{ with } a_1, ..., a_n \in A$$

If we let $E\varphi$ abbreviate $K_{a_1}\varphi \wedge K_{a_2}\varphi \wedge ... \wedge K_{a_n}\varphi$ for all $a_i \in \mathcal{A}$, we can see that this definition implies:

$$\mathcal{M}, s \models C\varphi \Leftrightarrow \text{s satisfies all the sentences } \varphi, E\varphi, EE\varphi, EEE\varphi...$$

So we have a modality corresponding to the notion of common knowledge given above. We can also define a restricted version of common knowledge. For any subgroup of agents G, we add formulas of the form:

$$C_G\varphi$$

The semantics for $C_G\varphi$ will be the same as above, except we only quantify over finite chains $w = w_o \sim_{a_1} w_1 \sim_{a_2} w_2 ... \sim_{a_n} w_n = t,$ with $a_1, ..., a_n \in G$.

Distributed Knowledge We can also expand the language $\mathcal{L}_{EL}$ to include the modality, D. For every formula φ of $\mathcal{L}_{EL}$, we add a formula:

$$D\varphi$$

which can be read as "It is distributed knowledge that φ." The truth of these new sentences in the language is defined as follows:

$$\mathcal{M}, w \models D\varphi \text{ iff } v \models \varphi \text{ for every } v \text{ such that } w \sim_a v \text{ for all } a \in \mathcal{A}$$

In other words, D corresponds to the intersection of all the $(\sim_a)_{a\in\mathcal{A}}$ relations. This notion is meant to capture the combined sum of all agents' information in the model. It can also be thought of as the potential knowledge of each agent in the group, if the agents were able to communicate their knowledge to each other.

Just as with common knowledge, we can also define this modality for subgroups. We add formulas of the form:

$$D_G\varphi$$

Truth for these formulas will be defined in the same manner as the D_G operators, except we only require φ be satisfied at every v such that $w \sim_a v$ for all $a \in G$

2.3 Dynamic Epistemic Logic

The definitions so far have allowed us to capture knowledge in a static state. An accessibility relation links the states the respective agent cannot distinguish from reality. However, to capture changes in what the agent knows, we incorporate an idea from Propositional Dynamic Logic (PDL). In PDL, "dynamic modalities" are formalized by the notion

$$[\alpha]\varphi$$

This formula can be read as: "If the action α is performed on the current world, then the sentence φ will become true after." In PDL the semantics for this modality are quite close to traditional modal logic – we check satisfaction by examining all worlds accessible via the relation corresponding to α.

Now, however, we instead apply the same idea of dynamics to *models* instead of worlds. This time, we define a function that maps epistemic models to another, and such functions *updates* or *actions* on models. Then we reinterpret the above formula as: "If the action α is performed on the current *model*, then the sentence φ will become true after." We interpret this "if" as a material conditional, so $[\alpha]\varphi$ is true of a model when α cannot be performed.[1]

The family of logics formed by adding such dynamic modalities to an epistemic logic are known as Dynamic Epistemic Logics.

Updates in Dynamic Epistemic Logic

We now formally introduce the idea of an update. The first component of an update α is a function that maps the initial model $\mathbf{S}$ to a new model $\mathbf{S}^\alpha$, meant to represent the model after a change has occurred. The second component is a binary "transition" relation, which links each state in the initial model to one or more states in the new model. The transition relation is meant to connect each state in the old model to the state or states which are identical in the new model. Here "identical" means that they represent the same state of affairs.

Definition An *epistemic update* α consists of:

1. A map $\mathbf{S} \mapsto \mathbf{S}^\alpha$ from an initial model $\mathbf{S} = (S, (\sim_a)_{a\in\mathcal{A}}, V)$ to a new model $\mathbf{S}^\alpha = (S', (\sim'_a)_{a\in\mathcal{A}}, V')$.
2. A binary transition relation $\overset{\alpha}{\underset{\mathbf{S}}{\longrightarrow}} \subseteq S \times S'$, pairing each world in the old model with one or more worlds in the new model.

In updates which inform all agents equally, the state space in the new model will just be the same set as in the original and the transition relation will map every state to its copy. In updates with private communication the set of states must be multiplied to be allow for the different perspectives, and the transition relation becomes more complicated.

[1] An action α can also be non-deterministic, so that multiple states can result from applying α to a model. In this case, all possible output states must satisfy φ for the entire expression to be true.

With the above definitions, we can formally define what it means for dynamic formulas to be true. Intuitively, a state satisfies a formula beginning with a dynamic modality if the corresponding state in the updated model satisfies the formula with the dynamic modality removed.

Definition (Truth for Dynamic Modalities)

$s \models_{\mathbf{S}} [\alpha]\varphi$ iff $t \models_{\mathbf{S}^\alpha} \varphi$ for all $t \in S^\alpha$ such that $s \longrightarrow^{\alpha}_{\mathbf{S}} t$.

Before continuing, we give a concrete example of an epistemic update. This update has been referred to as the "Tell All You Know" modality.[2] Although we will not be using this modality in this thesis directly, other modalities used in this thesis are closely related.

Definition For a given state s, let $s(a) := \{t \in S | (s,t) \in (\sim_a)\}$. Then the "Tell All You Know" update $!a$ maps any model $\mathbf{S} = (S, (\sim_a)_{a\in\mathcal{A}}, V)$ to a new model $\mathbf{S}^{!a} = (S', (\sim'_a)_{a\in\mathcal{A}}, V')$ given by:

$$S' := S \cap s(a)$$

$$s \sim'_a t \text{ iff } s \sim_a t, \text{ for all s, t} \in S'$$

$$V'(p) := V(p) \cap S'$$

The transition relation $\longrightarrow^{!a}$ relates any state $s \in \mathbf{S}$ satisfying φ to the same state in the model $\mathbf{S}^{!a}$.

Reduction Laws for "Tell All You Know"

We can add the following axiom schema to our previous ones to get a complete axiomatization for the new dynamic logic. Given any formula φ, these axioms allow us to push any occurring dynamic modalities deeper into the formula until they reach the atomic propositions and disappear. This results in a static formula equivalent to our original dynamic one.

$$[!a]p \Longleftrightarrow p$$
$$[!a]\neg\varphi \Longleftrightarrow \neg[!a]\varphi$$
$$[!a](\phi \wedge \psi) \Longleftrightarrow [!a]\varphi \wedge [!a]\psi$$
$$[!a]K_b\varphi \Longleftrightarrow D_{\{a,b\}}[!a]\varphi$$
$$[!a]D_G\varphi \Longleftrightarrow D_{G\cup\{a\}}[!a]\varphi$$

Since every dynamic formula thus has an equivalent static one, it follows that

[2]The "Tell All You Know" modality has been taken from lectures presented by Alexandru Baltag as part of the course Topics in Dynamic Epistemic Logic at the ILLC.

the language including the dynamic "Tell All You Know" modality is only as expressive as static EL. Since they give rise to this recursive method for reducing away any dynamic operators, these laws are also often called "recursion laws" or "reduction laws."

Chapter 3

Epistemic Access Logic

We now expand the epistemic logic introduced in the last section by adding communication graphs to our models. We can represent the knowledge of multiple agents as before, but by adding a directed graph we are able to represent one agent having access to the information of another. We can imagine that each agent's knowledge is stored on a database, where the agent can then look up whatever information he needs. An agent a's database is private, but others may have a's password and be able to access the information a has. We assume a has no way of knowing who can read his information or what they have learned. Only the agent accessing knows these facts. Thus information flow is entirely "one-way" in this framework: the agent with access gains information from the other agent without giving any information away.

However, it is important to keep in mind that this is a static logic. Formulas of this language describe what is true at a moment, i.e. in a single model. We will later model communication explicitly via a dynamic modality, but for now we can only stipulate that what we are representing is available information.

3.1 Introducing EAL

We now introduce the language of Epistemic Access Logic (EAL). The syntax for EAL is focused on capturing not just what an agent knows but what an agent *could* know if she pooled information with surrounding agents. Hence, in addition to the K_i modalities for knowledge, we introduce a modality K'_i for available information. Since we assume agents are able to independently pull information from connected neighbors in the communication graph, we can interpret $K'_i\varphi$ in English as: "Agent a has potential knowledge of φ." We also add sentences of the form "aEb" to express that agent a has the ability to access or *enter* agent b's store of information.

Definition (Syntax of $\mathcal{L}_{\mathrm{EAL}}$)

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \ \vee \ \varphi \mid K_a\varphi \mid K'_a\varphi \mid aEb$$

where $p \in$ Prop, a set of proposition symbols, and $a, b \in \mathcal{A}$, a set of agents. We use the standard modal logic abbreviations "$\langle K_a\rangle\varphi$" for $\neg K_a\neg\varphi$ and "$\langle K'_a\rangle\varphi$"

for $\neg K'_a \neg\varphi$.

Abbreviations

We also introduce the new abbreviation $a \sqsubseteq c$, which we read as "c is at least as connected as a." This abbreviation is meant to capture an intuitive idea of c having access to all the resources that a does. We define $a \sqsubseteq c$ as the abbreviation:

$$a \sqsubseteq c \Leftrightarrow \bigwedge_{b \in \mathcal{A}} (aEb \to cEb)$$

"

We can similarly define the abbreviation $a \not\sqsubseteq b$. Such formulas intuitively mean that b is not more connected than a. We define $a \not\sqsubseteq b$ as the abbreviation:

$$a \not\sqsubseteq b \Leftrightarrow \neg \bigwedge_{b \in \mathcal{A}} (aEb \to cEb)$$

.

Finally, we can define the strict notion of $a \sqsubset b$, i.e. "b is more connected than a," via these first two abbreviations. Let $a \sqsubset b$ be defined by:

$$a \sqsubset b \Leftrightarrow a \sqsubseteq b \wedge b \not\sqsubseteq a$$

.

Semantics

We now introduce the models that formulas of EAL will be interpreted on. They are essentially multi-agent epistemic models with a communication graph added: a set of agents $\mathcal{A}$ for the nodes and a set of edges E.

An *epistemic network model* $\mathcal{M}$ for a set of agents $\mathcal{A}$ is a tuple:

$$(W, (\sim_a)_{a \in \mathcal{A}}, E, V)$$

where $W \neq \emptyset$ is a set of states, for each $i \in \mathcal{A}$, $\sim_i$ is a binary equivalence relation on W, E is a function $E : W \to \mathcal{P}(\mathcal{A} \times \mathcal{A})$ which returns the set of edges for a given world. We also impose the following condition on the edge function E: If $(a, b) \in E(w)$ and $w \sim_a w'$, then also $(a, b) \in E(w')$. This condition ensures that agents know who they do and do not have access to. $V : \text{Prop} \to \mathcal{P}(W)$ is a valuation on proposition variables.

It is worth noting that the edge relation between agents has not been specified as symmetric. Agent a can have access to b's computer without the converse being true. Think of a hacker or government agency which can monitor the contents of people's computers. Asymmetry also allows us to represent inactive sources of information in the network. These might include computers, books, or any other repositories of information that agents may be relying upon. From a formal perspective, these sources will just be agents that lack any out-pointing lines of communication. Such sources will thus have no additional information available via communication, and will never update with information from another agent.

Thus far in our models we have epistemic relations representing what each agent knows, but no portion of the model corresponding to the information available to an agent. We define the additional set of relations $\sim_{\langle a \rangle}$ for all $a \in \mathcal{A}$ as follows:

$s \sim_{\langle a \rangle} s'$ iff $s \sim_b s'$ for all b such that $(a, b) \in E(s)$

Definition (Truth for $\mathcal{L}_{\text{EAL}}$):

$\mathcal{M}, w \models p$ iff $w \in V(p)$

$\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$

$\mathcal{M}, w \models \varphi \vee \psi$ iff $\mathcal{M}, w \models \varphi$ or $\mathcal{M}, w \models \psi$

$\mathcal{M}, w \models aEb$ iff $(a, b) \in E(w)$

$\mathcal{M}, w \models K_a\varphi$ iff for all v such that $w \sim_a v$ we have that $\mathcal{M}, w \models \varphi$

$\mathcal{M}, w \models K'_a\varphi$ iff for all v such that $w \sim_{\langle a \rangle} v$ we have that $\mathcal{M}, w \models \varphi$

3.2 A Database Example

In this example, each agent has a computer storing all of his or her information. Some agents also have the password to access the information on another agent's computer. We assume that access is secret, so that only the agent with the access knows this fact. Access is also not generally symmetric – I can know your password without you knowing mine. In the case given below, agents a and b have access to the computer of a third agent, d. Agent c has access to d's computer, but also to the computer of agent e. Agent d and e in turn have access to no other computers, and do not even know whether others are accessing their computers.

First, compare agents a and b. On the basis of knowledge alone, the agents seem to be in quite different positions. The set of formulas known by a is disjoint from that of b. However, since they both have access to d's more extensive information, they have the same potential knowledge. That is, $a \sqsubseteq b$ and $b \sqsubseteq a$. From this perspective, they are in a very similar positions epistemologically.

Now compare these agents to agent c. On the basis of knowledge, agent c is in a far worse position than either a or b: he actually knows nothing. On the other hand, if we look at potential knowledge c is in a superior position. Having access to both d and e, c has more potential knowledge than any other agent in this scenario. Since any information accessible by the first two agents is also accessible by c, $a \sqsubseteq c$ and $b \sqsubseteq c$. Agent e is, from this perspective, in the best epistemological position without knowing more than anyone else (or, in fact, knowing anything at all).

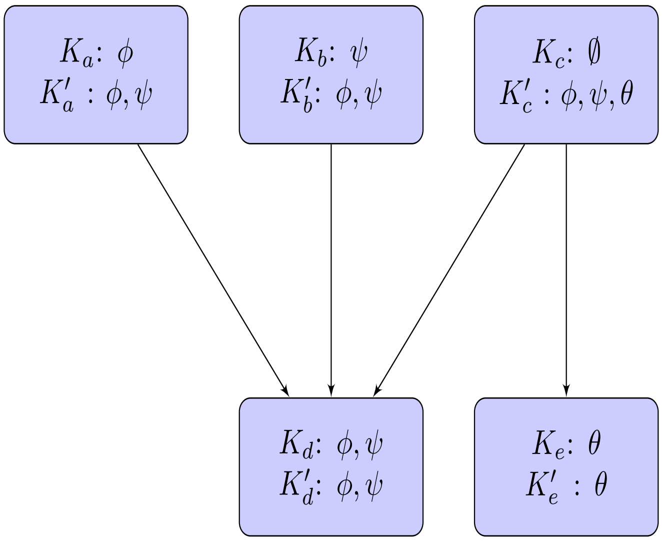

Figure 3.1: Agents are here represented as blocks labelled with the formulas the agent knows and potentially knows. Arrows represent access, with each arrow running from the accessing agent to the one being accessed.

Finally, just judging on the basis of knowledge the agents d and e seem to have an advantage over the others. Agent d knows more than any other agent, and e knows θ while nobody else does. However, if we look at potential knowledge the two agents are at quite a disadvantage. Neither one has access to another agent's computer, so their potential knowledge is identical to their current knowledge.

This example is meant to show how many crucial points in a scenario can depend upon the notion of potential knowledge. Describing the scene in terms of what agents know does not capture all the relevant facts. Agents who know none of the same facts, and thus seem to be in quite different epistemological positions, can actually have the same accessible knowledge. Agents who are completely ignorant can nonetheless have access to more information than anyone else.

3.3 Axiomatization of EAL

Modus ponens and necessitation for K_a are the only inference rules for EAL. In the context of the axioms below, necessitation for K'_a follows from necessitation for K_a.

(MP) From φ and $\varphi \to \psi$, derive ψ.

(N) From φ derive $K_a\varphi$.

The axioms for EAL are the validities of propositional logic, S5 axioms for K_a and K'_a, along with four axioms serving to connect K_a, K'_a and the formulas regarding access.

(Prop) All validities of propositional logic

(K) $K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$

(T) $K_a\varphi \to \varphi$

(4) $K_a\varphi \to K_aK_a\varphi$

(5) $\neg K_a\varphi \to K_a\neg K_a\varphi$

These axioms are standard knowledge axioms for K_a. They ensuring that K_a distributes over conditionals, that $K_a\varphi$ implies the truth of φ and that agents have positive introspection and negative introspection of K_a.

(K') $K'_a(\varphi \to \psi) \to (K'_a\varphi \to K'_a\psi)$

(T') $K'_a\varphi \to \varphi$

$(4')$ $K'_a\varphi \to K'_aK'_a\varphi$

$(5')$ $\neg K'_a\varphi \to K'_a\neg K'_a\varphi$

These axioms are standard knowledge axioms for K'_a. Just as with the previous five, they ensuring that K'_a distributes over conditionals, that $K'_a\varphi$ implies the truth of φ and that agents have positive introspection and negative introspection of K'_a.

(E1) aEa

The axiom (E1) ensures that agents can access their own knowledge. This ensures that the set of formulas known by any agent in the model forms a subset of the formulas the agent potentially knows.

(E2) $(aEb \wedge K_b\varphi) \rightarrow K'_a\varphi$

The axiom (E2) states that if a has access to b's information, then anything b knows is potentially known by a. In other words, (E2) guarantees that your potential knowledge contains the knowledge of those you have access to.

(E3) $aEb \rightarrow K_a aEb$

The axiom (E3) states that each agent knows the agents they have access to. Note that this does *not* imply that agents know who has access to themselves.

(E4) $\bigwedge_{b \in \mathcal{A} \setminus a} (\neg aEb) \rightarrow (K'_a\varphi \rightarrow K_a\varphi)$

Axiom (E4) specifies that if an agent has no other connections, then what he potentially knows is the same as what he knows.

(E5) $a \sqsubseteq c \rightarrow (K'_a \sqsubseteq K'_c)$

Axiom (E5), states that if c is at least as connected as a, then any information a has access to c can access as well.

3.4 Completeness of EAL

Proposition: EAL **is sound and complete for the class of epistemic network models**

To show that EAL is complete for epistemic network models, we shall use intermediary models called "network pseudo-models." These models will contain extra relations so that we can interpret formulas containing K'_a operators in the same way we interpret K_a operators. The strategy, in broad stroke, is to then connect satisfiability in these pseudo-models with satisfiability in epistemic network models. We prove that for every pseudo-model there exists an epistemic network model satisfying the same formulas, and furthermore one in which K'_a must have the semantic definition given earlier. We now introduce the notions needed for the proof.

Definition We say that a formula φ is *consistent* if its negation $\neg\varphi$ cannot be proven in EAL. Otherwise we say that φ is *inconsistent*.

Definition A formula φ is *satisfiable* if there is a network model and a state s such that $(\mathcal{M}, s) \models_{EAL} \varphi$. The formula φ is then said to be satisfiable in $\mathcal{M}$.

Definition A *(network) pseudo-model* for a set of $\mathcal{A}$ of n agents is an epistemic network model of $2n$ agents:

$$\overline{\mathcal{M}} = (S, (\sim_a)_{a\in\mathcal{A}}, (\sim_{\langle a\rangle})_{a\in\mathcal{A}}, E, V)$$

such that $\sim_a, \sim_{\langle a\rangle}$ are equivalence relations for all a.

Definition Where $\overline{\mathcal{M}}$ is a network pseudo-model, we say that $(\overline{\mathcal{M}}, s)$ *pseudo-satisfies* a formula φ if the pair $(\overline{\mathcal{M}}, s)$ satisfies φ in the normal way except that $(\overline{\mathcal{M}}, s) \models K'_a\varphi$ iff $(\overline{\mathcal{M}}, s') \models \varphi$ for all s' such that $(s, s') \in \sim_{\langle a\rangle}$. A pseudo-model $\overline{\mathcal{M}}$ *validates* a formula φ if $(\overline{\mathcal{M}}, s) \models \varphi$ for every state s of $\overline{\mathcal{M}}$.

Definition A pseudo-model of EAL is a pseudo-model which validates all the axioms of EAL.

Proposition: Every epistemic network model is a pseudo-model of EAL.

Note that every epistemic network model is trivially a pseudo-model, given how we have defined $\sim_{\langle a\rangle}$ for network models. So we just need to show that every network model validates the axioms of EAL. Take an arbitrary model $\mathcal{M}$ and state s of $\mathcal{M}$. First off, $\sim_a$ and $\sim_{\langle a\rangle}$ are equivalence relations by definition of a network model. Therefore they will each validate their respective axioms $K, T, 4, 5$ and $K', T', 4', 5'$. Then we prove that $(\mathcal{M}, s)$ satisfies the remaining axioms individually.

(E1) aEa

$(\mathcal{M}, s) \models aEa$ since we specified E to be reflexive in network models.

(E2) $(aEb \wedge K_b\varphi) \to K'_a\varphi$

Assume $(\mathcal{M}, s) \models (aEb \wedge K_b\varphi)$. Then $(\mathcal{M}, t) \models \varphi$ for all t such that $(s,t) \in \sim_b$. However, by definition $(s,t) \in \sim_{\langle a \rangle}$ iff $(s,t) \in \sim_a$ and $(s,t) \in \sim_b$ for all b such that $(a,b) \in E(s)$. Since $\mathcal{M}, s) \models aEb$, $(a,b) \in E(s)$ and so $(s,t) \in \sim_b \to (s,t) \in \sim_{\langle a \rangle}$. Then $(\mathcal{M}, t) \models \varphi$ for all t such that $(s,t) \in \sim_{\langle a \rangle}$. So $(\mathcal{M}, s') \models K'_a\varphi$.

(E3) $aEb \to K_a aEb$

Assume $(\mathcal{M}, s) \models aEb$. Then by our definition of a network model, we have that for every world s' such that $s \sim_a s'$, $(a,b) \in E(s')$. So $(\mathcal{M}, s) \models K_a aEb$.

(E4) $\bigwedge_{b \in \mathcal{A} \setminus a}(\neg aEb) \to (K'_a\varphi \to K_a\varphi)$

Assume $(\mathcal{M}, s) \models \bigwedge_{b \in \mathcal{A} \setminus a}(\neg aEb)$. So then if $(\mathcal{M}, s) \models K'_a\varphi$, we have that $\mathcal{M}, t \models \varphi$ for all t such that $(s,t) \in \sim_{\langle a \rangle}$. But by definition $(s,t) \in \sim_{\langle a \rangle}$ iff $(s,t) \in \sim_a$ and $(s,t) \in \sim_b$ for all b such that $(a,b) \in E(s)$. By our first assumption there are no such b except a himself so $\sim_{\langle a \rangle} = \sim_a$. So $(\mathcal{M}, s) \models K_a\varphi$.

(E5) $a \sqsubseteq c \to (K'_a\varphi \to K'_c\varphi)$

Assume $(\mathcal{M}, s) \models a \sqsubseteq c$. Then we have that $\bigwedge_{b \in \mathcal{A}} (aEb \to cEb)$. We can take an arbitrary pair $(s,s') \in \sim_{\langle c \rangle}$. By how $\sim_{\langle c \rangle}$ is defined for network models, we have that $(s,s') \in \sim_b$ for all b such that cEb. But then by our first assumption, we must have that $(s,s') \in \sim_b$ for all b such that aEb. So $(s,s') \in \sim_{\langle a \rangle}$. Thus we have that $(\mathcal{M}, s) \models (K'_a\varphi \to K'_c\varphi)$.

Definition A formula φ is pseudo-satisfiable if there is a network model $\overline{\mathcal{M}}$ and a state s of $\overline{\mathcal{M}}$ such that $(\overline{\mathcal{M}}, s) \models \varphi$. The formula φ is then said to be pseudo-satisfiable in $\overline{\mathcal{M}}$.

In order to prove completeness for EAL it suffices to prove the following propositions:

(1) If φ is consistent, then φ is pseudo-satisfiable.
(2) If φ is pseudo-satisfiable, then φ is satisfiable.

Proof of (1): If φ is consistent, then φ is pseudo-satisfiable.

We will use the canonical model construction to show that every consistent set of formulas is pseudo-satisfiable.

Definition: A set of formulas Φ is maximal consistent if Φ is consistent and any set of formulas properly containing Φ is inconsistent. If Φ is a maximal consistent set of formulas then we say it is an MCS.

Lindenbaum Lemma: For any consistent set Φ of formulas from $\mathcal{L}_{\mathsf{EAL}}$, there is an MCS Φ^+ such that $\Phi \subseteq \Phi^+$.

Proof: Let $\phi_0, \phi_1, ...$ be an enumeration of the formulas of $\mathcal{L}$. Define Φ^+ as the union of the chain of consistent sets:

$\Phi_0 = \Phi$

$$\Phi_{n+1} = \begin{cases} \Phi_n \cup \{\phi_n\} & \text{if this is consistent} \\ \Phi_n \cup \{\neg\phi_n\} & \text{otherwise} \end{cases}$$

This set Φ^+ is clearly a superset of Φ. Each Φ^{n+1} is consistent, since if adding the new formula ϕ_n would lead to inconsistency then $\neg\phi$ is by definition already provable from Φ^n and can be added without affecting consistency. Furthermore, since we assumed $\phi_0, \phi_1...$ was a complete enumeration, Φ^+ must contain either ϕ or $\neg\phi$ for every formula ϕ. Thus Φ^+ is an MCS, since for any $\phi \notin \Phi^+$ we have that $\neg\phi \in \Phi^+$. Any strict superset of Φ^+ will contain both ϕ and $\neg\phi$ for some ϕ and thus be inconsistent.

Definition: We define the canonical network pseudo-model for a set of agents $\mathcal{A}$ to be $\widehat{\mathcal{M}} = (\widehat{W}, (\widehat{\sim}_a)_{a\in\mathcal{A}}, (\widehat{\sim}_{\{a\}})_{a\in\mathcal{A}}, \widehat{E}, \widehat{V})$, where:

$\widehat{W} = \{\Phi \mid \Phi \text{ is an MCS }\}$

$\Phi \mathrel{\widehat{\sim}_a} \Phi'$ iff for all ψ, $(K_a\psi \in \Phi) \to (\psi \in \Phi')$

$\Phi \mathrel{\widehat{\sim}_{\langle a\rangle}} \Phi'$ iff for all ψ, $(K'_a\psi \in \Phi) \to (\psi \in \Phi')$

$\widehat{E}(\Phi) = \{(a, b) \mid aEb \in \Phi\}$

$\widehat{V}(p) = \{\Phi \mid p \in \Phi\}$

Truth Lemma: $(\widehat{\mathcal{M}}, \Phi) \models \varphi$ iff $\varphi \in \Phi$.

Proof: By induction on φ. The base case follows from the definition of V. The boolean cases follow from the properties of MCSs. Formulas for accessibility (of the form aEb) follow immediately from the definition of $\widehat{E}(\Phi)$.

So first assume $(\widehat{\mathcal{M}}, \Phi) \models \langle K_a\rangle\varphi$. Then we have that there exists a Φ' such that $\Phi \widehat{\sim}_a \Phi'$ and $(\widehat{\mathcal{M}}, \Phi) \models \varphi$. But then $\varphi \in \Phi'$ and so $\langle K_a\rangle\varphi \in \Phi$.

For the opposite direction, assume $\langle K_a\rangle\varphi \in \Phi$. Then by the same equivalences as above, it suffices to find an MCS Φ' such that $\Phi \mathrel{\widehat{\sim}_a} \Phi'$ and $\varphi \in \Phi'$. We have this fact by the Existence lemma.

The case of K'_a is similar. Assume $(\widehat{\mathcal{M}}, \Phi) \models \langle K'_a\rangle\varphi$. Then we have that there exists a Φ' such that $\Phi \mathrel{\widehat{\sim}_a} \Phi'$ and $(\widehat{\mathcal{M}}, \Phi) \models \varphi$. But then $\varphi \in \Phi'$ and so $\langle K'_a\rangle\varphi \in \Phi$.

For the other direction, assume $\langle K'_a\rangle\varphi \in \Phi$. It suffices to find an MCS Φ' such that $\Phi \mathrel{\widehat{\sim}_{\langle a\rangle}} \Phi'$ and $\varphi \in \Phi'$, which we know by the Existence lemma.

Corollary: EAL is sound and complete for the canonical pseudo-model.

Proof: Suppose Σ is a consistent set of formulas from $\mathcal{L}_{EAL}$. By Lindenbaum's Lemma there is an MCS Σ^+ extending Σ. By the Truth Lemma we have that

$(\widehat{\mathcal{M}}, \Sigma^+) \models \Sigma$.

Before proving (2), we will need a lemma which utilizes further definitions.

Definition Let $\overline{\mathcal{M}}$ be a network pseudo model and let s be a state of $\overline{\mathcal{M}}$. We define the $\mathcal{L}_{EAL}$*-type of* $(\overline{\mathcal{M}}, s)$ to be the set of formulas $\varphi \in \mathcal{L}_{EAL}$ such that $(\overline{\mathcal{M}}, s) \models \varphi$.

Definition Let $\overline{\mathcal{M}} = (S, \sim_1, ..., \sim_n, \sim_{\langle 1 \rangle}, ..., \sim_{\langle n \rangle}, E, V)$ be a network pseudo model with $s, t \in S$. We call a sequence $\langle v_1, i_1, v_2, i_2, ..., i_{k-1}, v_k \rangle$ where $k \geq 1$ a *path from s to t* if:

(1) $v_1 = s$
(2) $v_k = t$
(3) $v_1, ..., v_k$ are states
(4) $i_1, ..., i_{k-1}$ are generalized agents from the set $\{1, ..., n, \langle 1 \rangle, ..., \langle n \rangle\}$.
(5) $(v_j, v_{j+1}) \in \sim_{i_j}$ for $1 \leq j < k$.

Definition The *reduction of a path* $\langle v_1, i_1, v_2, i_2, ..., i_{k-1}, v_k \rangle$ is the path formed by replacing each maximal consistent subsequence $\langle v_q, i_q, v_{q+1}, i_{q+1}, ..., i_{r-1}, v_r \rangle$ where $i_q = i_{q+1} = ... = i_{r-1}$ by $\langle v_q, i_q, v_r \rangle$. A reduction of a path is a path, by transitivity of every $\sim_i$. A path is *reduced* if it equals its reduction.

Definition $\overline{\mathcal{M}}$ is *tree-like* if whenever s and t are states of $\overline{\mathcal{M}}$, then there is at most one reduced path from s to t in $\overline{\mathcal{M}}$.

Lemma (Unravelling for pseudo models): Let $\overline{\mathcal{M}}_1 = (W, (\sim_a)_{a \in \mathcal{A}}, (\sim_{\langle a \rangle})_{a \in \mathcal{A}}, E, V)$ be a pseudo model of EAL. Then there is an pseudo model $\overline{\mathcal{M}}_2$ of n agents such that:

(1) $\overline{\mathcal{M}}_2$ is tree-like and
(2) $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_2$ have the same $\mathcal{L}_{EAL}$-types

Proof: Fix a particular $w \in W$. Then let $\overline{\mathcal{M}}_2 = (\overrightarrow{W}, (\overrightarrow{\sim_a})_{a \in \mathcal{A}}, E, \overrightarrow{V})$, where:

(i) $\overrightarrow{W}$ is the set of all finite sequences $\langle w, R_1, w_1, R_2, ..., R_2, w_n \rangle$ such that $wR_1w_1R_2...R_nw_n$ where each $R_i \in \{\sim_a |a \in \mathcal{A}\} \cup \{\sim_{\langle a \rangle} |a \in \mathcal{A}\}$.

(ii) For $\overrightarrow{s_1}, \overrightarrow{s_2} \in \overrightarrow{W}$, define $\sim'_a$ such that $\overrightarrow{s_1} \sim'_a \overrightarrow{s_2}$ if there exists v such that $\overrightarrow{s_1}; R'_a; v = \overrightarrow{s_2}$, where ; denotes sequence concatenation. Then $\overrightarrow{\sim_a}$ is defined to be the reflexive, symmetric, transitive closure of $\sim'_a$.

(iii) E is the same set of edges from $\overline{\mathcal{M}}_1$.

(iv) $\overrightarrow{V}(p) = \{(w_1, w_2, ..., w_n) \in \overrightarrow{W} \mid w_n \in V(p)\}$

We first show that $\overline{\mathcal{M}}_2$ is tree-like. Essentially, we use the fact that for each state $\overrightarrow{t}$ there is at most one state $\overrightarrow{s}$ such that $\overrightarrow{s} \overrightarrow{\sim}_a \overrightarrow{t}$.

Define a *primitive path* P to be a sequence $\langle v_1, i_1, v_2, i_2, ..., i_{k-1}, v_k \rangle$ where $k \geq 1$ a *path from s to t* if:

(1) $v_1 = s$
(2) $v_k = t$
(3) $v_1, ..., v_k$ are states
(4) $i_1, ..., i_{k-1}$ are generalized agents from the set $\{1, ..., n, \langle 1 \rangle, ..., \langle n \rangle\}$, and
(5) $(v_j, v_{j+1}) \in \sim'_{i_j}$ or $(v_{j+1}, v_j) \in \sim'_{i_j}$ for $1 \leq j < k$.

We say that P is an *a-primitive path* if $i_j = a$ for $1 \leq j < k$. We say that P is *nonredundant* if there is no j such that $v_j = v_{j+2}$ and $i_j = i_{j+2}$. Intuitively, P is nonredundant if there is no part of the path going along an edge and then immediately back again. Note that there is at most one nonredundant primitive path from $\overrightarrow{s}$ to $\overrightarrow{t}$, since for each there is at most one sequence $\overrightarrow{s'}$ such that $\overrightarrow{s} + \overrightarrow{s'} = \overrightarrow{t}$. Furthermore, note that $(\overrightarrow{s}, \overrightarrow{t}) \in \overrightarrow{\eqsim}_a$ iff there is an i-primitive path from $\overrightarrow{s}$ to $\overrightarrow{t}$.

Take arbitrary states $\overrightarrow{s}$ and $\overrightarrow{t}$. Let $P = \langle \overrightarrow{v}_1, i_1, ..., i_{k-1}, \overrightarrow{v}_k \rangle$ and $P' = \langle \overrightarrow{v}'_1, i'_1, ..., i'_{k-1}, \overrightarrow{v}'_k \rangle$ be reduced paths from s to t. Then since $(v_j, v_{j+1}) \in K'_{i_j}$ for $1 \leq j < k$, there must be a reduced i-primitive path from v_j to v_{j+1}. Let $\widehat{P}$ denote the path obtained by substituting in this nonredundant primitive path for each (v_j, v_{j+1}) in P. Since each of the primitive paths are nonredundant, the resulting path $\widehat{P}$ is also nonredundant. We have also that P is the reduction of $\widehat{P}$. Similarly, let $\widehat{P}'$ also be a nonredundant path from $\overrightarrow{s}$ to $\overrightarrow{t}$. By uniqueness of nonredundant primitive paths, we know that $\widehat{P} = \widehat{P}'$, which means that the reductions of $\widehat{P}$ and $\widehat{P}'$ are the same as well. So $P = P'$, and so $\overline{\mathcal{M}}_2$ must be tree-like.

To show that $\overline{\mathcal{M}}_2$ satisfies the same formulas as $\overline{\mathcal{M}}_1$, we can show that $\overline{\mathcal{M}}_2$ is a bounded morphic image of $\overline{\mathcal{M}}_1$. For if we let $f : \overrightarrow{W} \to W$ be defined by $f(w, w_1, ..., w_n) = w_n$, we can see that f is surjective, has the back and forth properties, and maps any $\overrightarrow{s}$ to a state in $\overline{\mathcal{M}}_1$ satisfying the same propositional variables.

Proof of (2): If φ is pseudo-satisfiable, then φ is satisfiable.

Assume that $\varphi \in \mathcal{L}_{EAL}$ is pseudo-satisfiable. By unravelling, we can assume without loss of generality that there is a tree-like pseudo-network model $\overline{\mathcal{M}}$ and a state s such that $(\overline{\mathcal{M}}, s) \models \varphi$. Let $\overline{\mathcal{M}}$ be $(S, (\eqsim_a)_{a \in \mathcal{A}}, (\eqsim_{\langle a \rangle})_{a \in \mathcal{A}}, E, V)$. We define $\sim_a$ for $a \in \mathcal{A}$, by: $(s, t) \in \sim_a$ iff there exists a finite chain $v_1 ... v_k$ such that:

(i) $v_1 = s$
(ii) $v_k = t$
(iii) $v_j \eqsim_a v_{j+1}$ or $v_j \eqsim_{\langle b_j \rangle} v_{j+1}$ for some $\eqsim_{b_j}$ such that $(b_j, a) \in E(v_j)$

Let $\mathcal{M} = (S, (\sim_a)_{a \in \mathcal{A}}, E, V)$. Note that $\mathcal{M}$ and $\overline{\mathcal{M}}$ have the same state space S, edge relation E and valuation function V. To complete the proof it suffices to show the truth of three claims.

Claim I: $\sim_a$ is an equivalence relation for each $a \in \mathcal{A}$

Claim II:$\sim_a$ satisfies the condition that $s \sim_a t$ and $(a, b) \in E(s)$ implies $(a, b) \in E(t)$.

Claim III: For every pseudo-model $\overline{M}$ there exists a network model $\mathcal{M}$ such that: $(\mathcal{M}, s) \models \varphi \Leftrightarrow (\overline{M}, s) \models \varphi$

Proof of I: We first note that the condition holds for any b such that $(b, a) \in E(s)$. That is, we can show that $s \overline{\sim}_{\langle b \rangle} t$ and $(b, a) \in E(s)$ implies $(b, a) \in E(t)$. Taking axioms P1 ($K_a\varphi \to K'_a\varphi$) and P3 ($aEb \to K_a aEb$) together, we have that that $bEa \to K'_b bEa$. Then by the definition of $\overline{\sim}_{\langle b \rangle}$ in the canonical pseudo-model, we have exactly the above fact: $s \overline{\sim}_{\langle b \rangle} t$ and $(b, a) \in E(s)$ implies $(b, a) \in E(t)$.

From our definition of pseudo-models, we also already have that $s \overline{\sim}_a t$ and $(a, b) \in E(s)$ implies $(a, b) \in E(t)$. These two conditions together guarantee that $E(s) = E(t)$ for any $(s, t) \in \sim_a$, since we have that $(v_j, v_{v+1}) \in \overline{\sim}_a$ or $\overline{\sim}_{\langle b \rangle}$ and in either case $E(v_j) = E(v_{j+1})$. Thus, each $(v_j, v_{j+1}) \in \sim_a$ is taken from a fixed set of equivalence relations. Furthermore $\sim_a$ includes *all* paths formed from elements of these equivalence relations. So for any $(s, t) \in \sim_a$ we also have the the returning path t to s, since these equivalence relations must contain for each edge (v_j, v_{j+1}) the returning edge (v_{j+1}, v_j). These equivalence relations will also have to contain reflexive edges (v_j, v_j) for all relevant states, which means $\sim_a$ will also be reflexive. Finally, $\sim_a$ is transitive since we can compose any two paths. Therefore $\sim_a$ is itself an equivalence relation.

Proof of II: We must prove that $s \sim_a t$ and $(a, c) \in E(s)$ together imply that $(a, c) \in E(t)$. We prove this inductively, on the length of the path from s to t. Assume that the condition holds up to state v_j. Then (v_j, v_{j+1}) is either in $\sim_a$ or in $\sim_{\langle b \rangle}$ for some b s.t. $(a, b) \in E(v_j)$. In the first case, we have by definition of a pseudo-model that:

$$v_j \overline{\sim}_a v_{j+1} \wedge (a, c) \in E(v_j) \to (a, c) \in E(v_{j+1})$$

So in this case condition (ii) holds. Now consider the second case, when $(v_j, v_{j+1}) \in \sim_{\langle b \rangle}$ for some b such that $b \in E(v_j)$. Note that by axioms E2 and E3 we can derive as a theorem:

$$bEa \to (aEc \to K'_b aEc)$$

So then for any b such that $(b, a) \in E(v_j)$, we have that $(\overline{\mathcal{M}}, v_j) \models bEa$ and thus:

$$v_j \overline{\sim}_{\langle b \rangle} v_{j+1} \wedge (a, c) \in E(v_j) \to (a, c) \in E(v_{j+1})$$

which means that in this case too we have that condition (ii) holds. This completes the induction, thus (ii) holds for any pair $(s, t) \in \sim_a$. Therefore (ii) holds in general and $\mathcal{M}$ is an network model.

Proof of III: And now we can show, by induction on the structure of formulas in $\mathcal{L}_{EAL}$, that:

(3) $(\mathcal{M}, s) \models \psi$ iff $(\overline{\mathcal{M}}, s) \models \psi$

If ψ is a primitive proposition then (3) is immediate, since $\mathcal{M}$ and $\overline{\mathcal{M}}$ have the same valuation function V. The case where ψ is a Boolean combination of formulas for which the analog of (3) holds is also immediate. Since $\mathcal{M}$ and $\overline{\mathcal{M}}$ also have the same edge relation the case when ψ is of the form aEb is also immediate. Now consider the case where we have a ψ of the form $K_a\gamma$.

($\Rightarrow$) (Contrapositive) Assume first that $(\overline{\mathcal{M}}, s) \not\models K_a\gamma$. Thus, there is a state $t \in S$ such that $(s,t) \in \overline{\sim}_a$ and $(\overline{\mathcal{M}}, t) \not\models \gamma$. Since $\overline{\sim}_a \subseteq \sim_a, (s,t) \in \sim_a$. So $(\mathcal{M}, t) \not\models \gamma$, and thus $(\mathcal{M}, s) \not\models K_a\gamma$.

($\Leftarrow$) Assume now that $(\overline{\mathcal{M}}, s) \models K_a\gamma$. To show that $(\mathcal{M}, s) \models K_a\gamma$ we must show that $(\mathcal{M}, t) \models \gamma$ whenever $(s,t) \in \sim_a$. Assume that $(s,t) \in \sim_a$. By definition of $\sim_a$ there exists a finite chain $v_1, ..., v_k \in S$ such that:

(1) $v_1 =$ s
(2) $v_k =$ t
(3) $v_j \overline{\sim}_a v_{j+1}$ or $v_j \overline{\sim}_{\langle b_j \rangle} v_{j+1}$ for some b_j such that $(b_j, a) \in E(v_j)$

We now show, by induction on j, (where $1 \leq j < k$) that $(\overline{\mathcal{M}}, v_j) \models K_a\gamma$. The case $j = 1$ is by assumption. Assume inductively on j that $(\overline{\mathcal{M}}, v_j) \models K_a\gamma$ (where $1 \leq j < k-1$). We need to show that $(\overline{\mathcal{M}}, v_{j+1}) \models K_a\gamma$. Since $(\overline{\mathcal{M}}, v_j) \models K_a\gamma$, it follows by positive introspection that $(\overline{\mathcal{M}}, v_j) \models K_a K_a\gamma$. We know that either $(v_j, v_{j+1}) \in \overline{\sim}_a$ or $(v_j, v_{j+1}) \in \sim_{\langle a \rangle}$ or $(v_j, v_{j+1}) \in \overline{\sim}_{\langle b \rangle}$ for some b such that $(b,a) \in E(w)$. In the first case, since $(\overline{\mathcal{M}}, v_j) \models K_a K_a\gamma$, it follows that $(\overline{\mathcal{M}}, v_{j+1}) \models K_a\gamma$, as desired. In the second case, we have that $(v_j, v_{j+1}) \in \overline{\sim}_{\langle b \rangle}$ for some b such that $(b,a) \in E(w)$. So by assumption we have that $(\overline{\mathcal{M}}, s) \models K_a\gamma$ and since $(b,a) \in E(w)$, we have also that $(\overline{\mathcal{M}}, s) \models bEa$. Then as an instance of Axiom E2 we have that $(K_a\gamma \wedge bEa) \rightarrow K'_b(K_a\gamma)$. So we have that $\mathcal{M}, s \models K'_b(K_a\gamma)$ and by Axiom T' for K' we get that $\overline{\mathcal{M}}, s \models K_a\gamma$.

This completes the induction. It follows that $(\overline{\mathcal{M}}, t) \models K_a\gamma$. Therefore by axiom (T) we have that $(\overline{\mathcal{M}}, t) \models \gamma$. So by inductive assumption, $(M, t) \models \gamma$.

Now consider the case when ψ is of the form $K'\gamma$.

($\Rightarrow$) (Contrapositive) Assume $(\overline{\mathcal{M}}, s) \not\models K'_a\gamma$. So there exists a state t such that $(s,t) \in \sim_{\langle a \rangle}$ and $(\overline{\mathcal{M}}, s) \not\models \gamma$, and thus by inductive assumption $(\mathcal{M}, s) \not\models \gamma$. Note that for all b such that $(a,b) \in E(s)$, we have by definition of $\sim_{\langle a \rangle}$ that: $\sim_{\langle a \rangle} \subseteq \sim_b$. So then $(s,t) \in \bigcap\{\sim_b |(\mathcal{M}, s) \models aEb\}$. And thus $(\mathcal{M}, s) \not\models K'_a\gamma$.

($\Leftarrow$) Assume $(\overline{\mathcal{M}}, s) \models K'_a\gamma$. To show that $(\mathcal{M}, s) \models K'_a\gamma$, we must show that for every t such that $s \sim_{\langle a \rangle} t$, $(\mathcal{M}, s) \models \gamma$. So we take such an arbitrary state t such that $s \sim_b t$ for all b such that $(a,b) \in E(s)$. Then for each $\sim_b$ in $\mathcal{M}$ such that $(a,b) \in E(s)$, there exists a reduced path in $\overline{\mathcal{M}}$ (based on our initial definition for $\mathcal{M}$):

$$P_b = \langle s = v_{1_b}, R_{1_b}, v_{2_b}, R_{2_b}, ..., v_{k_b} = t \rangle$$

with each $R_{i_b} \in \{\sim_i \text{ and } \sim_{\langle i \rangle} | i \in \mathcal{A}\}$ and $v_{j_b} R_{c_{j_b}} v_{(j+1)_b}$ for some c_{j_b} such that $(c_{j_b}, b) \in E(s)$.

However, by uniqueness of reduced paths, these paths from s to t must all be identical. So there exists a unique reduced path:

$$P = \langle s = v_1, R_1, v_2, R_2, ..., v_k = t \rangle$$

with $R_{i_b} \in \{\sim_i \text{ and } \sim_{\langle i \rangle} | i \in \mathcal{A}\}$, $v_j R_{c_j} v_{j+1}$ and $(c_j, b) \in E(s)$ for all b such that $(a, b) \in E(s)$. So this gives us that $(\overline{\mathcal{M}}, v_j) \models \bigwedge_{b \in \mathcal{A}} (aEb \Rightarrow c_j Eb)$ for each j. Thus by Axiom E4N we have:

$(*)$ $(\overline{\mathcal{M}}, v_j) \models K'_a \gamma \to K'_{c_j} \gamma$ for all j

Subclaim: $(\overline{\mathcal{M}}, v_j) \models K'_a \gamma$ for all j.

Proof: Induction on j. For $j = 1$, the statement is true by initial assumption since $v_1 = s$ and $(\overline{\mathcal{M}}, s) \models K'_a \gamma$. For the induction step, suppose $(\overline{\mathcal{M}}, v_j) \models K'_a \gamma$. Then, by axiom 4: $(\overline{\mathcal{M}}, v_j) \models K'_a K'_a \gamma$. By $(*)$ with axiom E4 we have:

$$(\overline{\mathcal{M}}, v_j) \models K'_a \gamma \to K'_{c_j} K'_a \gamma$$

and taken with $v_j R_{c_j} v_{j+1}$ we have that:

$$(\overline{\mathcal{M}}, v_{j+1}) \models K'_a \gamma$$

This completes the proof of the subclaim.

Apply the subclaim to $j = k$ to get $(\overline{\mathcal{M}}, t) \models K'_a \gamma$. Then by axiom T' we have $(\overline{\mathcal{M}}, t) \models \gamma$. By inductive assumption $(\mathcal{M}, t) \models \gamma$ as well, which completes the proof.

Chapter 4

Iterated Access Logic

In the previous chapter we introduced a modality for the information an agent has access to via her immediate network connections. In this chapter we expand the language with modalities for capturing all the knowledge reachable by a fixed number of edges n. The operators will be of the form "K_a^n," and will express the combined knowledge of a and any agent within n edges of a. The agent's own knowledge will be recovered as K_a^0, and the information previously picked out as K_a' will now be described as K_a^1. In this logic, the n-th level of potential knowledge corresponds to what could be known after a n rounds of information travelling from each agent to those who access the agent. This can be thought of as the "degrees of separation" between the agent and an agent with the knowledge in question. Just as in the previous logic, we use the term aEb to express that agent a has direct access to b's information.

4.1 Introducing IAL

Definition (Syntax of Iterated Access Logic)

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \ \vee \ \varphi \mid K_a^n\varphi \mid aEb$$

As in the last chapter, we use "$\langle K_a^n \rangle\varphi$" to abbreviate $\neg K_a^n \neg\varphi$. In addition to this syntax, we recursively define abbreviations to capture connection via paths longer than a single edge. Intuitively, $aE^n c$ holds if there is a directed path from a to c of length at most n:

$$aE^0c := \top \text{ for } a = c, \bot \ \text{ for } a \neq c$$

$$aE^{n+1}c := \bigvee_{b \in \mathcal{A}} (aEb \wedge bE^n c)$$

The formulas of this language, $\mathcal{L}_{\mathsf{IAL}}$, are interpreted on the same network models as were used in interpreting $\mathcal{L}_{EAL}$. However, just as we needed to define an additional relation for our potential knowledge modality K_a', we now need to define relations corresponding to the each K_a^n. So we additionally define, for each relation $\sim_a$, all the relations $\sim_a^n$ to represent the information available to a within n edges. For a fixed k, the relation $\sim_a^k$ captures the distributed knowledge of agent a and those agents connected to her by a path of length at most k:

$(s,t) \in\sim_a^n$ iff $(s,t) \in \sim_b$ for all b such that $(a,b) \in E^n(s)$

Finally, we also generalize the idea of access-dominance for this context, using "$a \sqsubseteq_n b$" to abbreviate $(\bigwedge_{c\in\mathcal{A}} aEc \rightarrow bEb)$. Intuitively, $a \sqsubseteq_n b$ means that every agent a is connected to within n edges is also connected to b within n edges. We continue to use "$a \sqsubseteq b$" for the case of $a \sqsubseteq_1 b$. An important consequence of the semantics below is that $(a \sqsubseteq_n b) \rightarrow (a \sqsubseteq_{n+1} b)$ for $n \geq 1$. However one can even see this on an intuitive level: assuming I can reach all the same people you can in n steps, I can follow any further step you take.

We now define truth for formulas of $\mathcal{L}_{IAL}$. The semantics is essentially the same as in the previous logic except that the new relations are utilized in defining truth for formulas of the form $K_a^n\varphi$.

Definition (Semantics of $\mathcal{L}_{\text{IAL}}$):

$\mathcal{M}, w \models p$ iff $w \in V(p)$

$\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$

$\mathcal{M}, w \models \varphi \vee \psi$ iff $\mathcal{M}, w \models \varphi$ or $\mathcal{M}, w \models \psi$

$\mathcal{M}, w \models aEb$ iff $(a,b) \in E$

$\mathcal{M}, w \models K_a^n\varphi$ iff for all v such that $w \sim_a^n v$ we have that $\mathcal{M}, w \models \varphi$

4.2 A Second Database Example

To illustrate how these new modalities can be applied, we consider a particular model. As in the previous example, each agent in this scenario has a computer. Some agents can secretly access the computers of other agents, illustrated by directed arrows.

In the scenario, the agents to the right – b, c, and d – are the only agents with information. However, as with the last example, this does not imply that these agents have more information available to them. Although a possesses no knowledge of her own, a has more accessible information than any of the agents supplying the information in the first place. We have that $b \sqsubseteq a$, $c \sqsubseteq a$, $d \sqsubseteq a$, and $e \sqsubseteq a$.

Interestingly, however, agent a *does* increase the information that the agent b, c, and d in an extended sense. While "sharing" information with a does not result in any information directly from a, it is still beneficial because agents b, c, and e are able to gain information from each other after a has acquired it. Agent a serves as a middleman for information between these three agents.

The last agent in the scenario, agent e, is a silent observer in this exchange – a hacker, say. By accessing a's computer, e is able to have the same second-level accessibility as the information-contributing agents on the right side of the diagram. That is to say $b \sqsubseteq e$, $c \sqsubseteq e$, and $d \sqsubseteq e$. Agent e has tapped a computer that contains no information, but accessing this computer allows e to learn anything agent a learns in his role as a kind of "hub" for the other agents on the right. Hacking any of the agents on the right would not get e secondary-access to all of the formulas in this way.

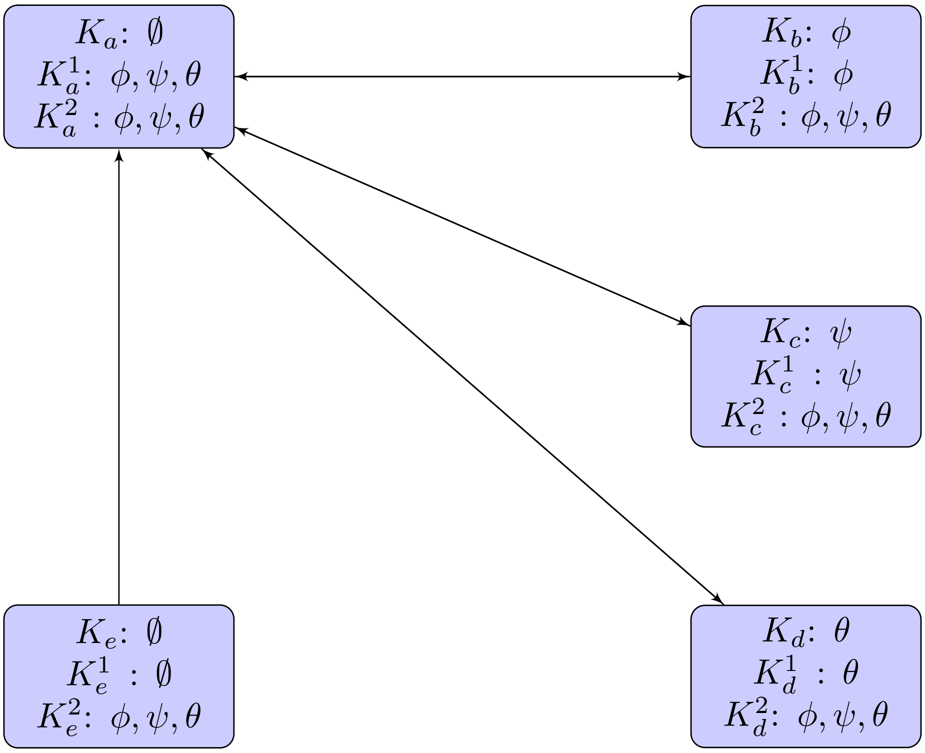

Figure 4.1: Agents are represented as boxes, labelled with the formulas each agent knows in the sense of K, K^1, and K^2. Arrows indicate access, originating from the agent with access.

4.3 Axiomatization of IAL

(MP) $\vdash_{IAL} \varphi, (\varphi \to \psi) \Rightarrow \vdash_{IAL} \psi$

(N) $\vdash_{IAL} \varphi \Rightarrow \vdash_{IAL} K_a\varphi$

(Prop) All validities of propositional logic

(KN) $K_a^n(\varphi \to \psi) \to (K_a^n\varphi \to K_a^n\varphi)$

(TN) $K_a^n\varphi \to \varphi$

(4N) $K_a^n\varphi \to K_a^nK_a^n\varphi$

(5N) $\neg K_a^n\varphi \to K_a^n\neg K_a^n\varphi$

(E1N) aEa

(E2N) $(aEb \wedge K_b^n) \to K_a^{n+1}\varphi$

(E3N) $aEb \to K_a aEb$

(E4N) $\bigwedge_{b\in A}(aE^nb \to cE^mb) \to (K_a^n\varphi \to K_a^m\varphi)$

4.4 Completeness of IAL

Proposition: IAL is sound and complete for epistemic network frames.

Proof: Much of the completeness proof of IAL resembles the completeness proof of EAL. As before, we define pseudo-models which have extra relations to pseudo-satisfy potential knowledge formulas from $\mathcal{L}_{\mathsf{IAL}}$ directly, as if they were just additional agents in the model. However, this time our pseudo-models will have additional relations $\overline{\sim_a^n}$ for every n and a (instead of just one additional relation per agent). We can then again build a canonical pseudo-model, taking states to be maximally consistent sets of formulas from $\mathcal{L}_{\mathsf{IAL}}$. The canonical valuation and canonical relations formed from each K_a^n operator give us a structure which can pseudo-satisfy any consistent formula of $\mathcal{L}_{\mathsf{IAL}}$. We then prove there exists an epistemic network model satisfying the same formulas.

Definition We say that a formula φ is *consistent* if its negation $\neg\varphi$ cannot be proven in IAL. Otherwise we say that φ is *inconsistent.*

Definition A formula φ is *satisfiable* if there is a network model and a state s such that $(\mathcal{M}, s) \models_{IAL} \varphi$. The formula φ is then said to be satisfiable in $\mathcal{M}$.

Definition A *network pseudo-model* for a set of $\mathcal{A}$ of k agents is an epistemic network model of countably many agents, where $n \geq 0$:

$$\overline{\mathcal{M}} = (S, (\sim_a^n)_{a\in\mathcal{A}}, E, V)$$

for which $\sim_a^n$ are equivalence relations for all a and all n.

Definition Where $\overline{\mathcal{M}}$ is a network pseudo-model, we say that $(\overline{\mathcal{M}}, s)$ *pseudo-satisfies* a formula φ if the pair $(\overline{\mathcal{M}}, s)$ satisfies φ in the normal way except that $(\overline{\mathcal{M}}, s) \models K_a^n \varphi$ iff $(\overline{\mathcal{M}}, s') \models \varphi$ for all s' such that $(s, s') \in \underline{\sim_a^n}$. A pseudo-model $\overline{\mathcal{M}}$ *validates* a formula φ if $(\overline{\mathcal{M}}, s) \models \varphi$ for every state s of $\overline{\mathcal{M}}$.

Definition A pseudo-model of IAL is a pseudo-model which validates all the axioms of IAL.

Proposition: Every model is a pseudo-model of IAL.

Proof: Compare with proof of validity of pseudo-models in EAL.

Definition A formula φ is pseudo-satisfiable if there is a network pseudo-model $\overline{\mathcal{M}}$ and a state s of $\overline{\mathcal{M}}$ such that $(\overline{\mathcal{M}}, s) \models \varphi$. The formula φ is then said to be pseudo-satisfiable in $\overline{\mathcal{M}}$.

In order to prove completeness for IAL it suffices to prove the following propositions:

(1) If φ is consistent, then φ is pseudo-satisfiable.
(2) If φ is pseudo-satisfiable, then φ is satisfiable.

Proof of (1): If φ is consistent, then φ is pseudo-satisfiable.

We will use the canonical model construction to show that every consistent set of formulas is pseudo-satisfiable.

Definition: A set of formulas Φ is maximal consistent if Φ is consistent and any set of formulas properly containing Φ is inconsistent. If Φ is a maximal consistent set of formulas then we say it is an MCS.

Lindenbaum Lemma: For any consistent set Φ of formulas from $\mathcal{L}_{\mathsf{IAL}}$, there is an MCS Φ^+ such that $\Phi \subseteq \Phi^+$.

Proof: Compare with proof from EAL

Definition: We define the canonical network pseudo-model for a set of agents $\mathcal{A}$ to be $\widehat{\mathcal{M}} = (\widehat{W}, (\widehat{\sim}_a)_{a \in \mathcal{A}}, (\widehat{\sim}_{\{a\}})_{a \in \mathcal{A}}, \widehat{E}, \widehat{V})$, where:

$\widehat{W} = \{\Phi \mid \Phi \text{ is an MCS }\}$

$\Phi \mathrel{\widehat{\sim}_a^n} \Phi'$ iff for all ψ, $(K_a^n \psi \in \Phi) \to (\psi \in \Phi')$

$\widehat{E}(\Phi) = \{(a, b) \mid aEb \in \Phi\}$

$\widehat{V}(p) = \{\Phi \mid p \in \Phi\}$

Proposition: The canonical network pseudo-model is a pseudo-model of IAL.

Proof: Compare with proof from EAL

Existence Lemma: For any state Φ, if $\langle K_a^n\rangle\varphi \in \Phi$ then there is a state $\Phi' \in \widehat{W}$ such that $\Phi \widehat{\sim}_a^n \Phi'$.

Suppose $\langle K_a^n\rangle\varphi \in \Phi$. Then we can construct a state Φ' such that $\Phi \widehat{\sim}_a^n \Phi'$ and $\varphi \in \Phi$. Let Φ'^- be $\{\varphi\} \cup \{\psi \mid K_a^n\psi \in \Phi\}$. Then Φ' is consistent. Assume for contradiction that this is not the case. Then there would be $\psi_1, ..., \psi_n$ such that $\vdash_{IAL} K_a^n(\psi_1 \wedge ... \wedge \psi_n) \rightarrow \neg\varphi$. And so we have $\vdash_{IAL} K_a^n(\psi_1 \wedge ... \wedge \psi_n) \rightarrow K_a^n\neg\varphi$. By propositional calculus we get $\vdash_{EAL} K_a^n(\psi_1 \wedge ... \wedge \psi_n) \rightarrow K_a^n\neg\varphi$. So now, $K_a^n\psi_1 \wedge ... \wedge \psi_n \in \Phi$. But this is impossible since Φ is an MCS containing $\langle K_a^n\rangle\varphi$. So Φ'^- must be consistent. Let Φ' be any MCS extending Φ'^-. By construction, $\varphi \in \Phi'$. For all formulas ψ, $K_a^n\psi \in \Phi$ implies $\psi \in \Phi'$. So by definition $\Phi \widehat{\sim}_a^n \Phi'$.

Truth Lemma: $(\widehat{\mathcal{M}}, \Phi) \models \varphi$ iff $\varphi \in \Phi$.

Proof: By induction on φ. The base case follows from the definition of V. The boolean cases follow from the properties of MCSs. Formulas for accessibility (of the form aEb) follow immediately from the definition of $\widehat{E}(\Phi)$. The only cases left to consider ar of the form $K_a^n\varphi$.

So first assume $(\widehat{\mathcal{M}}, \Phi) \models \langle K_a^n\rangle\varphi$. Then we have that there exists a Φ' such that $\Phi \widehat{\sim}_a^n \Phi'$ and $(\widehat{\mathcal{M}}, \Phi) \models \varphi$. But then $\varphi \in \Phi'$ and so $\langle K_a^n\rangle\varphi \in \Phi$.

For the opposite direction, assume $\langle K_a^n\rangle\varphi \in \Phi$. Then by the same equivalences as above, it suffices to find an MCS Φ' such that $\Phi \widehat{\sim}_a^n \Phi'$ and $\varphi \in \Phi'$. We have exactly this fact by the Existence lemma.

Corollary: IAL is sound and complete for the canonical pseudo-model.

Proof: Suppose Σ is a consistent set of formulas from $\mathcal{L}_{IAL}$. By Lindenbaum's Lemma there is an MCS Σ^+ extending Σ. By the Truth Lemma we have that $(\widehat{\mathcal{M}}, \Sigma^+) \models \Sigma$.

We now have to show that pseudo-satisfiability implies satisfaction in an epistemic network model. For any given formula φ, we know there exists a pseudo-model $\mathcal{M} = (W, (\overline{\sim^n}_a)_{a\in\mathcal{A}}, E, V)$ satisfying it. We can use unravelling as in the previous proof of completeness, however we must specify the state space as:

> $\overrightarrow{W}$ is the set of all finite sequences $\langle w, R_1, w_1, R_2, ..., R_2, w_n\rangle$ such that $wR_1w_1R_2...R_nw_n$ where each $R_i \in \{\sim_a^n |a \in \mathcal{A}\}$.

Then by unravelling we can assume without loss of generality that $\overline{\mathcal{M}}$ is tree-like. We now define an epistemic network model $\mathcal{M}$ which satisfies the same formulas that $\overline{\mathcal{M}}$ pseudo-satisfies. Define that $(s,t) \in \sim_a$ iff there exists a finite chain of states $v_1, ..., v_k$ such that:

(i) $v_1 = s$
(ii) $v_k = t$
(iii) $v_j \overline{\sim}_b^{n_j} v_{j+1}$ for some agent b and some $n_j \in \mathbb{N}$ such that $(b,a) \in E^{n_j}(v_j)$.

Then let $\mathcal{M} = (W, (\sim_a)_{a \in \mathcal{A}}, E, V)$. To finish the proof, it suffices to show the truth of three claims:

Claim I: Each $\sim_a$ is an equivalence relation.

Claim II: Each $\sim_a$ satisfies the condition that $s \sim_a t$ and $(a,b) \in E(s)$ implies $(a,b) \in E(t)$.

Claim III: For each pseudo-model $\overline{\mathcal{M}}$ there exists a network model $\mathcal{M}$ such that $(\mathcal{M}, s) \equiv (\overline{\mathcal{M}}, s)$.

Proof of Claim I and II: Follows by induction on j, since for each step we have the axiom $bEa \to K_b(bEa)$. Then as in the last proof, we are taking the intersection of several equivalence relations, giving us again an equivalence relation.

Proof of III: Need to show that for every pseudo-model $\overline{M}$ there exists a network model $\mathcal{M}$ such that:

$$(\mathcal{M}, s) \models \varphi \Leftrightarrow (\overline{M}, s) \models \varphi$$

By induction on structure of φ. If φ is a Boolean formula, this is trivial. Similarly if φ is of the form aEb. Now consider if φ is of the form $K_a^n \psi$.

($\Rightarrow$) (Contrapositive) Assume $(\overline{\mathcal{M}}, s) \not\models K_a^n \psi$. So there exists a state t such that $(s,t) \in \sim_a^n$ and $(\overline{\mathcal{M}}, s) \not\models \psi$, and thus by inductive assumption $(\mathcal{M}, s) \not\models \psi$. Note that for all b such that $(a,b) \in E^n(s)$, we have by definition of $\sim_a^n$ that: $\sim_a^n \subseteq \sim_b$. So then $(s,t) \in \bigcap\{\sim_b | (\mathcal{M}, s) \models aE^n b\}$. And thus $(\mathcal{M}, s) \not\models K_a^n \psi$.

($\Leftarrow$) Assume $(\overline{\mathcal{M}}, s) \models K_a^n \psi$. To show that $(\mathcal{M}, s) \models K_a^n \psi$, we must show that for every t such that $s \sim_a^n t$, $(\mathcal{M}, s) \models \psi$. So we take such an arbitrary state t, and by definition of $\sim_a^n$ we have that $s \sim_b t$ for all b such that $(a,b) \in E^n(s)$. Then for each $\sim_b$ in $\mathcal{M}$ such that $(a,b) \in E^n(s)$, there exists a reduced path in $\overline{\mathcal{M}}$ (based on our initial definition for $\mathcal{M}$):

$$P_b = \langle s = v_{1_b}, \sim_{c_{1_b}}^{m_{1_b}}, v_{2_b}, \sim_{c_{2_b}}^{m_{2_b}}, ..., v_{k_b} = t \rangle$$

with $v_{j_b} \sim_{c_{j_b}}^{m_{j_b}} v_{(j+1)_b}$ for some c_{j_b} such that $(c_{j_b}, b) \in E^{m_{j_b}}(s)$.

However, by uniqueness of reduced paths, these paths from s to t must all be identical. So there exists a unique reduced path:

$$P = \langle s = v_1, \sim_{c_1}^{m_1}, v_2, ..., \sim_{c_2}^{m_2}, ..., v_k = t \rangle$$

with $v_j \sim_{c_j}^{m_j} \underline{v_{j+1}}$ and $(c_j, b) \in E^{m_j}(s)$ for all b such that $(a,b) \in E^n(s)$. So this gives us that $(\overline{\mathcal{M}}, v_j) \models \bigwedge_{b \in \mathcal{A}} (aE^n b \Rightarrow c_j E^{m_j} b)$ for each j.

Thus by Axiom E4N we have:

($*$) $(\overline{\mathcal{M}}, s) \models K_a^n \psi \to K_{c_j}^{m_j} \psi$ for all j

Subclaim: $(\overline{\mathcal{M}}, v_j) \models K_a^n \psi$ for all j.

Proof: Induction on j. For $j = 1$, the statement is true by initial assumption since $v_1 = s$ and $(\overline{\mathcal{M}}, s) \models K_a^n \psi$. For the induction step, suppose $(\overline{\mathcal{M}}, v_j) \models K_a^n \psi$. Then, by axiom 4: $(\overline{\mathcal{M}}, v_j) \models K_a^n K_a^n \psi$. By $(*)$ with axiom E4N we have:

$$(\overline{\mathcal{M}}, v_j) \models K_a^n \psi \to K_{c_j}^n K_a^n \psi$$

and taken with $v_j \sim_{c_j}^{m_j} v_{j+1}$ we have that:

$$(\overline{\mathcal{M}}, v_{j+1}) \models K_a^n \psi$$

This completes the proof of the subclaim.

Apply the subclaim to $j = k$ to get $(\overline{\mathcal{M}}, t) \models K_a^n \psi$. Then by axiom T we have $(\overline{\mathcal{M}}, t) \models \psi$. By inductive assumption $(\mathcal{M}, t) \models \psi$ as well, which completes the proof.

All-Read Update

Our basic update is the action whereupon every agent in the network incorporates the information of her neighbors. We will refer to this update as "all-read" update, and denote the corresponding dynamic modality by $[READ]$. This can be imagined to be the passing of one round of inquiry, in which all agents first gather all the information they can find and then add it to their own databases. An idealizing assumption is that these actions occur strictly in this order, so that no agent is searching for information while another updates. We can imagine that agents can only search for information during the day, and update their own databases in the evenings.

Since we assume agents read *all* the information they have access to, the $READ$ update is deterministic. For a given epistemic network model, there is only one possible resulting model. We assume agents know this, and know when $READ$ has occurred. So then ignorance of agents in the resulting model always comes from ignorance in the original model, and we will have the same state space. Each possible world in the original model is mapped to a unique world in the resulting model.

Definition ($READ$ Update)

The update $READ$ maps any model $\mathbf{S} = (S, (\sim_a)_{a \in \mathcal{A}}, V)$ to a new model $\mathbf{S}^{READ} = (S', (\sim'_a)_{a \in \mathcal{A}}, V')$ given by:

$$S' := S$$

$$\text{For all } a \in \mathcal{A} : s \sim'_a t \text{ iff } s \sim_a t \text{ and } s \sim_b t \text{ for all } b \in \mathcal{A} \text{ such that } (a, b) \in E$$

$$V'(p) := V(p) \cap S'$$

The transition relation $\longrightarrow^{READ}$ relates any state $s \in \mathbf{S}$ satisfying φ to the same state in the model $\mathbf{S}^{READ}$.

This action naturally connects to the static modalities K_a^n, since updating with $[READ]$ makes each formula φ potentially known in the sense of $K_a^n\varphi$ to be known in the sense of $K_a^{n-1}\varphi$. We add the following reduction laws to the axioms for IAL to arrive at a full axiomatization for DIAL:

Reduction Laws for $[READ]$

$$[READ]p \Longleftrightarrow p$$

$$[READ]\neg\varphi \Longleftrightarrow \neg[READ]\varphi$$

$$[READ](\phi \wedge \psi) \Longleftrightarrow [READ]\varphi \wedge [READ]\psi$$

$$[READ]K_a^n\varphi \Longleftrightarrow K_a^{n+1}[READ]\varphi$$

$$[READ]aEb \Longleftrightarrow aEb$$

4.5 A Classroom Example

As an example, we can consider the test day scenario mentioned in the introduction. Below we depict a group of students sitting for a test, along with what they know and have access to. We assume students can read the answers off of the person sitting in front or horizontal to them, but not behind. We assume all the relevant information in this scenario regards the test questions, so that agents write all they know on their answer sheet.

To ensure that information change occurs in a synchronized manner, we also assume a teacher is watching over these students. Each time the teacher turns away, the students all read their neighbors answers. They are then able to write down these answers on their own paper. This dynamic corresponds to the $READ$ update introduced in the previous section, so we are able to model this interaction in the diagram below.

The first thing to note is that accessible information can change without a change in actual knowledge. Agent d has the same knowledge after $READ$ as before, but e has new information which is now available to d. Secondly, an agent's knowledge can change without a change in accessible information. Agent b learns ψ from a, but there is no new information on the computers b has access to. Knowledge and accessible information can both change after $READ$, as with agent e. Agent e learns ψ from c, but c in the meantime also learns. Finally, knowledge and accessible information can remain completely unchanged, as with agent a. Agent a has no outward-directed edges, so she does not have access to anyone else's information.

It is also worth noting that if $READ$ is applied enough times – if the teacher turns away repeatedly – students will soon stop gaining knowledge or accessible information. Among the (reduced) directed paths from one agent to another along edges of access, there will be some longest path of length n. Regardless of the information given originally, repetitions of $READ$ after n times will have no effect on the model.

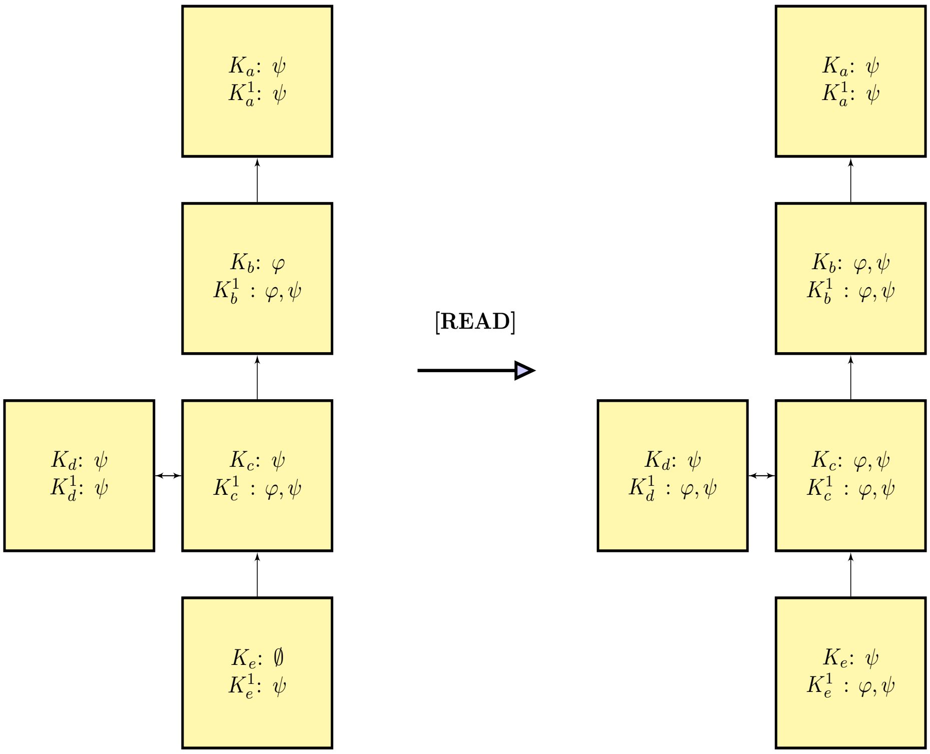

Figure 4.2: Diagram of a the classroom model $\mathcal{M}$ before and after update with $[READ]$. Students are represented by blocks, each labelled with the corresponding information the student knows and has access to.

Chapter 5

Restricted Access Logic

In the next two chapters, we explore less developed logics for which completeness has not been proven.

In the previous chapters, we have assumed that access between agents implies *full* access to the information of another or or none at all. We have this kind of access when we have a password to someone else's computer, or when we can peek at their answer sheet, but it is not typically the case. Most often, we only have access to a small portion of someone else's information. When relying on a friend or colleague, we cannot not ask for everything she knows or download everything contained on her computer. We typically only communicate about a limited set of issues which are especially useful to us. We can get closer to this picture by defining a more nuanced syntax and semantics which distinguishes access to knowledge on a formula-by-formula basis.

5.1 Introducing RAL

On the syntax side, we can replace the access relation P with a function that returns the relevant issues for any given pair. This set of issues can be thought of as the queries that the first agent is to the second, so we refer to the new models as "interrogative."

Definition (Syntax for Restricted Access Logic)

$\varphi \;:=\; p \mid \neg\varphi \mid \varphi \,\vee\, \varphi \mid K'_a\varphi \mid Q_{ab}\varphi$

where a, b are agents in $\mathcal{A}$. The new terms q_{ab} and $a(b) - c$ are added for the sake of completeness. Each formula of the form q_{ab} will serve as the canonical witness for the existence of a formula φ such that $Q_{ab}\varphi$. In other words, whenever $Q_{ab}\varphi$ we have that $Q_{ab}q_{ab}$ as well. Similarly, $a - c$ will serve as the canonical witness for the existence of a formula φ such that a can access φ from some agent b while c cannot. These formulas will be explained in more detail further on.

As mentioned above, the models for formulas of $\mathcal{L}_{RAL}$ will replace the access relation between agents with a function on pairs of agents, returning the potential information the first has access to via the second. We can represent a lack of access as access to the empty set, so that this function fully replaces the edge sets used in previous models.

Definition An *interrogative network model* based on a set of agents $\mathcal{A}$ is a quadruple:

$$(S, (\sim_a)_{a\in\mathcal{A}}, Q, V)$$

where as before $S \neq \emptyset$ is a set of states, for each $i \in \mathcal{A}$, $\sim_i$ is a binary equivalence relation on W, and V : Prop $\rightarrow \mathcal{P}(W)$ is a valuation. The Q function, the new component, gives the accessible information from each For a given state, $Q : W \rightarrow (\mathcal{A} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{L}_{ENL}))$ returns the function assigning a set of issues to each pair of agents for that state.

To capture the accessible information for agents in this more restricted framework, we also define supplementary relations for "what b knows in regards to φ." This will be denoted by $s \sim_b^{\varphi} t$, and will be defined by:

$$\begin{aligned}(s \sim_b^{\varphi} t) \Longleftrightarrow &(s \models K_b\varphi \wedge t \models K_b\varphi)\\ &\vee (s \models K_b\neg\varphi \wedge t \models K_b\neg\varphi)\\ &\vee (s \models \neg K_b\varphi \ \wedge \ s \models \neg K_b\neg\varphi \ \wedge \ t \models \neg K_b\varphi \ \wedge \ t \models \neg K_b\neg\varphi)\end{aligned}$$

Intuitively, this definition covers the three possible cases when a pair of states would be indistinguishable for b (restricted to a particular issue). Either b knows φ in both states, or b knows $\neg\varphi$ in both states, or b is ignorant in both states. We can now use these relations, which represent an agent's knowledge on a particular issue, to define a relation for accessible information. Intuitively, we combine a's information with the knowledge that other agents have about φ determined by Q. Let $\sim_a'$ be formally defined as:

$(w, v) \in\sim_a'$ iff $(w, v) \in\sim_a \cap \{\bigcap \sim_b^{\varphi} | \varphi \in Q_w(a, b)\}$.

Using these new relations, we define the semantics of $\mathcal{L}_{\mathsf{RAL}}$ as:

Definition (Semantics of $\mathcal{L}_{\text{RAL}}$):

$\mathcal{M}, w \models p$ iff $w \in V(p)$

$\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$

$\mathcal{M}, w \models \varphi \vee \psi$ iff $\mathcal{M}, w \models \varphi$ or $\mathcal{M}, w \models \psi$

$\mathcal{M}, w \models Q_{ab}\varphi$ iff $\varphi \in Q_w(a, b)$

$\mathcal{M}, w \models K_a'\varphi$ iff for all v such that $w \sim_a' v$ we have that $\mathcal{M}, w \models \varphi$

$\mathcal{M}, w \models q_{ab}\varphi$ iff there existsφ such that $\varphi \in Q_w(a, b)$

5.2 Axioms for RAL

It is possible to capture most of the axioms of EAL within the syntax of RAL, giving us a set of familiar axioms.

(Q1) $(Q_{ab}\varphi \wedge K_b\varphi) \rightarrow (K'_a K_b\varphi)$

(Q2) $(Q_{ab}\varphi \wedge K_b\neg\varphi) \rightarrow (K'_a K_b\neg\varphi)$

(Q3) $(Q_{ab}\varphi \wedge \neg K_b\varphi \wedge \neg K_b\neg\varphi) \rightarrow K'_a(\neg K_b\varphi \wedge \neg K_b\neg\varphi)$

Axioms (Q1)-(Q3) serve to connect an agent a's knowledge to that of a connection b, in each of the three knowledge conditions for b. These three axioms take over the role of (E2) in EAL

(Q4) $Q_{ab}\varphi \rightarrow K_a Q_{ab}\varphi$

Axiom (Q4) ensures that all agents know what formulas φ they are able to ask of every other agent. This is comparable to axiom (E3) from (EAL), which ensured that agents knew their outgoing edges.

(K) $K_a(\varphi \rightarrow \psi) \rightarrow (K_a\varphi \rightarrow K_a\varphi)$

(T) $K_a\varphi \rightarrow \varphi$

(4) $K_a\varphi \rightarrow \varphi$

(5) $\neg K_a\varphi \rightarrow K_a\neg K_a\varphi$

These axioms are standard knowledge axioms for K_a. They ensuring that K_a distributes over conditionals, that $K_a\varphi$ implies the truth of φ and that agents have positive introspection and negative introspection of K_a.

(K') $K'_a(\varphi \rightarrow \psi) \rightarrow (K'_a\varphi \rightarrow K'_a\varphi)$

(T') $K'_a\varphi \rightarrow \varphi$

(4') $K'_a\varphi \rightarrow \varphi$

(5') $\neg K'_a\varphi \rightarrow K'_a\neg K'_a\varphi$

These axioms are standard knowledge axioms for K'_a. They ensuring that K'_a distributes over conditionals, that $K_a\varphi'$ implies the truth of φ and that agents have positive introspection and negative introspection of K'_a.

(P1) $K_a\varphi \rightarrow K'_a\varphi$

For Future Work: A Complete Axiomatization of RAL

The main obstacle to having a complete axiomatization for RAL is that the syntax expresses access on a formula-by-formula basis. There is no way of expressing "a has no access to b" in this syntax without involving existential quantifiers. However, Skolemization offers a method for expressing just what we need. We can add to our interrogative models the functions q_{ab} and $a - c$, and specify these functions to return the witnesses we need for eliminating existential quantifiers. The function $q_a b$ will return, for a given state, a formula such that $Q_a b(\varphi)$ if such a formula exists. The function $a - c$ will return, for a given state, a formula φ such that a has access to and c does not. Then we can mimic the rest of the EAL axioms by writing:

(S1) $Q_{ab}\varphi \to Q_{ab}(q_{ab})$

(S2) $(a \sqsubseteq c) \to (Q_{ab}\varphi \to Q_{cb}\varphi)$

(S3) $(a \not\sqsubseteq c) \to \neg \bigvee_{b \in \mathcal{A}} (Q_{ab}(a - c) \land \neg Q_{cb}(a - c))$

(Q5) $\bigwedge_{b \in \mathcal{A}} (\neg Q_{ab} q_{ab}) \to (K'_a \varphi \to K_a \varphi)$

However, having a replica of every axiom of the complete logic EAL, does not make this proof system complete. So the whether there is a complete set of RAL axioms is still an open question.

Chapter 6

Iterated Friendship Logic

Epistemic Friendship Logic (EFL) is a recently developed framework for modelling social knowledge and interaction.[1] One distinguishing feature of EFL is that formulas in this language are interpreted relative to a state *and an agent*. This approach allows EFL to consider indexical propositions such as, "I am your father," which can be true when expressed by one individual but false when expressed by others. A second main feature is the use of agent nominals, i.e. propositions that are true only when evaluated at a particular agent. Rather than expressing formulas that hold of *some* agent, we are then able to express that a formula holds of any particular agent. In this chapter I introduce the standard logic of EFL, then a version with iterated versions of the modalities for friendship and knowledge. This version of EFL will be called "iterated friendship logic" (IFL).

6.1 Epistemic Friendship logic

The language $\mathcal{L}_{EFL}$ differs significantly from the languages in the previous chapters. First, there are two sets of primitive propositions rather than one. The set Prop is a set of general indexical propositions such as "I am in danger." Aside from being indexical, this is the typical set of primitive propositions. The second set ANOM is a set of indexical propositions asserting identity such as "I am n." These propositions enable us to assert things about specific agents in the model. A primitive proposition only true of n allows us to express things like "All of my friends know n is a bowling enthusiast."

Secondly, there are no modalities expressing the knowledge of a fixed agent (K_a) or formulas expressing that one agent can access the information of another (aEb). The new modalities used in $\mathcal{L}_{EFL}$ are K, F and A. K is a general knowledge operator – since formulas are evaluated at world-agent pairs, $K\varphi$ expresses that φ is known by the given agent at the given world. Intuitively, $K\varphi$ corresponds to the sentence "I know that φ." The F operator, applied to a formula φ, asserts that the formula φ holds of all agents who are friends with

[1] EFL is usually studied in conjunction with General Dynamic Dynamic Logic, a system of representing model updates. We will not be making use of GDDL in this chapter. Throughout this section I will be presenting the most recent version of this framework from [3], which has changed substantially since its original presentation in [4].

a. Formulas of the form $F\varphi$ can be approximated in natural language as "φ is true of all my friends." Similarly, the operator A expresses what is true of all agents in the model. We can gloss $A\varphi$ as "φ is true of everyone."

Definition(Syntax of Facebook Logic)

$\varphi := p \mid n \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K\varphi \mid F\varphi \mid A\varphi$

where the first kind of atom $p \in$ Prop is an indexical proposition and the second $n \in$ ANom is an indexical proposition asserting identity.

We also define the duals of our operators, in the usual manner of modal logic. $\langle K \rangle\varphi$ is defined as $\neg K \neg\varphi$, roughly "it is an epistemic possibility for me that φ." $\langle F \rangle$ is defined as $\neg F \neg$, meaning "I have a friend for whom φ holds."

Formulas of the language $\mathcal{L}_{\mathsf{EFLA}}$ are interpreted on models similar to the ones used in previous chapters with two main difference. First, a function g is included to handle the indexical nature of expressions. g is a function mapping each agent $a \in ANom$ to the agent $g(n) \in \mathcal{A}$ named by a. The set of agents $\mathcal{A}$ appears explicitly in the model for this reason. In previous models this set was implicit, with knowledge relations $\sim_a$ being the only objects of interest. Secondly, instead of a function E providing an access relation for each state w, we have a function F returning a *friendship* relation. The difference here is that the friendship relation is symmetric and irreflexive – your friends must be friends with you and you cannot be friends with yourself. This relation requires a much different interpretation.

We previously thought of connections to other agents as possession of a password to their database, or the ability to peek at their answer sheet. In epistemic friendship models, we have to interpret things differently. One way to interpret this language is by analogy with Facebook. In this setting, users can only access each others information after they have both confirmed that they are friends. Once this is established, they both have access to all the other's listed information. This is the kind of symmetric relationship depicted by EFL. Additionally, Facebook users are not allowed to add themselves as friends. The relationship here is irreflexive, as it is in EFL.

Definition An *epistemic friendship model* is a tuple:

$$< W, \mathcal{A}, (\sim_a)_{a \in A}, F, g, V >$$

where W is a set of states and A is a set of agents. $(\sim_a)_{a \in \mathcal{A}}$ is a family of equivalence relations $\sim_i$ for each $i \in \mathcal{A}$, representing each agent's ignorance between states. F is function $F : W \to (\mathcal{A} \times \mathcal{A})$ returning the a symmetric and irreflexive friendship relation for each state $s \in W$. g is an indexed set of agents g_n for each agent nominal $n \in \mathsf{ANom}$. V is a valuation function.

In addition to the operators defined above, we define hybrid-logic inspired operators $@a$ and $\downarrow n$. Intuitively, @ makes the evaluation of the inner formula depend on some fixed agent a rather than whatever agent is given. We can think of $@a\varphi$ as "φ as applied to a." If a proposition p stood for "I am poor," $@ap$ would assert that "Agent a is poor."

The $\downarrow n$ operator, by contrast, is designed to lock the *content* of a formula no matter who the given agent is. The $\downarrow n$ operator causes the inner formula n to actually refer to n even when evaluated at other agents m. With this operator, it is possible to refer back to the speaker after shifting focus to friends or other agents. For example, $\downarrow nFK\langle F\rangle n$ expresses the sentence "All my friends know they are friends with me," rather than "All my friends know they are friends with themselves."

Definition(Semantics for Epistemic Friendship Logic)

$$
\begin{array}{ll}
M, w, a \models n & \text{iff } a = n \text{ for } n \in \mathsf{ANom} \\
M, w, a \models p & \text{iff } (w, a) \in V(p), \text{ for } p \in \mathsf{Prop} \\
M, w, a \models K\varphi & \text{iff } M, w, a \models \varphi \text{ for every v } \in W \text{ such that } (w, v) \in \sim_a \\
M, w, a \models F\varphi & \text{iff } M, w, a \models \varphi \text{ for every b } \in A \text{ such that } (a, b) \in F(w) \\
M, w, a \models A\varphi & \text{iff } M, w, a \models \varphi \text{ for every b } \in A \\
M, w, a \models @n\varphi & \text{iff } M, w, n \models \varphi \\
M, w, a \models \downarrow n\varphi & \text{iff } [^n_a]M, w, a \models \varphi
\end{array}
$$

where $[^n_a]M$ is the result of changing M so that n now names a. In more precise terms, $[^n_a]M = \langle W, A, k, f, g[^n_a], V\rangle$ and $g[^n_a] = a$ if $m = n$ and g_m otherwise.

In the context of named agent models, i.e. models in which every agent in $\mathcal{A}$ has a corresponding name in $ANom$, formulas of the form $\downarrow n\varphi$ can be abbreviated as:

$$\downarrow n\varphi := \bigvee_{m \in \mathsf{ANom}} (m \wedge \varphi[^n_a])$$

where $\varphi[^n_a]$ is the result of replacing agent nominal m by n throughout φ.

6.2 Iterated Friendship Logic

We can iterated the modalities for knowledge and friendship to arrive at an iterated version of EFL. Formulas of the form $K^n\varphi$ will express, for a given agent a, the combined information of a and agents within n friendship edges of a.

Definition(Syntax of IFL)

$\varphi := p \mid n \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K^n\varphi \mid F^n\varphi \mid A\varphi$

where the first kind of atom $p \in$ Prop is an indexical proposition and the second $n \in$ ANom is an indexical proposition asserting identity. We define the notation "$\langle K^n\rangle\varphi$" to abbreviate $\neg K^n\neg\varphi$ and the notation "$\langle F^n$" to abbreviate $\neg F^n\neg\varphi$.

We additionally define the abbreviation $\langle F^n\rangle\varphi$ as:

$\langle F^0\rangle\varphi := \varphi$

$\langle F^{n+1}\rangle\varphi := \langle F^n\rangle\varphi \vee \langle F\rangle\langle F^n\rangle\varphi$

To define the semantics for these iterated knowledge modality, we first define the relations $\sim^n_a$ as before:

$(w,v) \in \sim_a^0$ iff $(w,v) \in \sim_a$

$(w,v) \in \sim_a^1$ iff $(w,v) \in \sim_a \cap \{\bigcap \sim_b | \varphi \in (a,b) \in F_w\}$.

$(w,v) \in \sim_a^2$ iff $(w,v) \in \sim_a^1 \cap \{\bigcap \sim_{b_2}^{\varphi} | \exists b_1$ such that $(a,b),(b_1,b_2) \in F(w)\}$

$(w,v) \in \sim_a^n$ iff $(w,v) \in \sim_a^{n-1} \cap \{\bigcap \sim_{b_n} | \exists b_1, ..., b_{n-1}$ such that $(a,b), ..., (b_{n-1}, b_n) \in F(w)\}$

Now we can easily define the semantics for IFL based on these supplementary relations:

Definition (Semantics for Iterated Friendship Logic)

$M,w,a \models n$	iff $a = n$ for $n \in \mathsf{ANom}$
$M,w,a \models p$	iff $(w,a) \in V(p)$, for $p \in \mathsf{Prop}$
$M,w,a \models K^n\varphi$	iff $M,w,a \models \varphi$ for every v $\in W$ such that $(w,v) \in \sim_a^n$
$M,w,a \models F\varphi$	iff $M,w,a \models \varphi$ for every b $\in A$ such that $(a,b) \in F^n(w)$
$M,w,a \models A\varphi$	iff $M,w,a \models \varphi$ for every b $\in A$
$M,w,a \models @n\varphi$	iff $M,w,n \models \varphi$
$M,w,a \models \downarrow n\varphi$	iff $[^n_a]M,w,a \models \varphi$

6.3 Axiomatization of IFL

(MP) From φ and $\varphi \to \psi$, derive ψ.

(N) From φ derive $K_a\varphi$.

(Prop) All validities of propositional logic

(KF) $K^n(\varphi \to \psi) \to (K^n\varphi \to K^n\varphi)$

(TF) $K^n\varphi \to \varphi$

(4F) $K^n\varphi \to K^nK^n\varphi$

(5F) $\neg K^n\varphi \to K^n\neg K^n\varphi$

(F1) $K^n\varphi \to K^{n+1}\varphi$

(F2) $(\langle F\rangle \wedge K^n)\varphi \to K^{n+1}\varphi$

(F3) $\langle F\rangle b \to K\langle F\rangle b$

(F4) $(\bigwedge_{b\in\mathcal{A}} \neg\langle F\rangle b) \to (K^n\varphi \to K\varphi)$

(F5) $\varphi \to F\langle F\rangle\varphi$

(F6) $\bigwedge\limits_{a\in\mathcal{A}} \neg(a \wedge \langle F\rangle a)$

(F7) $\bigwedge\limits_{b\in\mathcal{A}} \langle F\rangle b \to \langle A\rangle(c \wedge \langle F\rangle b)) \to (K\varphi \to \langle A\rangle(c \wedge K\varphi))$

Conjecture: IFL is sound and complete for epistemic friendship models.

For every axiom of the complete logic IAL, we have a translation into the syntax of EFL logic. This makes IFL a good candidate for completeness, but this is still an open question.

Chapter 7

Controlling Information Access

In the previous chapters, we developed several logics for reasoning about the information agents had access to. These languages were able to represent situations involving access, and in one case even represent the relevant information change. There is, however, the further component of access involving *choice*. The chapters thus far have not included the agency of the epistemic agents represented. In this chapter, I introduce a logic that connects agents' choices and information access. Agents can choose to communicate facts to those they are connected to in the network. We then develop a logic for reasoning about what agents could know with the help of others, assuming that individual agents can only receive information from their immediate network neighbors. The language will be particularly focused on capturing inferences regarding the combined ability of agents to supply or deny information. Inferences of this sort are a common feature of spy and mafia films, where characters must reason about the flow of vital information in communication networks. Sentences like: "Bob could learn φ without Aaron" or "Cathy can make sure Darren and Eric never know" will be formalized in the logic introduced.

This logic, which we will call the Logic of Communication Graphs with Coalitions (LCGC), is an expansion of the work in [1]. This paper introduced an epistemic logic (LCG) for reasoning about what agents could come to know given a particular communication graph. Formulas of LCG captured sentences such as "If Bob knows φ, then it is possible for Aaron to come to know φ." The logic introduced here, LCGC, expands the syntax with two coalition modalities and can additionally capture statements about the informative power of groups.

Defintion (Syntax for LCGC)

$\varphi := p \mid \neg\varphi \mid \varphi \vee \psi \mid K_i\varphi \mid \Diamond\varphi \mid [C]\varphi$

where $p \in \mathsf{At}$, a set of proposition variables, and C is a subset of the set of agents $\mathcal{A}$. The formulas of the form $\Diamond\varphi$ and $K_i\varphi$ are interpreted in a temporal manner – glossed as "It is possible for φ to become true" and "Agent i knows that φ will always be true." We also define $\Box\varphi = \neg\Diamond\neg\varphi$. Formulas of the form

$[C]\varphi$ can be read as "Group C is effective for φ," which intuitively means that the agents in C are capable of independently forcing φ to be true.

To begin explaining this logic, we first introduce the notion of a communication graph:

Definition A *communication graph model* is a tuple

$$\langle \mathcal{G}, \mathsf{At}, \overrightarrow{v} \rangle$$

where:

1. $\mathcal{G} = \langle \mathcal{A}, E \rangle$ is a graph consisting of a finite set of agents $\mathcal{A}$ and a set of edges $E \subseteq \{\mathcal{A} \times \mathcal{A} - \{(i,i) \in \mathcal{A}\}\}$, i.e. the set of all possible irreflexive edges between agents in $\mathcal{A}$. Intuitively, each pair $(i,j) \in E$ means that agent i can receive messages from agent j.

2. At is a finite set of proposition variables.

3. $\overrightarrow{v} = (v_1, v_2, ..., v_n)$ is a vector of partial boolean valuations. This represents the initial information of all the agents in $\mathcal{A}$. A partial valuation will be defined as a partial function $v_i : \mathsf{At} \to \{0,1\}$. The vector is required to be consistent in the sense that, for each $p \in \mathsf{At}, v_i(p) = v_j(p)$ for all $i, j \leq n$.

The above definitions capture the initial information state and the arrangement of agents in a communication graph. We will assume that all the information captured in a communication graph model is common knowledge, except for the values within each partial valuation v_i. The set of propositions p in the domain of each v_i, however, *is* common knowledge. So while other agents do not know the contents of others valuation, they know which issues the other agents possess knowledge about. We now turn to how information change is represented:

Definition A tuple (i, j, φ), where $\varphi \in \mathcal{L}_{DNF}(At)$ and $(i,j) \in E$, is called a *communication event*. Then for a given communication graph $\mathcal{G}$, $\Sigma_{\mathcal{G}} = \{(i,j,\varphi) | \varphi \in \mathcal{L}_{DNF}, (i,j) \in E\}$

Throughout this section we are assuming that communication events are fully known to participants, and "fair-game" private to other agents. That is, agents not involved in the event will consider it possible that the event took place. A possibility is represented as a particular sequence of communication events – a "history." We also introduce states, which represent the facts about the world. Formulas of the language are evaluated at pairs (w, H), where w is a state and H is a history.

Definition A *history* is a finite sequence of events from Σ_G. The empty history is represented by ϵ.

Definition Let W be the set of boolean valuations on At. Then an element $w \in W$ is called a *state*.

Not all histories are legal. An agent must know φ before being able to inform another agent of this fact. For an event (i, j, φ) to be a legal addition to a history H it must hold that $(i, j) \in E_G$ and that after the events in H, j actually knows φ. The issue of specifying which histories are legal will be solved by introducing a propositional symbol L, which is satisfied only when the pair (w, H) is legal. We may also write $L(w, H)$ in place of $M, w, H \models L$. As noted previously, whether a history is legal depends in part on what knowledge agents have gained previously. Knowledge in turn requires quantification over legal histories, so we will need to use mutual recursion. To capture a particular agent's perspective on events, we define *i-equivalence*.

Definition Let w be a state and H a finite history. Let λ_i for all $i \in \mathcal{A}$ be defined as follows, were ε denotes the empty string:

$$\lambda_i((j,k,\varphi)) = \begin{cases} (j,k,\varphi) & \text{if } i = j, k \\ \varepsilon & \text{otherwise} \end{cases}$$

Let $\lambda_i(H)$ be defined as the history resulting from applying λ_i to each communication event $e \in H$. Then the *i-equivalence relation* $\sim_i$ is defined as follows: $(w, H) \sim_i (v, H')$ iff $\lambda_i(H) = \lambda_i(H')$.

Definition (Semantics for LCGC)

$M, w, \epsilon \models L$

$M, w, H; (i, j, \varphi) \models L$ iff $M, w, H \models L, (i, j) \in E,$ and $M, w, H \models K_j\varphi$

$M, w, H \models p$ iff $w(p) = 1,$ where $p \in \mathsf{At}$

$M, w, H \models \neg\varphi$ iff $M, w, H \not\models \varphi$

$M, w, H \models \varphi \wedge \psi$ iff $M, w, H \models \varphi$ and $M, w, H \models \psi$

$M, w, H \models \Diamond\varphi$ iff $\exists H', H \preceq H', L(w, H'),$ and $M, w, H' \models \varphi$

$M, w, H \models K_i\varphi$ iff $\forall(v, H')$ if $(w, H) \sim_i (v, H')$,
and $L(v, H'),$ then $M, v, H' \models \varphi$

$M, w, H \models [C]\varphi$ iff $\exists\sigma_1 ... \sigma_n \in \Sigma_{w,H}(C)$ such that
$\forall H'$ such that $H \preceq H', L(w, H^+)$,
and $\{H'\} \cap \Sigma_C = \{\sigma_1, ..., \sigma_n\}$,
then $M, w, H' \models \varphi$

where $\Sigma_{w,H}(C)$ is the set of communication events (i, j, φ) such that $j \in C$ and $L(w, H; (i, j, \varphi))$.

7.1 Decidability of Model Checking in LCGC

To show that model checking in LCGC is decidable, we need to show that model checking a formula involves only a finite number of histories and that the checking of each history terminates in finitely many steps. [1]

Lemma (i): For a communication graph $\mathcal{G}$ with k propositions and n agents, there are at most $n \times n \times (2^{2^k})$ events in $\Sigma_{\mathcal{G}}$. So there are only finitely many communication events possible.

Proof: We assume all communicated formulas are in disjunctive normal form. Then the above limit is the result of calculating all possible pairings of agents and possible formulas.

Lemma (ii): We now show that every formula satisfiable by a pair (w, H) is also satisfiable in a history $(w, c(H))$ in which no communication is repeated after the first instance. In particular, (w, H) is legal iff $(w, c(H))$ is legal.

Proof: It suffices to show that these conditions hold between an arbitrary history H and any H' obtained from H by removing one copy of a communication event. Then for some communication event $e = (i, j, \theta)$ we can represent $H = H_1 e H_2 e H_3$ and $H' = H_1 e H_2 H_3$.

Take an arbitrary formula ψ. If ψ is a propositional formula or boolean combination then clearly (w, H) and (w, H') either both satisfy ψ or not, since w is unchanged.

Then consider the case of $\psi = \Diamond\phi$. If $w, H \models \psi$ then there exists some history H_4 such that $w, H; H_4 \models \phi$. By induction hypothesis $w, H', H_4 \models \phi$ and so $w, H' \models \Diamond\psi$. The same argument works for the converse.

Now suppose $\psi = K_r\phi$ and $w, H \models \psi$. Although $K_r\phi$ is technically defined via satisfaction of ϕ within the scope of legal, r-equivalent world-history pairs, we can avoid using L by showing the condition holds for *all* r-equivalent world-history pairs.

First take $r \neq i, j$. Then $w, H' \models \theta$ since H, H' are r-equivalent.Take $r = i$. Then $w, H \models K_i\phi$ iff for all v, H'' such that $(v, H'') \sim_i (w, H), v, H'' \models \phi$. But since e has already occurred in both H, H'', the possible v are the same. Since H and H'' are i-equivalent they have the same communication events, and in particular they both have two e instances. So we can eliminate the second e instance to form an H''' such that $v, H'' \models \phi$ iff $w, H''' \models \phi$ by induction hypothesis. So $w, H' \models K_i\phi$ as well. The same argument works for $r = j$.

Next, the fact that compression does not affect satisfaction of L, the proposition symbol for legality, follows from the previous arguments because L satisfaction is recursively defined in terms of legal formulas and knowledge formulas.

[1]This proof uses much of the proof from decidability for LCG given in [1]

Finally, the condition holds automatically for formulas of the form $[C]\phi$, since truth for $[C]\phi$ is defined in terms of truth for L and $\Diamond$ formulas.

Note that proving this last lemma implies that formula checking will always terminate, since an algorithm has to check at most:

$$\frac{|\Sigma_{\mathcal{G}}|!}{1!} + \frac{|\Sigma_{\mathcal{G}}|!}{2!} + ... + |\Sigma_{\mathcal{G}}| + 1$$

histories, which is finite. So LCGC is decidable.

Chapter 8

Conclusions and Further Work

This thesis has explored several ways of representing and reasoning about agents' access to information. Several themes have emerged:

1. The information accessible to an agent can, to a large extent, be treated as another variety of knowledge. This includes higher-order formulations ("I know that he could find out.") and conditional statements ("If Allen can find out so can I."). As discussed further on, there are additional features of knowledge, such as introspection, that could be added in future treatments of access.

2. Knowledge about social connections has, in itself, epistemic value. In the situations studied, knowledge of communication ties allowed agents to reason about what others could find out or disclose.

A clear area for future work is in creating logics which pay more attention to the interaction involved in information access. The kind of access studied in this thesis was a one-way transfer of information from a source to a learner, which is not how access typically proceeds. Access more commonly involves questions, permission or other kinds of intermediate steps. These in-between steps have epistemic consequences. At the very least, asking questions betrays my ignorance to others. At the extreme end, asking questions could inform someone of my bad character (e.g."Do you carry a lot of cash?") and actually make the knowledge I seek unavailable. While questions and other features of communication have been studied independently in several places [13], it is necessary to incorporate these approaches to get a nuanced picture of what information an agent can *actually* access.

A particularly important improvement would be to incorporate the introspective nature of access. In this thesis individuals have access to many facts, but in general do not know what these facts might be. Information access, as distinguished from other forms of communication, implies that the agent knows an answer will be forthcoming. This is especially true for contexts in which access plays a constant and crucial role. Reflecting on the situations in which we typically rely on information access – searching for directions, looking up a phone number – we know that the information is there to be had. This condition on access would be formalized along the lines of:

$$K'_a\varphi \Rightarrow K_a(K'_a\varphi \vee K'_a\neg\varphi).$$

These ideas are closely related to the philosophical disputes over internal versus external justifications of knowledge. The above condition might be seen as a kind of "internalist" condition on having information access, while this thesis represents a more externalist view.

Chapter 9

Bibliography

[1] E. Pacuit and R. Parikh. Reasoning about Communication Graphs. In *Interactive Logic, Texts in Logics and Games*, Volume 1, Johan van Benthem, Benedikt Loewe, Dov Gabbay (eds.), 2006.

[2] P. Girard, J. Seligman, F. Liu. General Dynamic Dynamic Logic. *Advances in Modal Logic* 2012: 239-260, 2011.

[3] J. Seligman, F. Liu, and P. Girard. Facebook and Epistemic Logic of Friendship in Burkhard C. Schipper ed, Theoretical Aspects of Rationality and Knowledge Proceedings of the 14th Conference Chennai, India, pp.229-238.

[4] J. Seligman, F. Liu, and P. Girard. "Logic in the Community" in M. Banerjee and A. Seth, editors, *Proceedings of the 4th Indian Conference on Logic and its Applications*,Volume 6521 of LNCS, pp. 178-188. Springer, 2011.

[5] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge.* The MIT Press, 1995. Cambridge, Massachusetts. London, England

[6] A. Baltag and S. Smets. "The Logic of Conditional Doxastic Actions". In R. van Rooij and K. Apt (eds.), *Texts in Logic and Games, Special Issue on New Perspectives on Games and Interaction*, Volume 4, pp. 9–31, Amsterdam University Press, 2008.

[7] J. Bruggeman. *Social Networks: An Introduction.* Routledge, New York, 2008.

[8] R. Fagin, J. Halpern, and M. Vardi What Can Machines Know? On the Properties of Knowledge in Distributed Systems *Journal of the Association for Computing Machinery.* Vol 39. No 2, April 1992 pp 328–376.

[9] A. Clark and D. Chalmers The Extended Mind *Analysis* 58(1):7-19. 1998.

[10] M. Pauly "A Modal Logic for Coalitional Power in Games" *J. Logic Computat.,* Vol. 12 No. 1, pp. 149-166, Oxford University Press, 2002.

[11] A. Baltag, L. Moss and S. Solecki, The Logic of Public Announcements, Common Knowledge, and Private Suspicions. *Technical Report* SEN-R9922, CWI, Amsterdam (1999)

[12] P. Blackburn, M. de Rijke, Y. Venema. *Modal Logic*. Volume 53 of Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2002

[13] J. van Benthem, S. Minica. Toward a Dynamic Logic of Questions. J. Philosophical Logic 41(4): 633-669 (2012)