Collaborative Knowability

MSc Thesis *(Afstudeerscriptie)*

written by

Klarise Marais

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

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
October 30, 2025	Dr Malvin Gattinger (chair)
	Dr Aybüke Özgün (supervisor)
	Dr Alexandru Baltag (supervisor)
	Dr Maria Aloni
	Dr Marianna Girlando

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

Formal epistemologists design and study formalisms that represent concepts or processes from mainstream epistemology. The purpose of this is either to contribute to mainstream epistemology literature by showing relationships between or implications of theories and concepts, or to be used in implementations. In this thesis, we attempt to create a formalism that has the potential to serve both purposes. On the philosophical side, we attempt to use our formalism to show consequences of theories on group knowledge, and on the technological side, we aim to make our formalism suited to teach artificial systems about learning/collaboration. We create our formalism by adding elements relevant to collaboration to existing dynamic epistemic logic literature, and using this setting to formally define a notion of group knowledge/knowability called potential collaborative knowledge. We then use this notion to formally show how it is possible for a group to know/learn more than the sum of what individual members know/learn, benefitting the philosophical literature by showing implications of concepts/theories, and showing how potential collaborative knowledge is a generalisation of distributive knowledge, a commonly studied type of group knowledge. We then work towards axiomatising our formalism by proposing an axiomatisation and proving soundness. We end with a discussion of some possible future directions of research.

Acknowledgements

I would like to start by thanking my supervisors, Aybüke Özgün and Alexandru Baltag. To Aybüke, for her thoroughness and reliability, her careful and clear explanations, and her emotional support and skilful management of sensitive situations throughout my thesis writing. To Alexandru, for all the time and effort he put into developing many of the central ideas present in this thesis, and for his presence and enthusiasm for the project. I would also like to extend a special thanks to both of my supervisors for all the additional hours they put in in the last few weeks of my thesis. I would not have been able to complete this thesis without them.

Next, I would like to thank the members of the ILLC for the skills and knowledge they imparted to me that made writing this thesis possible. Here, I would like to highlight my academic mentor, Dick de Jongh, for his guidance and mentorship during my studies.

I would also like to thank my defence committee for taking the time to read my thesis and for formulating thoughtful questions.

This thesis would also not have been possible without the financial support of the EW Beth foundation and my family, particularly my grandparents, Oupa Pierrie and Ouma Linda, and my mom. Here I would also like to thank Nick Bezhanishvili for his assistance in getting my funding extended.

Lastly, I would like to thank everyone in my life who supported me emotionally during my thesis and helped to keep me (mostly) sane. To my family, specifically my mom, brother, dad, and step-dad, for our calls and weekly online board game sessions. To all my friends who, in addition to their company and general support, helped motivate me to work on my thesis by meeting with me to work together: Simon, Fangjing, Swapnil, Jan-Luc, Billie and Stefano. To my roommates, Lorenz and Louise, for our chats and for helping to make our apartment a pleasant and comfortable environment to be in. To my other friends who also helped keep me afloat: Tracy, Sindi, Lina, Zoleka, Dimakatso and Amanda. A special thank you to Pedro here, for regularly cooking for me and providing a warm place to come back to during the busy periods of my thesis, for holding space for all of my emotional ups and downs, and for encouraging me to be gentle on myself and to take my struggles seriously.

Contents

Chapter 1

Introduction

In formal epistemology, the goal is usually to design and study formalisms that represent concepts or processes from the mainstream epistemology literature. For example, in Bayesian epistemology, the concept of degrees of belief is studied, along with the way these degrees of belief change when presented with new information or evidence. As another example, in dynamic epistemic logic, learning processes are studied by considering how actions can affect the knowledge of agents. There are two types of purposes that these formalisms can serve. One is to contribute to mainstream epistemology, and the other is to aid in the development of new technologies. The contribution to mainstream epistemology usually comes in the form of helping to clarify concepts, relationships between alternative proposals, or to show consequences of concepts and theories [HB20].

In this thesis, we develop a notion of group knowledge or knowability, which we call potential collaborative knowledge, and motivate this notion both from a philosophical and technological perspective. Informally, we say a group has potential collaborative knowledge of a statement when the members can collaborate to each individually learn it. For the philosophical contribution, we consider the literature on group agency and group knowledge. The idea that groups can be viewed as agents in their own right has become more prevalent in social epistemology [Bro24; Lac20; LP11; Rol08]. This view allows for non-metaphorical speech on group knowledge and group knowability. Once groups are accepted as agents, it also becomes possible to compare the knowledge of the group to the knowledge of the individual members. Two opposing perspectives on this comparison are summativism and non-summativism [Bro24; Lac20; HR23]. Briefly, summativism is the view that group attitudes are a function of individual attitudes, and non-summativism views group attitudes as not reducible to individual attitudes, stating that the group can have knowledge or beliefs not held by any individual members.

Our philosophical contribution is to show that, under any perspective on actual group knowledge which takes groups as agents, the group has the ability to learn more than any individual member can learn alone. If we additionally view our account of potential collaborative knowledge as a notion of group knowledge, then we also contribute to the philosophical literature by presenting a form of non-summative group knowledge, and thereby showing which features an account of group knowledge could have to be non-summative. An implication of this contribution is that there are aspects of knowledge/learning that are irreducibly social, and therefore it is important to consider social elements of knowledge and knowledge formation in epistemology and the study of learning.

The implementation-based motivation is to create a formal framework that is well-suited to modelling scenarios involving collaborative learning or discovery. Our formalism builds on existing formalisms by providing ways to express which agents are executing an action, and by generalising a well-studied notion of group knowledge, called distributive knowledge. This allows our framework to capture aspects of collaboration and knowledge formation that were not captured by the existing frameworks that we build on.

The existing technical literature that we build on to produce our formalism comes from the field of dynamic epistemic logic. We use Kripke semantics to model the knowledge states of agents, following Jaakko Hintikka [Hin62], and we use data-exchange event models, developed by Alexandru Baltag and Sonja Smets [BS21; BS24], to model the actions that individuals and subgroups can take in collaboration processes. We build on these event models by assigning a specific set of "active agents" to a few instances of these models, and by using them in our formal definition of potential collaborative knowledge.

Our main contributions can then be summarised as consisting of three parts. The first is building on existing dynamic epistemic logic to create a formal setting capable of representing collaboration processes and the knowledge they produce. The second is using this setting to demonstrate how groups can collaborate so that each member is able to learn more than they could learn individually, using an example. This could be interpreted as presenting a form of non-summative group knowledge. The third is studying validities of our system and proving soundness of the static component of our system with respect to a proposed axiomatisation.

This thesis is structured as follows: In Chapter 2, we explain the motivation for our formalism in more detail and provide the necessary technical background for this thesis. We explain the motivation in five steps. First, we explain the purpose of modelling in general. Second, we provide necessary philosophical background on group agency and summativism vs. non-summativism. Third, we explain how, if potential collaborative knowledge is viewed as a notion of group knowledge, it is a non-summative one. Fourth, we explain how we can view potential collaborative knowledge as a notion of group knowability to demonstrate the significance of collaboration. Fifth and last, we explain how our formalism could have uses in artificial intelligence. In the technical background, we introduce multi-agent epistemic logic, data-exchange logic, and generated submodels.

In Chapter 3, we design the formal setting needed to model potential collaborative knowledge, formally define the notion in this setting, formalise an example that both demonstrates the functioning of the operator and supports our philosophical points, and study some properties of potential collaborative knowledge. For the formal setting, we require specific types of event models. We first define and explain these, as well as study some of their properties. We then create a syntax for our particular types of events and define meta-syntactic notions for them, including active subgroups and capabilities of subgroups. To complete the design of our formal system, we define the syntax and semantics for our main language $\mathcal{L}_{\mathcal{K}_H^G}$, and define potential collaborative knowledge as an abbreviation in the language. The example we formalise is of a collaborative effort to observe a planet, and the main property we show is that potential collaborative knowledge is a generalisation of distributive knowledge.

In Chapter 4, we prove soundness for our conjectured axiomatisation with respect to the semantics we define over $\mathcal{L}_{\mathcal{K}_H^G}^{static}$, which is our main language $\mathcal{L}_{\mathcal{K}_H^G}$, without the dynamic operator. Some validities that form part of our conjectured axiomatisation are all the validities of S4 for our main operator, $\mathcal{K}_H^G$, monotonicity validities for our main operator, $\mathcal{K}_H^G$, and an axiom schema capturing the essence of the functioning of our potential collaborative distributive knowledge operator, $\mathcal{K}_H^G$.

Finally, in Chapter 5, we summarise our work and discuss implications and future directions.

Chapter 2

Motivation and Background

In this chapter, we motivate the technical work done in this thesis and provide the relevant technical background for the rest of the thesis. Our motivations are both philosophical and implementation-based, but we go into more detail on the philosophical motivations. In the technical background, we give brief introductions to multi-agent epistemic logic, data-exchange logic, and generated submodels.

2.1 Motivation

In this thesis, we introduce a notion of group knowledge or knowability that we call potential collaborative group knowledge. Our goal with this is both to contribute to the philosophical literature, and to introduce a formalism that is suited for applications in the development of artificial intelligence. In this section, we explain our motivation and contribution in more detail.

2.1.1 The purpose of formalism

Formal epistemologists create formalisms to represent concepts or processes encountered in epistemology and study these formalisms. For example, Bayesian epistemologists study degrees of belief using probabilities, logicians working in epistemic logic study models that represent the belief and knowledge states of agents, and logicians working in dynamic epistemic logic study models that represent how belief or knowledge states change when events occur. The questions I address in this section are: Why create and study these formalisms, and why our formalism in particular?

There are two types of purposes these formalisms can serve: a philosophical purpose and an implementation purpose. The philosophical purpose is to clarify philosophical concepts, to show the consequences of these concepts, and to shed light on the relationships between alternative proposals. Sven Ove Hansson highlights four possible benefits of formalism in philosophy [Han00]. The first is demonstrating the interdefinability of concepts. An example he gives of this was the interdefinability of "must", "may", and "forbidden" in deontic logic. The second is to make implicit assumptions explicit. The third is that formal structures can help us study intricate concepts or ideas that are difficult to handle and follow informally. The example Hansson gives here is of Tarski's semantical analysis of the concept of truth. The fourth is that formalisation can help us see more consequences of our theories and concepts. The example Hansson gives here is on non-monotonic reasoning. The formalisation of non-monotonic inference has enabled researchers to link non-monotonic reasoning to theories of rational choice.

In this thesis, we focus on the fourth benefit, namely, demonstrating more consequences of theories and concepts, including relationships between theories. One example of research in formal epistemology benefiting the philosophical literature in this way was Baltag, Nick Bezhanishvili, Aybüke Özgün and Smets' analysis of theories of knowledge using evidence models [Bal+22]. In this paper, the authors fulfil the purpose of showing relationships between concepts by formally studying the relationships between knowledge, belief, defeasibility theories of knowledge, and Stalnaker's view on subjective certainty. In addition, they fulfil the purpose of helping to study intricate concepts by showing a subtle distinction between true and misleading, and true and non-misleading evidence.

This thesis aims to benefit the philosophical literature by providing a formal account of collaboration, demonstrating the significance and importance of collaboration and possibly contributing to the summativism vs. non-summativism debate by defining a type of non-summative group knowledge. Briefly, summativists claim that group knowledge can be reduced to individual knowledge, and non-summativists deny this, claiming that there are cases where groups can know propositions that no individual members know [Bro24; Lac20; HR23]. We demonstrate how, under any account of group knowledge, groups have the ability to collaborate to learn more than all the logical consequences of what each member can learn individually. This is the fourth type of benefit described by Hansson, since we show a consequence of both summative and non-summative theories on group knowledge. Defining a type of non-summative group knowledge also contributes to the philosophical literature by showing an example of features an account of group knowledge can have to be non-summative.

The implementation purpose is to be used in the development of programs and artificial intelligence systems. Previously, dynamic epistemic logic has been used to enable artificial systems to have more sophisticated decision-making processes. Thanks to dynamic epistemic logic, we now have a way to teach artificial systems to reason about the epistemic states of other agents and how these states change after actions. For example, Lasse Dissing and Thomas Bolander have implemented DEL-based reasoning in a humanoid robot to enable it to perform tasks associated with having a theory of mind, such as identifying false beliefs in others [HB20]. We argue that our formalism, which builds on previous work in dynamic epistemic logic, can have similar uses.

2.1.2 Groups as entities above and beyond their members

Before we go into more detail on the contribution of this thesis, we provide an overview of different positions on groups and group knowledge. In philosophy, groups are usually defined as more than a mere collection of individuals. For a collection of individuals to be a group, they usually need to be united by some common purpose or goal [Lac20; LP11]. For example, a chess club is a group because they have the shared purpose of playing chess and potentially winning chess games against other chess clubs. One idea in the philosophical literature is that groups are agents in their own right [Bro24; Lac20; LP11; Rol08]. This view can be called realism about group agents. Within the realist literature, there is still a debate about whether speech about group agents is readily reducible to speech about the individual members [LP11].

The alternative to realism about group agents is eliminativism. Elimativism aims to make sense of everyday speech that treats groups as agents, such as "the government claims that crime is decreasing". One variation of eliminativism suggests that group agency speech is purely metaphorical, and another variation suggests that speech treating groups as agents is always misconceived. Then, the statement that the government made a claim is either just metaphorical talk attempting to convey information about what the individuals involved have done or claimed, or a fundamentally misconceived statement. Elimativism about group agency was popular among utilitarians, with Jeremy Bentham as one of the founders of the idea [Ben70].

Christian List and Philip Pettit are defenders of a realist view on group agency [LP11]. They argue that agents are entities satisfying three criteria. The first is that they have representational states depicting the environment. The second is that they have motivational states for how they desire the environment to be. The third and last is that they can process their representational and motivational states in a way that enables them to take action to make their environment better match their motivational states. List and Pettit argue that when collections of individuals have a shared purpose or goal, they have motivational states, and since the individual members can observe and interact with the environment, the other criteria are satisfied too. They therefore conclude that groups can be viewed as agents.

Once you view groups as agents in their own right, it becomes possible to speak non-metaphorically of the knowledge and beliefs groups have as agents, and to compare this to the knowledge and beliefs of the individual members. There is a large debate in the literature on group epistemology on this comparison: that is, on the relationship between the knowledge and beliefs of the group and the knowledge and beliefs of individual members [OGG24]. Two opposing positions in this debate are summativism and non-summativism. Summativism is the view that the attitudes (such as knowledge and belief) of groups can be reduced to some function of the attitudes of the individual members. One example of a summativist view is that the group holds a belief if and only if some or all members of the group hold that belief [Qui75]. Non-summativism is the view that group attitudes are not reducible to the attitudes of individual members, and that the group can have knowledge or beliefs that none of its members has. One example of a non-summative account of belief, founded by Margaret Gilbert, is the joint acceptance account, which states that a group believes a proposition if and only if all the members come to jointly accept this proposition [Gil87]. This is non-summative because, for various reasons (such as practical constraints and compromise), group members may come to jointly accept something none of them accepts individually. An example of a non-summative account about all group attitudes is functionalism, which states that group attitudes are defined by their functional role (i.e. their actual and potential causal relation to inputs, outputs, and other mental states). Some defenders of this account are List and Pettit, Jessica Brown, and Alexander Bird [Bir14; Bro24; LP11]. An example of how a group can have knowledge that none of its members has under this account is of published results in the scientific community that are not known by any living member of the community. Bird argues that published results play the functional role of being used in investigations and decision making when the topic comes up, and therefore fulfil the functional role of knowledge [Bir14].

With this background in mind, we are now in a position to clarify our contribution. Before doing so, it is important to clarify that, since this thesis is about group knowledge and group knowability, and we intend these terms to be taken literally, we assume a realist understanding of group agents.

2.1.3 Potential collaborative knowledge as non-summative group knowledge

The idea behind the notion of group knowledge/knowability that we present in this thesis, potential collaborative knowledge, is that the group has potential collaborative knowledge of a proposition if and only if the members can come to know that proposition by executing a series of allowed sharing and/or observation actions. Here, a sharing action is an action in which one or more agents share their information with other agents. In this thesis, we consider only cases in which agents share all that they know with other agents (rather than sharing selected information). An observation action is an action in which one or more agents learn a fact about the world, possibly using some knowledge they have. We also specify for each case which sharing and observation actions are allowed in that case.

Our notion can be seen as an extension of distributive knowledge, which is another notion of group knowledge studied in formal epistemology. Informally, a statement is distributive knowledge if the members of the group can learn it by sharing all their information with each other. Equivalently, a statement is distributive knowledge for a group if it is a logical consequence of the combined individual knowledge of the members. If we allow only and all, sharing actions for a given group, then potential collaborative knowledge and distributive knowledge become equivalent.

In the philosophical literature on group knowledge, several examples meant to support a non-summative account of group knowledge can be explained as cases of distributive knowledge. One example of this is Avram Hiller and Wolfe Randall using the example of a US Navy ship, the USS Palau, in which no member knows the location of the ship, but the crew as a whole can navigate the ship [HR23]. Hiller and Randall take the location of the ship as a form of non-summative group knowledge. The location of the ship here can be seen as a form of distributive knowledge as well, since it is a logical consequence of all the combined knowledge of the individual crew members.

However, typically, distributive knowledge is viewed as a summative form of group knowledge by formal epistemologists [Bod14], despite technically allowing the group to know statements that are not known by any individual members. There are two different ways we can make sense of this. The first is to observe that the core idea behind non-summativism is that the knowledge the group has is not a function of the knowledge of the individuals. In other words, what the group knows as an agent goes beyond what can be obtained just by looking at the individual members' knowledge. However, distributive knowledge can be obtained from the knowledge of individual members! Namely, by pooling all of it together and closing under logical consequence. The second way to make sense of formal epistemologists viewing distributive knowledge as summative is to observe that, generally in formal epistemology, some idealising assumptions are made about knowledge. Specifically, knowledge is taken to be closed under logical entailment [Hol16]. Therefore, in formal epistemology, agents also know the logical consequences of all their knowledge. The same would then apply to group agents. Hence, if you take one member of a group knowing a statement as sufficient for the whole group knowing it, most formal epistemologists would naturally close under logical entailment, yielding distributive knowledge.

Now, the notion we propose, potential collaborative group knowledge, when viewed as a notion of group knowledge, is non-summative in a more unambiguous way. Namely, that statements can be potential group knowledge even when they are not entailed by any combination of individual members' knowledge. We show this formally in the thesis, but to demonstrate the intuition behind this, we give two examples. The first is of a wildlife research team, in which one member knows which calls each type of bird in the area makes, another knows where different species of birds are likely to be found, and the last member knows how to use sound-amplifying equipment. Now, if the forest were very large and the birds were high enough that it is difficult to hear their calls, all three members are needed to learn whether a particular bird species of interest lives in the forest. No member alone would be able to find out whether the bird species lives in the forest. Therefore, assuming the bird species does indeed live in the forest, this fact is potential collaborative knowledge for the group, but it is not potential collaborative knowledge for any individual (i.e. no individual can perform actions alone to learn it).

The second example is of a group of two astronomy researchers collaborating to discover a planet. One researcher is able to learn how to use the telescope, since she understands the language of the manual, and the other member is able to look through the telescope to observe the planet, since he has good eyesight. The existence of the planet is therefore potential collaborative knowledge. However, no individual member can discover the planet alone. We formalise the second example in this thesis.

The notion of group knowledge/knowability we present, therefore, cannot be reduced to individual member knowledge/knowability. Our examples also demonstrate how, in the spirit of non-summativism, groups can be viewed as entities above and beyond their individual members. Therefore, our account can be seen to contribute to the philosophical literature on group knowledge by showing examples of features an account of group knowledge can have that would lead to it being a non-summative account. This is an example of the fourth type of benefit explained in Subsection 2.1.1 (showing consequences of theories and concepts), because it shows that a consequence of taking group knowledge to be what can be discovered via observation and sharing is that group knowledge is non-summative.

2.1.4 Demonstrating the importance of collaboration

In the last section, we argued that our account can be viewed as a form of non-summative group knowledge. However, it is possible to object here that the notion we introduce in this thesis is one of group knowability rather than group knowledge. The objection would be that we cannot say that a group knows a proposition just because individual members can learn it via communication and observations. Instead, we can say the group is able to learn that proposition. Then, since the summativism vs. non-summativism debate is about the nature of group *knowledge*, summativists can still agree that, in some instances, collaboration is necessary to produce knowledge [Fag12]. The conclusion would therefore be that we cannot link our notion to the summativism vs. non-summativism debate.

Our response is that, since the notion we introduce is an extension of distributive knowledge, which is often viewed as a form of group knowledge rather than group knowability, it is reasonable to view our notion as of the same type. The only difference between our notion and distributive knowledge is that agents are allowed to take more types of actions (observation and communication actions, rather than just communication actions) to have the knowledge actualised in the individual members. It can therefore be viewed as a matter of opinion if this change is enough to make our generalisation of distributive knowledge into a notion of group knowability, or if it can still be seen as a notion of group knowledge.

However, we claim that even in the case where potential collaborative knowledge is viewed as a notion of group knowability rather than group knowledge, our thesis still contributes to social epistemology. To clarify, when we view potential collaborative knowledge as a notion of group knowability, we understand it informally as follows:

> *A group has potential collaborative knowledge of a statement when the group can perform observation and sharing actions to reach actual group knowledge.*

Where a statement is actual group knowledge when all members know it. The difference between this interpretation and the interpretation in the previous subsection is that the previous subsection took potential collaborative knowledge as the notion of actual group knowledge being defined, rather than seeing it as defining the conditions to reach actual group knowledge.

We can use this interpretation of potential collaborative knowledge to show that under any realist account of group knowledge (i.e. any account that accepts the literal use of the term "group knowledge"), the group is able to collaborate to learn more than the logical consequences of what the members can learn alone. To explain this, note that non-summativism already entails that groups can learn more than individual members. To see why, consider the case where the group and its members are unable to take actions to learn more. Then what they are able to learn is exactly everything they already know. If, in some situations, the group knows more than what the individual members know, then the group can come to know more than the individual members in this case. This means that only summativists can deny that the group can learn more than individual members can.

We therefore only have to consider summative accounts of group knowledge. Recall that when we view our notion as a notion of knowability, we view it as specifying which conditions need to be fulfilled for there to be a possible future in which the group knowledge is *actualised*. We focus on the case where actualisation of group knowledge means that all members gain the knowledge. However, our formalism can accommodate alternatives, and we discuss some of these in the thesis as well! Our reason for our choice of focus is that it is uncontroversial among summativists that all members knowing a statement, with compatible justification, is sufficient for the group to know it. It is necessary to specify that the justification should be compatible because some epistemologists claim that, if individual members know a statement for incompatible/inconsistent reasons, then the group cannot be said to know the statement [Lac20]. Hence, showing that groups are able to collaborate so that each member learns a statement that no member could have learnt individually also shows that groups can learn more than individual members can, under any summativist account. We therefore show exactly this: that there are situations where, by collaborating, all members of the group can learn facts that no member could learn alone. Note that we take the specification that the fact is learnt via collaboration to entail that the reasons the members come to know the fact are compatible.

To link this back to the benefits of modelling described in Subsection 2.1.1, we have argued that even when our notion of potential collaborative knowledge is viewed as an account of knowability rather than knowledge, we still demonstrate a consequence of all realist views on group knowledge. This shows that this thesis still benefits the philosophical literature in the fourth way described by Hansson.

An implication of showing that some knowledge creation inherently requires collaboration is that there are irreducibly social dimensions to knowledge formation. This suggests that attention needs to be paid to collaboration in the study of knowledge formation processes. This thesis can therefore also be seen as highlighting the importance of studying social factors and interactions in epistemology and the study of learning.

2.1.5 Applications in artificial intelligence

Our notion of potential collaborative knowledge gives a new way to model collaboration and the knowledge that can be achieved via collaboration. We build on distributive knowledge by also considering ways the group can interact with the external world to gain information, and by specifying which actions can be done by the group, rather than having a predetermined selection of actions for each group. Our notion of group knowledge/knowability, therefore, has more flexibility than distributive knowledge, and encompasses more types of actions. In addition, to be able to consider the capabilities of different subgroups, we introduce the notion of an "active subgroup" of an event. This allows us to include who is executing the action in the formalism.

Our formalism would therefore be particularly appropriate for modelling instances of scientific discovery or group investigations where each agent and subgroup has limited capabilities and/or some members are capable of making observations. Our notion of potential collaborative knowledge allows us to reason about what individual members can come to know by collaborating in situations such as these. Therefore, this can be used in designing artificial systems that can reason about what groups could learn given the abilities of individual members and subgroups.

2.2 Technical background

We now move on to providing some necessary technical background for the content developed in this thesis. We focus on multi-agent epistemic logics and data-exchange logic.

2.2.1 Multi-agent epistemic logics

Multi-agent epistemic logic aims to formally study the knowledge of individuals and collections of individuals. Here, I provide a basic introduction to epistemic logic and define notions relevant for the rest of my thesis. We mostly follow notation contained in the book Dynamic Epistemic Logic by van Ditmarsch, van der Hoek, and Kooi [DHK08]. Throughout this thesis, we assume $N = \{1, ..., n\}$ is a finite, non-empty index set, representing a finite set of agents.

The language for a multi-agent epistemic logic $\mathcal{L}_K$ is given as follows (where $\mathcal{P}$ is a countable set of propositional variables):

$$\varphi := p \mid \varphi \wedge \psi \mid \neg\varphi \mid K_i\varphi$$

$$p \in \mathcal{P}, i \in N$$

In the context of epistemic logic, we interpret $K_i\varphi$ as meaning "agent i knows φ". Each K_i operator is also a box operator of standard modal logic. The idea to interpret knowledge as a Kripke modality was first implemented by Hintikka [Hin62].

Then we define the abbreviations for implication, $\varphi \to \psi := \neg(\varphi \wedge \neg\psi)$, disjunction $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, and bi-implication $\varphi \leftrightarrow \psi := \neg(\neg\varphi \wedge \psi) \wedge \neg(\varphi \wedge \neg\psi)$. We call formulas involving only a propositional variable and possibly negations literals. We also call formulas with only propositional variables and possibly also negation and/or conjunction Boolean formulas.

To be in a position to define the semantics of each of these operators, we need to first introduce Kripke frames and models from general modal logic. These definitions come from the book "Modal Logic" by Blackburn, de Rijke, and Venema [BRV01]. Later, we focus in on the frames and models that are typically included in epistemic contexts.

Definition 1 (Kripke frame)**.** *A* Kripke frame $\mathbf{F}$ *is a tuple* $(S, (R_i)_{i \in N})$*, where:*

1. *S is a non-empty set of possible worlds.*

2. *$R_i \subseteq S \times S$ is a binary relation for each $i \in N$.*

Definition 2 (Kripke model)**.** *A* Kripke model $\mathbf{M}$ *is a tuple* $(S, (R_i)_{i \in N}, \|.\|)$*, where:*

1. *$(S, (R_i)_{i \in N})$ is a Kripke frame.*

2. *$\|.\| : \mathcal{P} \to S$ is a valuation function.*

We then evaluate the truth of formulas in models as follows (cf. [BRV01]):

Definition 3 (Kripke semantics for epistemic logic)**.** *Let* $\mathsf{M} = (S, (R_i)_{i\in N}, \|.\|)$ *be a Kripke model and let* $s \in S$ *be a state. Then:*

$$\begin{array}{lll} \mathsf{M}, s \vDash p & \text{iff} & s \in \|p\| \\ \mathsf{M}, s \vDash \neg\varphi & \text{iff} & \mathsf{M}, s \nvDash \varphi \\ \mathsf{M}, s \vDash \varphi \wedge \psi & \text{iff} & \mathsf{M}, s \vDash \varphi \textit{ and } \mathsf{M}, s \vDash \psi \\ \mathsf{M}, s \vDash K_i\varphi & \text{iff} & \textit{for all } s' \in S,\ \textit{if } sR_is',\ \textit{then } \mathsf{M}, s' \vDash \varphi \end{array}$$

Definition 4 (Validity, cf. [BRV01])**.** *A formula* φ *is* valid in a Kripke model $\mathsf{M} = (S, (R_i)_{i\in N}, \|.\|)$, *denoted* $\mathsf{M} \vDash \varphi$*, exactly when, for all* $s \in S$*, we have* $\mathsf{M}, s \vDash \varphi$*. A formula* φ *is* valid in a class of models *if it is valid in every model in that class. Lastly, a formula* φ *is* valid *in general, denoted* $\vDash \varphi$*, if it is valid in all Kripke models.*

We define the notions of soundness and completeness standardly (cf. [BRV01], Section 4.1).

Definition 5 (Normal modal logics)**.** *A* normal modal logic *is a set of formulas* Σ *containing all the axioms of classical propositional logic along with:*

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

And is closed under the following inference rules:

- *Modus Ponens: From* $\varphi \to \psi$ *and* φ*, infer* ψ*.*
- *Necessitation: From* φ*, infer* $K_i\varphi$*.*

We call the smallest normal modal logic **K***.*

Theorem 6 (Soundness and completeness for **K**, cf. [BRV01])**.** *The normal modal logic* **K** *is sound and complete with respect to the class of all Kripke frames.*

Some normal modal logics have additional axioms, and are sound and complete with respect to more specific classes of frames. The table below shows some axioms commonly used for epistemic logics as well as the frame property describing the class of frames with respect to which each of them are sound and complete. Note that in the definitions of the frame properties, we quantify over all $i \in N$ and all states $s, t, u \in S$.

Name	Name in Epistemic Logic	Axiom	Frame Property
T	Factivity	$K_i\varphi \to \varphi$	Reflexivity: sR_is
4	Positive Introspection	$K_i\varphi \to K_iK_i\varphi$	Transitivity: $(sR_it \wedge tR_iu) \to sR_iu$
5	Negative Introspection	$\neg K_i\varphi \to K_i\neg K_i\varphi$	Euclideanness: $(sR_it \wedge sR_iu) \to tR_iu$

Table 2.1: Common axioms and their corresponding frame properties

.

Factivity is a common axiom in epistemic logic because it corresponds to the factivity of knowledge (i.e. that only true propositions can be known); it is generally agreed that knowledge is factive. Similarly, Positive Introspection corresponds to the property that if an agent knows a proposition, they know that they know it, which is also generally accepted, despite some objections. Most notably, Timothy Williamson was against Positive Introspection [Wil00]. Negative Introspection represents the idea that if an agent does not know a proposition, they know that they do not know it. This axiom is more controversial (see, for example [Hin62; Voo93], for arguments against it).

The smallest normal modal logic satisfying **T** and **4** is called S4, and the smallest modal logic satisfying **T**, **4**, and **5** is called S5. Note that if we have reflexivity and transitivity, euclideanness and symmetry (i.e. $sR_it \to tR_is$) are equivalent. Therefore, S4 is sound and complete with respect to the class of reflexive and transitive frames, and S5 is sound and complete with respect to the class of frames in which all relations are equivalence relations.

For our purposes, we adopt the following definition of a multi-agent epistemic model (sometimes called an epistemic model for short):

Definition 7 (Multi-agent epistemic model)**.** *A* multi-agent epistemic model $\mathsf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$ *(or S5 model) is a Kripke model where all relations* $\sim_i$ *are equivalence relations, i.e. reflexive, transitive, and symmetric.*

We use the notation $\sim_i$ rather than R_i because we are dealing only with equivalence relations. This definition of epistemic models takes epistemic models to be Kripke models based on frames validating the axioms of S5 (K, Factivity, Positive Introspection, and Negative Introspection). Intuitively, this means that knowledge is factive and all agents have knowledge of what they do and do not know. However, other choices are possible. For example, given the controversy around Negative Introspection, it is also common to define epistemic models as transitive and reflexive models (i.e. based on frames satisfying the axioms of S4). Intuitively, this means agents have knowledge of what they know, but not of what they do not know.

Another modality (in addition to the knowledge modality K_i) that is often studied in multi-agent epistemic logics is the distributive knowledge modality. The syntax for multi-agent epistemic logic with the distributed knowledge modality D added is:

$$\varphi := p \mid \varphi \wedge \psi \mid \neg\varphi \mid K_i\varphi \mid D\varphi$$

$$p \in \mathcal{P}, i \in N$$

In terms of the semantics, distributed knowledge is the Kripke box modality for the relation obtained by intersecting all the other relations, that is, $\bigcap_{i\in N} \sim_i$. The following remark explains a helpful abbreviation we will use throughout the thesis.

Remark 8. *For the rest of this thesis, for* $H \subseteq N$ *we use the abbreviation* $a \sim_H b$ *to mean* a *and* b *are indistinguishable for all* $i \in H$*. In other words,* $\sim_H$ *is the intersection of all* $\sim_i$*s where* $i \in H$*:* $\sim_H = \bigcap_{i\in H} \sim_i$*. Note that* $\sim_{\{i\}} = \sim_i$*.*

The semantics for distributed knowledge using the abbreviation in the remark is given below.

Definition 9. *Given a multi-agent epistemic model* $\mathsf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$ *and a state* $s \in S$*, the semantics of* $D\varphi$ *is given by:*

$$\mathsf{M}, s \models D\varphi \qquad \text{iff} \quad \textit{for all } s' \in S, \textit{ if } s \sim_N s', \textit{ then } \mathsf{M}, s' \models \varphi.$$

Epistemic logic with distributive knowledge has been axiomatised. Its complete axiomatisation consists of all axioms and inference rules of classical propositional logic, the axioms and inference rules of S5 for each K_i modality, the axioms and inference rules of S5 for the D modality, and, for each $i \in N$, the axiom $K_i\varphi \to D\varphi$.

We can also consider a distributed knowledge operator restricted to a subgroup G, denoted D_G. Given a multi-agent epistemic model $\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$ and a state $s \in S$, the semantics for $D_G\varphi$ is given as follows:

$$\mathbf{M}, s \vDash D_G\varphi \quad \textit{iff} \quad \text{for all } s' \in S, \text{ if } s \sim_G s', \text{ then } \mathbf{M}, s' \vDash \varphi.$$

Hence D_G is the Kripke box modality for the relation obtained by intersecting all the other relations for members of G, that is, $\bigcap_{i\in G} \sim_i$.

This concludes the relevant parts of static epistemic logic. Next, we cover the dynamic counterpart we use in this thesis.

2.2.2 Data-exchange logics

Dynamic epistemic logics add events/actions to epistemic logics (cf. [BMS98], [BMS16], [DHK08]). Here, we consider multi-agent standard dynamic epistemic logic with data-exchange (which makes use of data-exchange event models). This logic was developed by Baltag and Smets, and all the content in this section comes from the paper "Logics for Data Exchange and Communication" [BS24]. The reason we use this dynamic extension of epistemic logic rather than standard dynamic epistemic logic is that in our semantics presented in this thesis, we also consider events in which agents share all they know with other agents.

The syntax for this logic $\mathcal{L}_{de}$ is given recursively as:

$$\varphi := p \mid \varphi \wedge \psi \mid \neg\varphi \mid D_G\varphi \mid [e]\varphi$$

In addition, we define $\langle e\rangle$ as the dual of $[e]$. In other words, $\langle e\rangle\varphi$ is the abbreviation for $\neg[e]\neg\varphi$.

The event models used in data-exchange logic are defined as follows:

Definition 10 (Data-exchange event model)**.** *A* data-exchange event model *is a tuple* $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$, *where:*

1. E *is a non-empty set of actions (also called events).*

2. $\sim_i$ *is an equivalence relation for each agent* $i \in N$*, called its indistinguishability relation.*

3. $\bullet(\bullet) : \mathsf{E} \to (N \to \mathcal{P}(N))$ *is a function which assigns to each event* $e \in \mathsf{E}$ *an access map* $e(\bullet) : N \to \mathcal{P}(N)$*, which satisfies two conditions for all* $i \in N$ *and* $e, f \in \mathsf{E}$*:*

   (a) $i \in e(i)$.

   (b) If $e \sim_i f$ *then* $e(i) = f(i)$.

4. $pre : \mathsf{E} \to \mathcal{L}_{de}$ *is a precondition map which assigns a precondition formula* pre_e *to each event* e.

These event models are the event models of standard dynamic epistemic logic $(\mathbf{E}, (\sim_i)_{i\in N}, pre)$ with an added access map $\bullet(\bullet)$. The idea behind the access map is that it specifies for each event e, which agents gain access to which other agents' databases in that event, where $e(i)$ is the set of all agents whose databases i can access.

We can extend the access map for a given event e to map subgroups to sets of agents by, for a given subgroup G, taking the union of all outputs of the access maps for agents in G, $e(G) = \bigcup_{i \in G} e(i)$. Here, intuitively, $e(G)$ is the set of all agents whose databases some member of G can access.

To represent an access map where i accesses j's data (as well as their own) and all agents except for i access only their own data, we write $i : j$.

Remark 11 (Notation for indistinguishability relations)**.** *Throughout this thesis, we use the notation $a \sim_i b$ to mean "a and b are indistinguishable to agent i in the model they both belong to". In other words, we do not usually specify which model the relation belongs to when it is clear from context. When the model is not clear from context, we clarify it by writing the model as a superscript above the relation. For example, we write $\sim_i^{\mathbf{M}}$ for $\mathbf{M}$'s indistinguishability relation.*

The product update of an epistemic model with an event model represents how the knowledge states of agents are changed by the event(s) that the event model represents. It is defined as follows:

Definition 12 (Product update)**.** *Let $\mathbf{M} = (S, (\sim_i)_{i \in N}, \|.\|)$ be a multi-agent epistemic model and let $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i \in N}, \bullet(\bullet), pre)$ be a data-exchange event model. Then the product update, $\mathbf{M} \otimes \mathbf{E} = (S \otimes \mathsf{E}, (\sim'_i)_{i \in N}, \|.\|')$ is defined as follows:*

1. *The domain of worlds consists of all consistent pairs:*

$$S \otimes \mathsf{E} = \{(s, e) \in S \times \mathsf{E} \mid \mathbf{M}, s \vDash pre_e\}$$

2. *The indistinguishability relations are given as follows:*

$$(s, e) \sim'_i (t, f) \iff s \sim_{e(i)} t \text{ and } e \sim_i f$$

3. *The valuation stays the same:*

$$(s, e) \in \|p\|' \iff s \in \|p\|.$$

The idea here is that the new knowledge of each agent is obtained by adding the knowledge of each agent who shares all they know with them.

Let $\mathbf{M}$ be an epistemic model and $\mathbf{E}$ be a data exchange event model. The semantics of all operators also present in $\mathcal{L}_{de}$ is defined as before. The semantic clause for $[e]\varphi$ is defined as follows:

$$\mathbf{M}, s \vDash [e]\varphi \quad \textit{iff} \quad \text{if } (s, e) \in S \otimes \mathsf{E}, \text{ then } \mathbf{M} \otimes \mathbf{E}, (s, e) \vDash \varphi$$

In some cases, we will also want to sequentially compose events. For this, we also define the composition of two data-exchange event models.

Definition 13 (Sequential composition of events)**.** *Let $\mathbf{E}_1 = (\mathsf{E}_1, (\sim_i^1)_{i \in N}, \bullet(\bullet)_1, pre^1), \mathbf{E}_2 = (\mathsf{E}_2, (\sim_i^2)_{i \in N}, \bullet(\bullet)_2, pre^2)$ be two data exchange event models. Then $\mathbf{E}_1; \mathbf{E}_2 = (\mathsf{E}_1 \times \mathsf{E}_2, (\sim'_i)_{i \in N}, \bullet(\bullet)', pre')$ is defined as follows:*

1. *The domain is Cartesian product $\mathsf{E}_1 \times \mathsf{E}_2$, we denote the pair (e, f) as $(e; f)$ or $e; f$.*

2. *for all $i \in N$, $(e; f), (e'; f') \in \mathbf{E}_1 \times \mathbf{E}_2$:*

$$(e; f) \sim'_i (e'; f') \iff e \sim^1_{f(i)} e' \text{ and } f \sim^2_i f'.$$

3. The access map $\bullet(\bullet)'$ is given as follows:

$$(e;f)(i) = e(f(i)) = \bigcup_{j \in f(i)} e(j).$$

4. For all $(e;f) \in \mathbf{E}_1 \times \mathbf{E}_2$, $pre'_{(e;f)} = \langle e \rangle pre_f$.

We note that this does indeed represent the composition of events:

Proposition 14 (Sequential Composition of Event Models is Sequential Composition of Updates)**.**

$$\mathbf{M} \otimes (\mathbf{E}_1; \mathbf{E}_2) = (\mathbf{M} \otimes \mathbf{E}_1) \otimes \mathbf{E}_2$$

The following remark explains some notation used throughout this thesis.

Remark 15 (Notation for Event Compositions)**.** *Let $m \in \omega$ be a natural number. We use the notation $\mathbf{E}^m$ to mean $\mathbf{E}; ...; \mathbf{E}$ (m times). We also define the sets E^m and E^* to mean the set of composite events of m many compositions of elements of E, and of finitely many compositions of elements of E, respectively. Note that in the context of an epistemic model $\mathbf{M}$ with domain S, we can equivalently view $\mathbf{M} \otimes \mathbf{E}^m$ as an abbreviation for $\mathbf{M} \otimes \mathbf{E} \otimes ... \otimes \mathbf{E}$ (m times) and $S \otimes \mathsf{E}^m$ as an an abbreviation for $S \otimes \mathsf{E} \otimes ... \otimes \mathsf{E}$ (m times).*

Note that since sequential composition of event models is sequential composition of updates, given an epistemic model $\mathbf{M}$ and an event model $\mathbf{E}$, in the case where e is the composite event $(e_1; ...; e_m)$, (s, e) is equivalent to $(s, e_1, ..., e_m)$ (i.e., they are both the same state in $\mathbf{M} \otimes \mathbf{E}^m$).

2.2.3 Generated submodels

One other concept that will be important for us later is that of a generated submodel. The content in this subsection is adapted from [BRV01].

Definition 16 (Generated submodel for epistemic model)**.** *Let $\mathbf{M} = (S, (R_i)_{i \in N}, \|.\|)$ and $\mathbf{M}' = (S', (R'_i)_{i \in N}, \|.\|')$ be two Kripke models. We say that $\mathbf{M}'$ is a* submodel *of $\mathbf{M}$ if and only if the following hold:*

1. $S' \subseteq S$.

2. For each $i \in N$, R'_i is the restriction of R_i to S': $R'_i = R_i \cap (S' \times S')$ for all $i \in N$.

3. The valuation in $\mathbf{M}'$ is the restriction of the valuation in $\mathbf{M}$ to S': for each $p \in \mathcal{P}$, $\|p\|' = \|p\| \cap S'$.

We say that $\mathbf{M}'$ is a generated submodel *of $\mathbf{M}$ if and only if $\mathbf{M}'$ is a submodel of $\mathbf{M}$ and the following closure condition holds:*

$$\text{if } s \in S' \text{ and } sR_it, \text{ then } t \in S'.$$

For a state $s \in S$, the submodel of $\mathbf{M}$ generated by a state $s \in S$ (denoted $\mathbf{M}_s$) is the smallest generated submodel of $\mathbf{M}$ with s in its domain.

The definition below specifies how the concept of a generated submodel is extended to event models. It is essentially identical to the previous definition, just with the added conditions on the preconditions and access maps, and without the condition on valuations.

Definition 17 (Generated submodel for event model)**.** *Let* $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ *and* $\mathsf{E}' = (\mathsf{E}', (\sim_i)_{i\in N}, \bullet(\bullet)', pre')$ *be two data-exchange models. We say that* E' *is a* submodel *of* E *if and only if the following hold:*

1. $\mathsf{E}' \subseteq \mathsf{E}$.
2. *For each* $i \in N$, $\sim'_i$ *is the restriction of* $\sim_i$ *to* E': $\sim'_i = \sim_i \cap (\mathsf{E}' \times \mathsf{E}')$ *for all* $i \in N$.
3. *The preconditions for events in* E' *stay the same: for all* $e \in \mathsf{E}'$, $pre'_e = pre_e$.
4. *The access maps for events in* E' *stay the same: for all* $e \in \mathsf{E}'$, *and all* $i \in N$ $e(i)' = e(i)$.

We say that E' *is a* generated submodel *of* E *if and only if* E' *is a submodel of* M *and:*

$$\textit{if } e \in \mathsf{E}' \textit{ and } e \sim_i f, \textit{ then } f \in \mathsf{E}'.$$

For an event $e \in \mathsf{E}$*, the submodel of* E *generated by a state* $e \in \mathsf{E}$ *(denoted* E_e*) is the smallest generated submodel of* E *with* e *in its domain.*

We also have that modal formulas are invariant under generated submodels. The following proposition states this formally:

Proposition 18 (Invariance Under Generated Submodels)**.** *Let* $\mathsf{M}' = (S', (R'_i)_{i\in N}, \|.\|')$ *be a generated submodel of* $\mathsf{M} = (S, (R_i)_{i\in N}, \|.\|)$*, and let* $s \in S'$ *be a state. Then for all formulas* $\varphi \in \mathcal{L}_K$*:*

$$\mathsf{M}, s \vDash \varphi \iff \mathsf{M}', s \vDash \varphi.$$

The following proposition is useful later, and is a corollary to Invariance Under Generated Submodels (Proposition 18).

Proposition 19 (Properties of Generated Submodels of Event Models)**.** *Let* M *be an epistemic model and* E *be a data-exchange event model. Additionally, let* $s \in S$ *be a state and* $\mathsf{E}_{e_1}, ..., \mathsf{E}_{e_m}$ *be submodels generated by* $e_1, ..., e_m$*, respectively. Then both of the following hold:*

1. $(s, e_1, ..., e_m) \in \mathsf{M} \otimes \mathsf{E}^m \iff (s, e_1, ..., e_m) \in \mathsf{M} \otimes \mathsf{E}_{e_1} \otimes ... \otimes \mathsf{E}_{e_m}$
2. $\mathsf{M} \otimes \mathsf{E}^m, (s, e_1, ..., e_m) \vDash \varphi \iff \mathsf{M} \otimes \mathsf{E}_{e_1} \otimes ... \otimes \mathsf{E}_{e_m}, (s, e_1, ..., e_m) \vDash \varphi$

Proof. This follows directly from Invariance Under Generated Submodels and the fact that if E' is a generated submodel of E, then $\mathsf{M} \otimes \mathsf{E}'$ is a generated submodel of $\mathsf{M} \otimes \mathsf{E}$. □

Chapter 3

Potential Collaborative Knowledge

In this chapter, we formally introduce our notion of group knowledge/knowability, namely, potential collaborative knowledge. For this, we first formally define actions using the data-exchange event models introduced in Subsection 2.2.2. Next, we define our syntax and semantics. The main operator in our syntax is not potential collaborative knowledge, but another operator that we call potential collaborative distributive knowledge. We use this as the primary operator instead to have a more general framework and for technical reasons relating to the axiomatisation, but we still define our main notion of philosophical interest, potential collaborative knowledge, in terms of it. After defining the semantics, we discuss some aspects of the semantics, including ways to model alternative understandings of group knowledge using our formalism. Thereafter, we formalise the telescope example introduced in Subsection 2.1.3. This serves the purpose of both demonstrating the workings of our semantics and justifying our main philosophical points explained in Section 2.1. We end this section by formally comparing potential collaborative knowledge to distributive knowledge, thereby formally justifying our earlier claims that potential collaborative knowledge is a generalisation of distributive knowledge.

3.1 Observation and sharing actions

Since we are interested in the formation of new knowledge via collaboration, we need to formally represent actions which members can take during a collaboration process. We have decided to focus on observation and information-sharing actions. We have also decided to model all these events using data-exchange event models because this allows us to have uniform notation. The preconditions of our event models use the syntax of multi-agent epistemic logic with distributed knowledge for subgroups $\mathcal{L}_D$. This is defined recursively as follows:

$$\varphi := p \mid \top \mid \varphi \wedge \psi \mid \neg\varphi \mid D_G\varphi$$

Where $G \subseteq N$. Here, we define $K_i\varphi$ as $D_{\{i\}}\varphi$.

3.1.1 Permissible event models

We start by formalising observation actions, and we consider two types of observations. The first type is where an agent i observes an ontic fact p or $\neg p$ with the help of a finite group of agents G (including i) without needing any knowledge. We call these *basic observations*, and they are represented by $(i :_G p)$ or $(i :_G \neg p)$ (depending on the outcome of the observation). The second type is where an agent

i observes an ontic fact using some knowledge they have, ψ, to be able to make the observation. We call these *informed observations*, and we write these as $(\psi/i : p)$ (representing a successful observation of p where i knows ψ), $(\psi/i : \neg p)$ (representing a successful observation of $\neg p$ where i knows ψ), or $(\psi/i : -_p)$ (representing a failed attempt at observing the truth of p). An example of an observation requiring some background knowledge is an agent looking at a clock to attempt to learn the time in their home country while travelling. To be able to learn the time in their home country, the agent needs to know which time zone they are in.

Next, we define the event models we use to represent basic and informed observations, as well as the active subgroup of each event model, which, intuitively, is the set of agents needed to perform that observation.

Definition 20 (Observation event models and active subgroups)**.** *An observation is represented by an event model called an* observation event model*. In the case of a* basic observation *of the truth of a propositional variable p by agent i with the help of a subgroup G, the corresponding event model $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ is given as follows:*

1. $\mathsf{E} = \{(i :_G p), (i :_G \neg p)\}$.
2. *For all $e, f \in \mathsf{E}$, $e \sim_i f$ iff $e = f$.*
3. *For all $e, f \in \mathsf{E}$, $e \sim_j f$, for all $j \neq i$.*
4. *The access map is trivial: $e(j) = \{j\}$ for all $e \in \mathsf{E}$ and all $j \in N$.*
5. $pre((i :_G p)) = p$ *and* $pre((i :_G \neg p)) = \neg p$.

For each basic observation event $e = (i :_G l)$, its active subgroup *is $G \cup \{i\}$.*

In the case of an informed observation *of the truth of a propositional variable p by an agent i, requiring i to have knowledge of the boolean formula ψ, the corresponding event model $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ is given as follows:*

1. $\mathsf{E} = \{(\psi/i : p), (\psi/i : \neg p), (\psi/i : -_p)\}$.
2. *For all $e, f \in \mathsf{E}$, $e \sim_i f$ iff $e = f$.*
3. *$e \sim_j f$ for all $e, f \in \mathsf{E}$, for all $j \neq i$.*
4. *The access map is trivial: $e(i) = \{i\}$ for all $e \in \mathsf{E}$ and all $i \in N$.*
5. *The preconditions have both an ontic and an epistemic component. The epistemic component is represented as a Boolean formula ψ. The ontic component is the object of the observations p. The preconditions are then the following:*
   (a) $pre_{(\psi/i:p)} = p \wedge K_i\psi$ *(representing a successful observation of p).*
   (b) $pre_{(\psi/i:\neg p)} = \neg p \wedge K_i\psi$ *(representing a successful observation of $\neg p$).*
   (c) $pre_{(\psi/i:-_p)} = \neg K_i\psi$ *(representing a failed observation attempt).*

For each informed observation event $e = (\psi/i : l)$ or $e = (\psi/i : -_p)$, its active subgroup *is $\{i\}$.*

The model transformed by the observation is then given by the product update $\mathsf{M} \otimes \mathsf{E}$. The access maps are included here despite both being trivial, because we consider them along with event models with non-trivial access maps and want to include both of these types together in one big event model.

We can equivalently view these observations as model transformations defined directly on the epistemic model. The following remark explains how a basic observation action transforms a multi-agent epistemic model.

Remark 21. *Let $\mathsf{M} = (S, (\sim_i)_{i \in N}, \|.\|)$ be a multi-agent epistemic model. The transformed model after a basic observation E by agent i with the help of subgroup $G \subseteq N$ of the truth of a propositional variable p, $\mathsf{M} \otimes \mathsf{E}$, is isomorphic to the model $\mathsf{M}_{(i:_G p?)} = (S', (\sim'_i)_{i \in N}, \|.\|')$ where:*

1. *The state set stays the same: $S' = S$. This is because the preconditions of the events in E are mutually contradictory and their disjunction is a tautology (w.r.t. the classical propositional logic).*
2. *Previously indistinguishable states become distinguishable for i when they disagree on the truth of p: For all $s, t \in S$, we have $s \sim'_i t$ if and only if $s \sim_i t$ and one of the following are satisfied:*
   (a) $\mathsf{M}, s \vDash p$ *and* $\mathsf{M}, t \vDash p$
   (b) $\mathsf{M}, s \vDash \neg p$ *and* $\mathsf{M}, t \vDash \neg p$

   This is because the events $(i :_G p)$ and $(i :_G \neg p)$ are distinguishable for i and have preconditions p and $\neg p$, respectively.
3. *The relations stay the same for all other agents: For all $j \in N$ such that $j \neq i$, and all $s, t \in S$, we have $s \sim'_j t$ if and only if $s \sim_j t$. This is because both events are indistinguishable for all other agents.*
4. *The valuation stays the same: For all $s \in S'$, $s \in \|q\|'$ if and only if $s \in \|q\|$.*

Note that for basic observations, if the subgroup G has more agents than just i, then the observation is a collaborative effort by construction. There are different reasons why an agent may need the help of others to make an observation. One is just that the action needed to make the observation is itself complicated and cannot be done by only one agent. Another is that agents require resources (such as a telescope or a submarine) that are owned by other agents, and that they therefore need the collaboration of these other agents. Note that here we did not update the knowledge of any active agents other than i, since we have assumed that i is the only agent that sees the outcome of the observation.

The following definition explains how an informed observation action transforms a multi-agent epistemic model.

Remark 22. *Let $\mathsf{M} = (S, (\sim_i)_{i \in N}, \|.\|)$ be a multi-agent epistemic model. The transformed model after an informed observation E by an agent $i \in N$ of the truth of a propositional variable p requiring the agent to have knowledge of the boolean formula ψ, $\mathsf{M} \otimes \mathsf{E}$, is isomorphic to $\mathsf{M}_{(\psi/i:p?)} = (S', (\sim'_i)_{i \in N}, \|.\|')$, where:*

1. *The state set stays the same: $S' = S$. Again, because the preconditions of the events in E are mutually contradictory and their disjunction is a tautology (w.r.t. the classical propositional logic).*

2. *In the case where i has required knowledge for the observation, previously indistinguishable states become distinguishable when they disagree on the truth of p: For all* $s, t \in S$*, we have* $s \sim'_i t$ *if and only if* $s \sim_i t$ *and one of the following are satisfied:*

   (a) $\mathbf{M}, s \vDash \neg K_i\psi$ *and* $\mathbf{M}, t \vDash \neg K_i\psi$

   (b) $\mathbf{M}, s \vDash p \wedge K_i\psi$ *and* $\mathbf{M}, t \vDash p \wedge K_i\psi$

   (c) $\mathbf{M}, s \vDash \neg p \wedge K_i\psi$ *and* $\mathbf{M}, t \vDash \neg p \wedge K_i\psi$

   This is because the events $(\psi/G : -_p)$*,* $(\psi/G : p)$ *and* $(\psi/G : \neg p)$ *are distinguishable for i and have preconditions* $\neg K_i\psi$*,* $p \wedge K_i\psi$*, and* $\neg p \wedge K_i\psi$*, respectively.*

3. *For all other agents, the valuation stays the same: For all* $j \in N$ *such that* $j \neq i$*, and all* $s, t \in S$*, we have* $s \sim'_j t$ *if and only if* $s \sim_j t$*. Again, because all events are indistinguishable for all other agents.*

4. *The valuation stays the same: For all* $s \in S'$*,* $s \in \|q\|'$ *if and only if* $s \in \|q\|$*.*

The other type of action we consider in the collaboration processes we model is sharing. We have decided to represent only cases of one agent j sharing all they know with another agent i, represented by $!(i : j)$. This is sufficient, since wider sharing of all an agent knows can still be represented by iterative agent-to-agent sharing. For example, if we have three agents, Nuha, Thabo, and Sindi, then Nuha sharing all she knows with both Thabo and Sindi simultaneously is equivalent to Nuha sharing all she knows first with Thabo, and then with Sindi. We are also assuming information is not explicitly hidden from group members since we are working in a collaborative context (we therefore do not need to represent secret sharing).

Definition 23 (Sharing event models and active subgroups)**.** *An instance of sharing is represented by a* sharing event model*. In the case of agent j sharing all they know with i, the corresponding event model* $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ *is given as follows:*

1. $\mathsf{E} = \{!(i : j)\}$*. The idea behind this notation is that this is the event where agent j shares everything they know with agent i.*

2. $\sim_k = \{(!(i : j), !(i : j))\}$ *for all* $k \in N$*.*

3. $!(i : j)(i) = \{i, j\}$ *and* $!(i : j)(k) = \{k\}$ *for all* $k \neq i$*.*

4. $pre_{!(i:j)} = \top$*.*

For each sharing event $e = !(i : j)$*, its* active subgroup*, is* $\{i, j\}$*.*

We can again equivalently view sharing as a model transformation defined directly on epistemic models.

Remark 24. *Let* $\mathbf{M} = (S, (\sim_k)_{k\in N}, \|.\|)$ *be a multi-agent epistemic model. The transformed model after agent j shares all they know with agent i (i.e. after sharing action* $\mathbf{E}$*),* $\mathbf{M} \otimes \mathbf{E}$*, is isomorphic to* $\mathbf{M}_{!(i:j)} = (S', (\sim'_k)_{k\in N}, \|.\|')$*, where:*

1. *The state set stays the same:* $S' = S$*.*

2. *Previously indistinguishable states become distinguishable for i when they are distinguishable for j: For all $s, t \in S$, we have $s \sim_i' t$ if and only if $s \sim_i t$ and $s \sim_j t$. Equivalently, $\sim_i' = \sim_i \cap \sim_j$.*

3. *For all $k \in N$ such that $k \neq i$, and all $s, t \in S$, we have $s \sim_k' t$ if and only if $s \sim_k t$.*

4. *The valuation stays the same: For all $s \in S'$, $s \in \|p\|'$ if and only if $s \in \|p\|$.*

The following remark is helpful for thinking about states in our product updates.

Remark 25. *Let the epistemic model $\mathsf{M} = (S, (\sim_i)_{i \in N}, \|.\|)$ be given. Note that for any basic event model E, taking the product update of M with that event model does not change the state set. This means that for any basic event model E, for any state $s \in \mathsf{M}$, there is exactly one event $e \in \mathsf{E}$ such that $(s, e) \in \mathsf{M} \otimes \mathsf{E}$. Therefore, in the case where we update with a basic event model, and the model is clear from context, we may write s instead of (s, e). In addition, if E is a permissible event model and $\mathsf{E}_{e_1}, ..., \mathsf{E}_{e_m}$ are the submodels generated from $e_1, ..., e_m$, respectively, then for each $s \in S$, there is a unique chain of events $e_1' \in \mathsf{E}_{e_1}, ..., e_m' \in \mathsf{E}_{e_m}$ such that $(s, e_1', ..., e_m') \in \mathsf{M} \otimes \mathsf{E}_{e_1} \otimes ... \otimes \mathsf{E}_{e_m}$. We can therefore write s instead of $(s, e_1', ..., e_m')$ if the generated submodels we have updated with are clear.*

The next notion we define is the core notion of this section, and is used in the semantics for our main language.

Definition 26 (Permissible events and event models)**.** *A* basic event model *is either an observation or sharing event model. A* permissible event model *is an event model $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i \in N}, \bullet(\bullet), pre)$ where each generated submodel of E is a basic event model. A* permissible event set E *is the domain of a permissible event model $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i \in N}, \bullet(\bullet), pre)$. A* permissible event *$e$ is any event that is a member of some permissible event set (also called a* permissible action*).*

Note that since permissible event models are S5, their generated submodels form a partition of the domain.

Remark 27. *A permissible event set consists only of events of the form $(i :_G l)$, $(\psi/j : p)$, $(\psi/j : \neg p)$ $(\psi/j : -_p)$, or $!(j : k)$ (for $G \subseteq N$, literal l, propositional letter p, boolean formula ψ, and $i, j, k \in N$). Additionally, it is closed in the sense that if one of $(i :_G p)$, or $(i :_G \neg p)$ is in the set, they both are. Similarly, if one of $(\psi/j : p)$, $(\psi/j : \neg p)$, or $(\psi/j : -_p)$ is in the set, they all are.*

The following remark shows that the domain of a permissible event set fully determines the rest of the event model. This means that, for permissible event models, we do not need to specify the full event model – it is enough to just specify the domain.

Remark 28. *Each permissible event set E uniquely determines an event model E_{E} in which each generated submodel is a basic event model. This event model is defined as $\mathsf{E}_{\mathsf{E}} = (\mathsf{E}, (\sim_i)_{i \in N}, \bullet(\bullet), pre)$, where:*

1. *For $e, f \in \mathsf{E}$, $e \sim_i f$ iff one of the following hold:*
   (a) *$e = f$.*
   (b) *There is a subgroup $G \subseteq N$, a member $j \in G$ where $j \neq i$, and a propositional variable p such that e and f are both one of $(j :_G p)$, or $(j :_G \neg p)$.*

(c) *There is an agent $j \in N$ where $j \neq i$, a boolean formula ψ, and a propositional variable p such that both e and f are one of $(\psi/j : p)$, $(\psi/j : \neg p)$, or $(\psi/j : -_p)$.*

2. *If $e =!(i : j)$ for some $i, j \in N$, then $e(i) = \{i, j\}$ and $e(k) = \{k\}$ for all agents $k \neq i$. Otherwise $e(k) = \{k\}$ for all agents $k \in N$.*

3. *pre_e is defined by cases as follows:*

   (a) *If $e = (i :_G l)$ for some $G \subseteq N$, $i \in N$, and literal l, then $pre_e = l$.*

   (b) *If $e = (\psi/i : l)$ for some $i \in N$, Boolean formula ψ, and literal l, then $pre_e = K_i\psi \wedge l$.*

   (c) *If $e = (\psi/i : -_p)$ for some $i \in N$, Boolean formula ψ, and propositional variable p, then $pre_e = \neg K_i\psi$.*

   (d) *If $e =!(i : j)$ for some $i, j \in N$, then $pre_e = \top$.*

3.1.2 Properties of permissible events

The event models we introduced in the previous section were designed to satisfy some specific properties that make them easy to work with. In this subsection, we explain and prove these properties. The last property proven in this section, Concatenation of Events, is particularly important for the rest of the thesis.

The following remark is helpful in understanding the nature of product updates with permissible event models.

Remark 29. *Let $\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$ be an epistemic model and $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$ be a permissible event model. Let $e_1, ...e_m \in \mathsf{E}$ be events. For any $(s, e_1, ..., e_m), (s', e'_1, ..., e'_m) \in S \otimes \mathsf{E}_{e_1}, ..., \mathsf{E}_{e_m}$, by Remark 25 we can write $s \sim_i^{\mathbf{M}\otimes\mathbf{E}_{e_1}\otimes...\otimes\mathbf{E}_{e_m}} s'$ instead of $(s, e_1, ..., e_m) \sim_i^{\mathbf{M}\otimes\mathbf{E}_{e_1}\otimes...\otimes\mathbf{E}_{e_m}} (s', e'_1, ..., e'_m)$ (without loss of information). In addition, note that for all $(s, e_1, ..., e_m) \in \mathbf{M} \otimes \mathbf{E}^m$, the only states $(s', e'_1, ..., e'_m) \in \mathbf{M} \otimes \mathbf{E}^m$ that could possibly satisfy $(s, e_1, ..., e_m) \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s', e'_1, ..., e'_m)$ in $\mathbf{M} \otimes \mathbf{E}^m$ are elements of $\mathbf{M} \otimes \mathbf{E}_{e_1} \otimes ... \otimes \mathbf{E}_{e_m}$, since the generated submodels form a partition. We also generally have invariance under generated submodels by Proposition 18. Hence, for all $(s, e_1, ..., e_m), (s', e'_1, ..., e'_m) \in S \otimes \mathsf{E}^m$, we have:*

$$(s, e_1, ..., e_m) \sim_i^{\mathbf{M}\otimes\mathbf{E}_{e_1}\otimes...\otimes\mathbf{E}_{e_m}} (s', e'_1, ..., e'_m) \iff (s, e_1, ..., e_m) \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s', e'_1, ..., e'_m)$$

Therefore, for all $s, s' \in S$ and $e_1, ..., e_m \in \mathsf{E}$ with $(s, e_1, ..., e_m) \in S \otimes \mathsf{E}^m$:

$$s \sim_i^{\mathbf{M}\otimes\mathbf{E}_{e_1}\otimes...\otimes\mathbf{E}_{e_m}} s' \iff (s, e_1, ..., e_m) \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s', e'_1, ..., e'_m) \qquad (\diamond)$$

for the only chain of events $e'_1, ..., e'_m$ that can be paired with s' in $\mathbf{M} \otimes \mathbf{E}_{e_1} \otimes ... \otimes \mathbf{E}_{e_m}$.

This allows us to see $\sim_i^{\mathbf{M}\otimes\mathbf{E}_{e_1}\otimes...\otimes\mathbf{E}_{e_m}}$ as a relation on S for any events $e_1, ..., e_m \in \mathsf{E}$.

The first lemma is about what happens to indistinguishability relations in product updates. In the lemma, we view relations in the product update as relations on S as explained in the previous remark. We therefore use the equivalence in $(\diamond)$ throughout the lemma.

Lemma 30. *Let $\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$ be a multi-agent epistemic model and let $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ be a premissible event model. Let $e_1, ..., e_m, f_1, ...f_k \in \mathsf{E}$ be events. Then for all $i \in N$:*

$$\sim_i^{\mathbf{M}\otimes\mathbf{E}_{e_1}\otimes...\otimes\mathbf{E}_{e_m}\otimes\mathbf{E}_{f_1}\otimes...\otimes\mathbf{E}_{f_k}} \subseteq \sim_i^{\mathbf{M}\otimes\mathbf{E}_{f_1}\otimes...\otimes\mathbf{E}_{f_k}}$$

Proof. Note first that since updating with one of our event models gives an epistemic model where the state space is the same and the indistinguishability relation for each agent is a subset of the original relation for that agent, we have $\sim_i^{\mathsf{M}\otimes\mathsf{E}_{e_1}\otimes...\otimes\mathsf{E}_{e_m}} \subseteq \sim_i^{\mathsf{M}}$. We proceed by induction on k. We already have the base case ($k = 0$). Now let $k = j + 1$. Then by our induction hypothesis, $\sim_i^{\mathsf{M}\otimes\mathsf{E}_{e_1}\otimes...\otimes\mathsf{E}_{e_m}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_j}} \subseteq \sim_i^{\mathsf{M}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_j}}$, for all $i \in N$. Let $(s, e_1, ..., e_m, f_1, ...f_{j+1})$, $(s', e'_1, ..., e'_m, f'_1, ...f'_{j+1}) \in S \otimes \mathsf{E}^{m+j+1}$, and assume $s \sim_i^{\mathsf{M}\otimes\mathsf{E}_{e_1}\otimes...\otimes\mathsf{E}_{e_m}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_{j+1}}} s'$. Then by the definition of product updates, $s \sim_{f_{j+1}(i)}^{\mathsf{M}\otimes\mathsf{E}_{e_1}\otimes...\otimes\mathsf{E}_{e_m}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_j}} s'$ and $f_{j+1} \sim_i f'_{j+1}$. Then, by our induction hypothesis, $s \sim_{f_{j+1}(i)}^{\mathsf{M}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_j}} s'$. Then again by the definition of product updates, $s \sim_i^{\mathsf{M}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_{j+1}}} s'$. This completes our proof. □

The following states that for an epistemic model and a data-exchange event model, preconditions representing successful observations or sharing are preserved under product updates.

Lemma 31 (Preservation of Preconditions)**.** *Let* $\mathsf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$ *be a multi-agent epistemic model and let* $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ *be a permissible event model. Let* $s \in S$*,* $m, k \in \omega$ *and* $e, e_1, ..., e_m, f_1, ..., f_k \in \mathsf{E}$ *be such that* $(s, e) \in S\otimes\mathsf{E}$*,* $(s, f_1, ..., f_k) \in S\otimes\mathsf{E}^k$*, and* $(s, e_1, ..., e_m, f_1, ..., f_k) \in S \otimes \mathsf{E}^{k+m}$ *be arbitrary. Then for any Boolean formula* ψ *and agent* $i \in N$*, both of the following hold:*

1. $\mathsf{M}, s \vDash \psi \iff \mathsf{M} \otimes \mathsf{E}, (s, e) \vDash \psi$
2. $\mathsf{M} \otimes \mathsf{E}^k, (s, f_1, ..., f_k) \vDash K_i\psi \implies \mathsf{M} \otimes \mathsf{E}^{k+m}, (s, e_1, ..., e_m, f_1, ..., f_k) \vDash K_i\psi$

Proof. 1. This proof is done by induction on the complexity of the formula.
For propositional variables, this is trivial, since for any $p \in \mathcal{P}$:

$$s \in \|p\|_{\mathsf{M}} \iff (s, e) \in \|p\|_{\mathsf{M}\otimes\mathsf{E}}$$

For $\psi = \neg\chi$:

$$\begin{aligned} \mathsf{M}, s \vDash \neg\chi &\iff \mathsf{M}, s \nvDash \chi \\ &\overset{IH}{\iff} \mathsf{M} \otimes \mathsf{E}, (s, e) \nvDash \chi \\ &\iff \mathsf{M} \otimes \mathsf{E}, (s, e) \vDash \neg\chi \end{aligned}$$

The case for conjunction follows similarly.

2. Assume $\mathsf{M} \otimes \mathsf{E}^k, (s, f_1, ..., f_k) \vDash K_i\psi$. Then for all $(s' f'_1, ..., f'_k) \sim_i (s, f_1, ..., f_k)$ we have $\mathsf{M} \otimes \mathsf{E}^k, (s', f'_1, ..., f'_k) \vDash \psi$. Now let $(s', e'_1, ..., e'_m, f'_1, ..., f'_k) \in S\otimes\mathsf{E}^{m+k}$ such that $(s', e'_1, ..., e'_m, f'_1, ..., f'_k) \sim_i (s, e_1, ..., e_m, f_1, ..., f_k)$ be arbitrary. By Remark 29, ($\diamond$), we can write this without loss of information as $s' \sim_i^{\mathsf{M}\otimes\mathsf{E}_{e_1}\otimes...\otimes\mathsf{E}_{e_m}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_k}} s$. Now, by ($\diamond$) and Lemma 30, $\sim_i^{\mathsf{M}\otimes\mathsf{E}_{e_1}\otimes...\otimes\mathsf{E}_{e_m}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_k}} \subseteq \sim_i^{\mathsf{M}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_k}}$. This proves that $s' \sim_i^{\mathsf{M}\otimes\mathsf{E}_{f_1}\otimes...\otimes\mathsf{E}_{f_k}} s$. Again by Remark 29, we can equivalently write this as $(s', f''_1, ..., f''_k) \sim_i^{\mathsf{M}\otimes\mathsf{E}^k} (s, f_1, ..., f_k)$ for the unique sequence of events $f''_1, ..., f''_k$ that can be paired with s' in $\mathsf{M} \otimes \mathsf{E}_{f_1} \otimes ... \otimes \mathsf{E}_{f_k}$. Then $\mathsf{M} \otimes \mathsf{E}^k, (s', f''_1, ..., f''_k) \vDash \psi$. Now from Part 1, for all $e'_1, ...e'_m$ such that $(s', e'_1, ..., e'_m, f''_1, ..., f''_k) \in S \otimes \mathsf{E}^{k+m}$, we have $\mathsf{M} \otimes \mathsf{E}^{k+m}, (s', e'_1, ..., e'_m, f''_1, ..., f''_k) \vDash \psi$. But now again by Remark 29, we must have that $f''_1 = f'_1, ..., f''_k = f'_k$. This proves $\mathsf{M} \otimes \mathsf{E}^{k+m}, (s', e'_1, ..., e'_m, f'_1, ..., f'_k) \vDash \psi$. Therefore, $\mathsf{M} \otimes \mathsf{E}^{k+m}$, $(s, e_1, ..., e_m, f_1, ..., f_k) \vDash K_i\psi$, as desired.

□

The next lemma is our main lemma of interest, and will be important later on! It states that if we have two sequences of events not containing failed observation attempts, we can concatenate them into one.

Lemma 32 (Concatenation of Events)**.** *Let* $\mathsf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$ *be a multi-agent epistemic model and* $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ *be a basic event model, and let* $s \in S$ *be a state. Let* $m, k \in \omega$ *and* $e_1, , ..., e_m, f_1, ..., f_k \in \mathsf{E}$ *be such that none of* $e_1, , ..., e_m, f_1, ..., f_k$ *are failed observation attempts. If* $(s, e_1, ..., e_m) \in S \otimes \mathsf{E}^m$ *and* $(s, f_1, ..., f_k) \in S \otimes \mathsf{E}^k$*, then* $(s, e_1, , ..., e_m, f_1, ..., f_k) \in S \otimes \mathsf{E}^{m+k}$*.*

Proof. Assume $(s, e_1, ..., e_m) \in S \otimes \mathsf{E}^m$ and $(s, f_1, ..., f_k) \in S \otimes \mathsf{E}^k$. We prove this by induction on k. For the base case, let $k = 0$. Then since $(s, e_1, ..., e_m) \in S \otimes \mathsf{E}^m$, we are done. Now let $k = j + 1$. Then by the induction hypothesis, $(s, e_1, ..., e_m, f_1, ..., f_j) \in S \otimes \mathsf{E}^{m+j}$. Since we assume that $(s, f_1, ..., f_{j+1}) \in S \otimes \mathsf{E}^{j+1}$, we have $\mathsf{M} \otimes \mathsf{E}^j, (s, f_1, ..., f_j) \vDash pre_{f_{j+1}}$. Now note that $pre_{f_{j+1}}$ is either $\top$, l, or $l \wedge K_i\psi$ (for some literal l, agent i, and Boolean formula ψ). If $pre_{f_{j+1}}$ is Boolean, then by repeated applications of Part 1 of Preservation of Preconditions, we get $\mathsf{M} \otimes \mathsf{E}^{m+j}, (s, e_1, ..., e_m, f_1, ..., f_j) \vDash pre_{f_{j+1}}$. If $pre_{f_{j+1}} = l \wedge K_i\psi$, then by repeated applications of Part 1 of Preservation of Preconditions, we get $\mathsf{M} \otimes \mathsf{E}^{m+j}, (s, e_1, ..., e_m, f_1, ..., f_j) \vDash l$, and by Part 2 of Preservation of Preconditions, we get $\mathsf{M} \otimes \mathsf{E}^{m+j}, (s, e_1, ..., e_m, f_1, ..., f_j) \vDash K_i\psi$. Therefore, $\mathsf{M} \otimes \mathsf{E}^{m+j}, (s, e_1, ..., e_m, f_1, ..., f_j) \vDash pre_{f_{j+1}}$. This proves $(s, e_1, , ..., e_m, f_1, ..., f_{j+1}) \in S \otimes \mathsf{E}^{m+j+1}$. Therefore, by induction, $(s, e_1, , ..., e_m, f_1, ..., f_k) \in S \otimes \mathsf{E}^{m+k}$ for all $k \in \omega$. □

3.2 Syntax and semantics

Now that we have introduced our event models, we are in a position to define our syntax and semantics. Before defining our syntax for the main language, we also define some important meta-syntactic notions that help us to formally discuss the capabilities of different agents and subgroups. After defining our semantics, we also discuss some aspects of our semantics, explain some choices we made, and show how our semantics could also accommodate other understandings of group knowledge.

3.2.1 Syntax

The language for events $\mathcal{L}_e$ is given recursively as follows (where $G \subseteq N$, $i, j, k \in N$, l is a literal, p is a propositional variable, and ψ is a boolean formula):

$$e := (i :_G l) \mid (\psi/j : l) \mid (\psi/j : -_p) \mid !(j : k) \mid (e; f)$$

Where $(e; f)$ is the composition of events e and f (in that order). Note that each item in the event syntax is one of the permissible events we had defined in Subsection 3.1.1, or a composition of such events.

Next, we define some meta-syntactic notions involving our events. Firstly, recall that we had already defined the *active subgroup* for each of our permissible events e (i.e. events that belong to some permissible event set E). Now, given an event $e \in \mathcal{L}_e$, we define the set $\mathcal{A}_e$. Intuitively, this set matches the active subgroup of e when viewed as an event in some event model. However, it is different in the sense that the e in $\mathcal{A}_e$ comes from the syntax.

Definition 33 (Active subgroups)**.** *We define* $\mathcal{A}_e$ *inductively as follows:*

1. $\mathcal{A}_{(i:_G l)} = G \cup \{i\}$
2. $\mathcal{A}_{(\psi/j:l)} = \{j\}$
3. $\mathcal{A}_{(\psi/j:-_p)} = \{j\}$
4. $\mathcal{A}_{!(j:k)} = \{j, k\}$
5. $\mathcal{A}_{(e;f)} = \mathcal{A}_e \cup \mathcal{A}_f$

Since we are modelling collaboration and what can be achieved by a given group, we also want a way to speak of the capabilities a given subgroup has. In the following definition, we define for each subgroup $G \subseteq N$ a set $\mathcal{A}_G$ that contains the syntactic labels for the actions that members of G are able to take on their own (i.e. without help from other members of N).

Definition 34 (Capabilities of subgroup)**.** *Let $G \subseteq N$ be a non-empty subgroup. The* capabilities of subgroup *G, is the following set:*

$$\mathcal{A}_G = \{e \in \mathcal{L}_e \mid G \supseteq \mathcal{A}_e\}$$

The following proposition shows that $\mathcal{A}_G$ satisfies some intuitive properties.

Proposition 35. *Let M be an epistemic model, and E be an event model in which every generated submodel is a basic event model. The following hold:*

1. *For all $e \in \mathsf{E}^*$, we have $e \in \mathcal{A}_G$ for some subgroup G.*
2. *For all $G, G' \subseteq N$, if $G \subseteq G'$ then $\mathcal{A}_G \subseteq \mathcal{A}_{G'}$.*
3. *For all $e, e' \in \mathcal{L}_e$, $(e; e') \in \mathcal{A}_G$ if and only if $e \in \mathcal{A}_G$ and $e' \in \mathcal{A}_G$.*

Proof. 1. For any event $e \in \mathsf{E}^*$, the set $\mathcal{A}_e$ was defined to be a non-empty subset of N. Therefore, all $e \in \mathsf{E}^*$ satisfy $e \in \mathcal{A}_N$ (and $N \subseteq N$).

2. Let $G, G' \subseteq N$ be arbitrary such that $G \subseteq G'$. Assume $e \in \mathcal{A}_G$. Then by definition, $G \supseteq \mathcal{A}_e$. Since $G' \supseteq G$, we have that $G' \supseteq \mathcal{A}_e$. This proves that $e \in \mathcal{A}_{G'}$. Therefore, $\mathcal{A}_G \subseteq \mathcal{A}_{G'}$, as desired.

3. Let $e, e' \in \mathcal{L}_e$ be events. Then:

$$\begin{aligned}
(e; e') \in \mathcal{A}_G \quad &\iff G \supseteq \mathcal{A}_{(e;e')} = \mathcal{A}_e \cup \mathcal{A}_{e'} \\
&\iff G \supseteq \mathcal{A}_e, \mathcal{A}_{e'} \\
&\iff e \in \mathcal{A}_G \text{ and } e' \in \mathcal{A}_G
\end{aligned}$$

□

Now, recall that in Subsection 3.1.2, we showed Event Concatenation, but that this did not apply to failed observation attempts. We therefore define the set $\mathcal{A}_G^{-fobs}$ to intuitively include all the labels for actions G is capable of performing without failed observation attempts.

To define this set, consider the restriction of our event language to exclude failed observations $\mathcal{L}_e^{-fobs}$ defined recursively as follows (where $G \subseteq N$, $i \in G$, $j, k \in N$, l is a literal, and ψ is a boolean formula):

$$e := (i :_G l) \mid (\psi/j : l) \mid !(j : k) \mid (e; f)$$

Definition 36 ($\mathcal{A}_G^{-fobs}$).

$$\mathcal{A}_G^{-fobs} := \mathcal{A}_G \cap \mathcal{L}_e^{-fobs}$$

This definition will be used in our semantics.

We had previously defined both preconditions and access maps for each of our permissible events. We now also define preconditions and access maps as meta-syntactic notions for each of our events in $\mathcal{L}_e$. These are intended to align with how we defined preconditions and access maps previously for events in our models. We also define them as meta-syntactic notions because this allows us use them in our axiomatisation. More broadly, doing this allows us to speak of preconditions and access maps without depending on an event model.

We start with preconditions. For the composite case, we define a function $p_e : \mathcal{L}_D \to \mathcal{L}_D$ as follows (for l a literal and ψ a Boolean formula):

1. $p_e(\top) = \top$
2. $p_e(p) = p$
3. $p_e(\neg\varphi) = \neg p_e(\varphi)$
4. $p_e(\varphi \wedge \chi) = p_e(\varphi) \wedge p_e(\chi)$
5. $p_e(D_H\varphi) = D_{e(H)}(\bigvee_{e' \sim_H e} pre_{e'}) \to p_e(\varphi)$

We can now define the precondition for each event in our syntax.

Definition 37. *We define pre_e inductively as follows:*

1. $pre_{(i:_G l)} = l$
2. $pre_{(\psi/j:l)} = l \wedge K_j\psi$
3. $pre_{(\psi/j:-_p)} = \neg K_j\psi$
4. $pre_{!(j:k)} = \top$
5. $pre_{(e;f)} = pre_e \wedge p_e(pre_f)$

Now, recall from the background on data-exchange logic, the precondition $pre_{(e;f)}$ for a composite event $(e; f)$ is defined as $\langle e\rangle pre_f$, which in that context had the semantics (given an epistemic model $\mathbf{M}$, a state $s \in \mathbf{M}$, and a data-exchange model $\mathbf{E}$):

$$\mathbf{M}, s \vDash \langle e\rangle\varphi \iff \mathbf{M}, s \vDash pre_e \text{ and } \mathbf{M} \otimes \mathbf{E} \otimes ... \otimes \mathbf{E}, (s, e) \vDash \varphi$$

We prove that our definition of $pre_{(e;f)}$ matches this.

Proposition 38. *Let $\mathbf{M}$, a state $s \in \mathbf{M}$, and a data-exchange model $\mathbf{E}$. In addition, let $e, f \in \mathbf{E}$ be events. Then:*

$$\mathbf{M}, s \vDash pre_{(e;f)}\varphi \iff \mathbf{M}, s \vDash pre_e \textit{ and } \mathbf{M} \otimes \mathbf{E} \otimes ... \otimes \mathbf{E}, (s, e) \vDash pre_f$$

Proof. We prove this by strong induction on the number of permissible events that make up $e = (e_1, ..., e_r)$ and $f = (f_1; ...; f_m)$. Our induction hypothesis is then, for all epistemic models $\mathbf{M}$ and all event models $\mathbf{E}$, if $r + m < k$, then:

$$\mathbf{M}, s \vDash pre_{(e;f)}\varphi \iff \mathbf{M}, s \vDash pre_e \text{ and } \mathbf{M} \otimes \mathbf{E}^r, (s, e) \vDash pre_f$$

Let $r + m = k$. We proceed by cases on the structure of f_m. Note that if f_m is a Boolean formula the statement follows trivially since the truth of Boolean formulas is unaffected by product updates. The most interesting case is if f is an informed observation. For the case of a successful informed observation, $f_m = (\psi/j : l)$. In this proof we use the abbreviation g for $(e; f_1; ...; f_{m-1})$. We then get:

$$\begin{aligned}
\mathbf{M}, s \vDash pre_{(e;f_1;...;f_m)}\varphi \quad &\iff \mathbf{M}, s \vDash pre_g \wedge p_g(pre_{f_m}) \\
&\iff \mathbf{M}, s \vDash pre_g \wedge p_g(l \wedge K_j\psi) \\
&\iff \mathbf{M}, s \vDash pre_g \text{ and } \mathbf{M}, s \vDash l \wedge p_g(K_j\psi) \\
&\iff \mathbf{M}, s \vDash pre_g \text{ and } \mathbf{M}, s \vDash l \wedge D_{g(j)}(\bigvee_{e' \sim_j e} pre_{e'}) \to p_g(\psi) \\
&\iff \mathbf{M}, s \vDash pre_g \text{ and } \mathbf{M} \otimes \mathbf{E}^{r+m-1}, (s, g) \vDash l \text{ and, if } \mathbf{M}, s \vDash D_{g(j)}(\bigvee_{e' \sim_j g} pre_{e'}), \\
&\qquad \text{then } \mathbf{M}, s \vDash p_g(\psi) \\
&\iff \mathbf{M}, s \vDash pre_g \text{ and } \mathbf{M} \otimes \mathbf{E}^{r+m-1}, (s, g) \vDash l \text{ and, if for all } s' \sim_{g(j)} s, \text{ there is a} \\
&\qquad e' \sim_j g \text{ such that } \mathbf{M}, s \vDash pre_{e'}, \text{ then } \mathbf{M} \otimes \mathbf{E}^{r+m-1}, (s, g) \vDash \psi \\
&\iff \mathbf{M}, s \vDash pre_g \text{ and } \mathbf{M} \otimes \mathbf{E}^{r+m-1}, (s, g) \vDash l \text{ and for all } (s', e') \sim_j (s, g) \\
&\qquad \text{we have } \mathbf{M} \otimes \mathbf{E}^{r+m-1}, (s, g) \vDash \psi \\
&\iff \mathbf{M}, s \vDash pre_g \text{ and } \mathbf{M} \otimes \mathbf{E}^{r+m-1}, (s, g) \vDash l \wedge K_j\psi \\
&\iff \mathbf{M}, s \vDash pre_g \text{ and } \mathbf{M} \otimes \mathbf{E}^{r+m-1}, (s, g) \vDash pre_{f_m} \\
&\overset{IH}{\iff} \mathbf{M}, s \vDash pre_e \text{ and } \mathbf{M} \otimes \mathbf{E}^r, (s, e) \vDash pre_{(f_1;...;f_{m-1})} \text{ and} \\
&\qquad (\mathbf{M} \otimes \mathbf{E}^r) \otimes \mathbf{E}^{m-1}, ((s, e), (f_1; ...; f_{m-1})) \vDash pre_{f_m} \\
&\overset{IH}{\iff} \mathbf{M}, s \vDash pre_e \text{ and } \mathbf{M} \otimes \mathbf{E}^r, (s, e) \vDash pre_{(f_1;...;f_{m-1};f_m)} \\
&\iff \mathbf{M}, s \vDash pre_e \text{ and } \mathbf{M} \otimes \mathbf{E}^r, (s, e) \vDash pre_f
\end{aligned}$$

The case for $f_k = \neg K_j\psi$ follows similarly. □

We now move on to access maps.

Definition 39. *We define $e(\bullet)$ inductively as follows:*

1. $(i :_G l)(j) = \{j\}$ *for all* $j \in N$.
2. $(\psi/i : l)(j) = \{j\}$ *for all* $j \in N$.
3. $(\psi/i : -_p)(j) = \{j\}$ *for all* $j \in N$.
4. $!(j : k)(j) = \{j, k\}$, *and* $!(j : k)(i) = \{i\}$ *for all* $i \neq j \in N$.
5. $(e; f)(i) = \bigcup_{j \in f(i)} e(j)$.

Note that the access map in the composite case is defined to match the definition of the access map of composite events given in the background.

The main logical language $\mathcal{L}_{\mathcal{K}_H^G}$ is defined recursively as follows (where e is given by the above event syntax):

$$\varphi := p \mid cap(e) \mid \varphi \wedge \psi \mid \neg\varphi \mid D_G\varphi \mid [e]\varphi \mid \mathcal{K}_H^G\varphi$$

$$p \in \mathcal{P}, i \in N, H, G \subseteq N$$

Here, $cap(e)$ is interpreted as "the group N is capable of doing e" (and can be viewed as a propositional variable with additional structure), and $\mathcal{K}_H^G\varphi$ is interpreted as the subgroup H, with the collaboration of the subgroup G, has the potential to come to distributively know φ (after executing some actions). We sometimes refer to this as the potential collaborative distributive knowledge operator (i.e. has the potential to become distributive knowledge via collaboration with agents).

Since each event has an active subgroup, $cap(e)$ can be equivalently viewed as a capability of any superset of its active subgroup $\mathcal{A}_e$ (rather than just the whole group N). We then define the following meta-syntactic notation:
Let e be an event and $G \supseteq \mathcal{A}_e$. Then $cap_G(e)$ can be equivalently written instead of $cap(e)$.

In this thesis, our focus will be on the $\mathcal{K}_H^G$ modality, but we also need the distributive knowledge modality D_G for each $G \subseteq N$ for our proposed axiomatisation. Note here that the actual knowledge modality K_i for each $i \in K$ can be defined as an abbreviation of a special case of the distributed knowledge modality as follows:

$$K_i\varphi := D_{\{i\}}\varphi$$

This is important because we use the actual knowledge modality in the preconditions of events. We then have that K_i (actual knowledge of i) and $\mathcal{K}_{\{i\}}^{\{i\}}$ (what i can learn alone via actions) are different, since it is possible for an agent to be able to come to know things that they do not already know by executing actions.

In addition to these basic operators, we define $\vee$, $\rightarrow$, and $\leftrightarrow$ the usual way. We also define a few other abbreviations involving the $\mathcal{K}_H^G$ operator:

Definition 40 (Abbreviations $\mathcal{K}_i^G\varphi$ and $\mathcal{K}_G\varphi$)**.** *Let $i \in N$ be an agent and $G \subseteq N$ be a subgroup.*

1. $\mathcal{K}_i^G\varphi := \mathcal{K}_{\{i\}}^G\varphi$.
2. $\mathcal{K}_G\varphi := \bigwedge_{i \in G} \mathcal{K}_i^G\varphi$.

We read $\mathcal{K}_G\varphi$ as "the subgroup G has potential collaborative knowledge of φ". Recall in Section 2.1, we had two ways of viewing this operator. The first was as a notion of group knowledge, where the group knows a statement when they can collaborate so that all members learn it. The second was as a notion of group knowability, where a statement is knowable when, by collaborating, the subgroup G can come to know φ. Recall that the notion of group knowledge used here is that the group knows a statement if and only if every agent in the group knows it. In this thesis, we take $\mathcal{K}_G$ to be the main notion of philosophical interest. We work out an example of a scenario using this notion, and compare this notion to that of distributive knowledge. The reason we have made $\mathcal{K}_H^G$ our primitive modality is that it is needed for our proposed axiomatisation, and because its generality makes our framework flexible enough to also be able to model alternative definitions of collaborative group knowledge, as explained in Subsection 3.2.4. We explain the need for this operator in our axiomatisation in Section 4.4.

3.2.2 Semantics

Next, we introduce the semantics. For this, we first need to define the structures in which our formulas are evaluated. We call these "epistemic group setups". The definition of our structures makes use of the notion of permissible event models defined in Definition 26.

Definition 41 (Epistemic group setup)**.** *An* epistemic group setup $\mathbb{S}$ *is a tuple* (M, E)*, where* $\mathsf{M} = (S, (\sim_i)_{i\in N}, \| \bullet \|)$ *is a multi-agent epistemic model and* $\mathsf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre)$ *is a permissible event model.*

Remark 42. *Note that an* epistemic group setup $\mathbb{S}$ *can equivalently be seen as a tuple* $(S, (\sim_i)_{i\in N}, \| \bullet \|, \mathsf{E})$*, where* $\mathsf{M} = (S, (\sim_i)_{i\in N}, \| \bullet \|)$ *is an multi-agent epistemic model and* E *is a permissible event set. We can recover the whole event model just from* E *as explained in Remark 28.*

The semantics for the main language $\mathcal{L}_{\mathcal{K}^G_H}$ is given below. Recall that the notation $\mathbf{E}^m$ and E^m was explained in Remark 15.

Definition 43 (Semantics)**.** *Given an epistemic group setup* $\mathbb{S} = (\mathsf{M} = (S, (\sim_i)_{i\in N}, \| \bullet \|),\ \mathsf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$*, the satisfaction conditions are:*

$\mathbb{S}, s \vDash p$	iff	$s \in \|p\|$
$\mathbb{S}, s \vDash cap(e)$	iff	$e \in \mathsf{E}^*$
$\mathbb{S}, s \vDash \neg\varphi$	iff	$\mathbb{S}, s \nvDash \varphi$
$\mathbb{S}, s \vDash \varphi \wedge \psi$	iff	$\mathbb{S}, s \vDash \varphi$ *and* $\mathbb{S}, s \vDash \psi$
$\mathbb{S}, s \vDash D_G\varphi$	iff	*for all* $s' \in S$, *if* $s \sim_G s'$, *then* $\mathbb{S}, s' \vDash \varphi$
$\mathbb{S}, s \vDash [e]\varphi$	iff	*there exists* $m \in \omega$ *such that* $e \in \mathsf{E}^m$, *and, if* $(s, e) \in S \otimes \mathsf{E}^m$, *then* $(\mathsf{M} \otimes \mathbf{E}^m, \mathsf{E}), (s, e) \vDash \varphi$
$\mathbb{S}, s \vDash \mathcal{K}^G_H\varphi$	iff	*there exists* $m \in \omega$ *and* $e \in \mathcal{A}^{-fobs}_G$ *with* $(s, e) \in S \otimes \mathsf{E}^m$ *such that, for all* $s' \in S, e' \in \mathsf{E}^m$ *such that* $(s', e') \in S \otimes \mathsf{E}^m$, *if* $(s, e) \sim^{\mathsf{M}\otimes\mathbf{E}^m}_H (s', e')$, *then* $\mathbb{S}, s' \vDash \varphi$

The role of the set $\mathcal{A}^{-fobs}_G$ in the semantic clause for $\mathcal{K}^G_H\varphi$ is to specify that these events cannot be executed with the help of members outside of G. For something to be able to become distributive knowledge for H with the collaboration of G, there cannot be collaboration from members outside of G. The interpretation of the semantics for the $\mathcal{K}^G_H$ operator is then that there are actions, other than failed observations, executable by members in G where, after they are executed, it becomes distributive knowledge for H that φ used to be the case before the events were executed. We exclude failed observations for technical reasons (so that we have Event Concatenation), but intuitively, it is natural that agents would not use information gathered from failed observation attempts to learn anything. Note that in the semantics, we allow $m = 0$. This is because if something is already distributive for H, it should also be potential collaborative distributive knowledge for H with respect to G.

Also note here that the semantics for the dynamic modality $[e]$ is different from the semantics for standard dynamic epistemic logic because we include $e \in \mathsf{E}^*$ in the antecedent of the semantic clause. This means that the semantics for the dual $\langle e\rangle$ (given an epistemic group setup $\mathbb{S}$ and state $s \in S$)

becomes:

$$\mathbb{S}, s \vDash \langle e \rangle \varphi \quad \textit{iff} \quad \text{there exists } m \in \omega \text{ such that } e \in \mathsf{E}^m \text{ and } (s,e) \in S \otimes \mathsf{E}^m \text{ and } (\mathbf{M} \otimes \mathbf{E}^m, \mathbf{E}), (s,e) \vDash \varphi$$

This is similar to the semantic clause found in protocol dynamic epistemic logic.

The semantics for $\mathcal{K}_H^G \varphi$ given above can be equivalently restated by shifting the scope of the quantifiers. The version we presented is the logically simpler version and most clearly connected to the intuition for the semantics, but when working with the semantics in proofs, it is often easier to think of the other version. To explain the two versions, let $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i \in N}, \| \bullet \|), \mathbf{E} = (\mathsf{E}, (\sim_i)_{i \in N}, \bullet(\bullet), pre))$ be an epistemic group setup, and $s \in S$ a state. First note that the semantics for $\mathcal{K}_H^G$ written in first order logic is:

$$\mathbb{S}, s \vDash \mathcal{K}_H^G \varphi \quad \textit{iff} \quad (\exists m \in \omega)(\exists e \in \mathcal{A}_G^{-fobs}) \Big[(s,e) \in S \otimes \mathsf{E}^m \wedge \Big((\forall s' \in S) (\forall e' \in \mathsf{E}^m) [((s',e') \in S \otimes \mathsf{E}^m \wedge (s,e) \sim_H^{\mathbf{M} \otimes \mathbf{E}^m} (s',e')) \to \mathbb{S}, s' \vDash \varphi] \Big) \Big]$$

Now, the more workable logically equivalent version can be obtained by changing the scope of the quantifiers for the second composite event to only the antecedent of the implication:

$$\mathbb{S}, s \vDash \mathcal{K}_H^G \varphi \quad \textit{iff} \quad (\exists m \in \omega)(\exists e \in \mathcal{A}_G^{-fobs}) \Big[(s,e) \in S \otimes \mathsf{E}^m \wedge \Big((\forall s' \in S) [(\exists e' \in \mathsf{E}^m)((s',e') \in S \otimes \mathsf{E}^m \wedge (s,e) \sim_H^{\mathbf{M} \otimes \mathbf{E}^m} (s',e'))] \to \mathbb{S}, s' \vDash \varphi \Big) \Big]$$

Note here that $(\exists e' \in \mathsf{E}^m)((s',e') \in S \otimes \mathsf{E}^m \wedge (s,e) \sim_H^{\mathbf{M} \otimes \mathbf{E}^m} (s',e'))$ is a condition specifying which states $s' \in S$ we need to consider given a composite event e. We can also, less formally, state this condition as: there is a composite event $e' \in \mathsf{E}$ such that $(s',e') \in S \otimes \mathsf{E}^m$ and $(s,e) \sim_H^{\mathbf{M} \otimes \mathbf{E}^m} (s',e')$. When using this version of the semantics, we can therefore think of narrowing down the class of states we need to consider for each agent first. To make this simpler and more explicit, we define the following set:

Definition 44. *Let $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i \in N}, \| \bullet \|), \mathbf{E} = (\mathsf{E}, (\sim_i)_{i \in N}, \bullet(\bullet), pre))$ be an epistemic group setup, $G \subseteq S$ a subgroup, and $s \in S$ a state. Let $m \in \omega$ and $e \in \mathcal{A}_G^{-fobs}$ be such that $(s,e) \in S \otimes \mathsf{E}^m$. Then define:*

$$S_{s,e}^{\mathbb{S},H} = \{ s' \in S \mid \exists e' \in \mathsf{E}^m \textit{ s.t. } (s',e') \in S \otimes \mathsf{E}^m \textit{ and } (s,e) \sim_H^{\mathbf{M} \otimes \mathbf{E}^m} (s',e') \}$$

This is interpreted as the set of states that remain possible for the members of H after combining their knowledge, after the composite event e has been executed. Using this new notation, the semantics for the $\mathcal{K}_H^G$ operator can be equivalently stated as:

$$\mathbb{S}, s \vDash \mathcal{K}_H^G \varphi \quad \textit{iff} \quad \text{there exists } m \in \omega \text{ and } e \in \mathcal{A}_G^{-fobs} \text{ with } (s,e) \in S \otimes \mathsf{E}^m$$

such that, for all $s' \in S^{\mathbb{S},H}_{s,e}$ we have $\mathbb{S}, s' \vDash \varphi$

3.2.3 Removing dynamic modalities

Our main semantics includes the dynamic modalities $[e]$, for each event e in our epistemic group setup. This is beneficial because it provides an intuitive semantics for collaboration processes. However, we show that it is possible to remove this dynamic modality from our language and semantics, as well as all mention of the full event models. This is helpful because it means that reduction laws are not a necessary component of all axiomatisations involving our main operator, $\mathcal{K}^G_H$. For the remainder of this subsection, we assume an epistemic group setup $\mathbb{S} = (S, (\sim_i)_{i\in N}, \|.\|, \mathsf{E})$ is given.

To show how we can remove dynamic modalities and references to event models, we first derive another equivalent statement of our semantics for $\mathcal{K}^G_H\varphi$ which does not involve taking product updates. For the derivation, first observe that $(s, e) \in S \otimes \mathsf{E}^m$ can be equivalently stated as $\mathbb{S}, s \vDash pre_e$, and that $(s, e) \sim_H (s', e')$ can be equivalently stated as $s' \sim_{e(H)} s$ and $e \sim_H e'$. This semantics is:

$$\mathbb{S}, s \vDash \mathcal{K}^G_H\varphi \quad \textit{iff} \quad \begin{array}{l}\text{there exists } e \in \mathcal{A}^{-fobs}_G \text{ such that } \mathbb{S}, s \vDash pre_e \text{ and, for all } s' \sim_{e(H)} s, e' \sim_H e,\\ \text{if } \mathbb{S}, s' \vDash pre_{e'}, \text{ then } \mathbb{S}, s' \vDash \varphi\end{array}$$

Next, observe that by again shifting the scope of the for all quantifier to the antecedent of the implication, this becomes:

$$\mathbb{S}, s \vDash \mathcal{K}^G_H\varphi \quad \textit{iff} \quad \begin{array}{l}\text{there exists } e \in \mathcal{A}^{-fobs}_G \text{ such that } \mathbb{S}, s \vDash pre_e \text{ and, for all } s' \sim_{e(H)} s,\\ \text{satisfying that there exists } e' \sim_H e \text{ such that } \mathbb{S}, s' \vDash pre_{e'}, \text{ we have } \mathbb{S}, s' \vDash \varphi\end{array}$$

Finally, observe that we can then simplify the existential clause into a disjunction as follows:

$$\mathbb{S}, s \vDash \mathcal{K}^G_H\varphi \quad \textit{iff} \quad \begin{array}{l}\text{there exists } e \in \mathcal{A}^{-fobs}_G \text{ such that } \mathbb{S}, s \vDash pre_e \text{ and, for all } s' \sim_{e(H)} s,\\ \mathbb{S}, s' \vDash (\bigvee_{e' \sim_H e} pre_{e'}) \to \varphi\end{array} \qquad \text{(static semantics)}$$

Now, note that this semantics does not make use of our dynamic modality, $[e]$, or event models. It only uses an event set, preconditions and access maps, which we had defined in the Syntax subsection as a meta-syntactic notion, boolean operators, and the indistinguishability relations of our epistemic model.

It is therefore possible to remove our dynamic modality $[e]$ from the language, and to only consider the static logic $\mathcal{L}^{static}_{\mathcal{K}^G_H}$ defined recursively as follows (where $e \in \mathcal{L}_e$ is an event in the event syntax):

$$\varphi := p \mid cap(e) \mid \varphi \wedge \psi \mid \neg\varphi \mid D_G\varphi \mid \mathcal{K}^G_H\varphi$$

$$p \in \mathcal{P}, i \in N, H, G \subseteq N$$

Then semantics for $\mathcal{K}^G_H$ is defined by the semantic clause derived in this subsection, and the semantics for all other operators are defined as in Subsection 3.2.2.

The following remark is important to make sense of our methodology in Chapter 4.

Remark 45. *Note that even when working with our logic over $\mathcal{L}^{static}_{\mathcal{K}^G_H}$, we can still make use of event models in semantic proofs. This is because we have all of the following:*

1. *A permissible event set fully determines a permissible event model, as explained in Remark 28.*
2. *All of our meta-syntactic notions were defined to match the semantic notions, as pointed out in Subsection 3.2.1.*
3. *The semantics of $\mathcal{K}_H^G\varphi$ in $\mathcal{L}_{\mathcal{K}_H^G}^{static}$ is equivalent to its semantics in $\mathcal{L}_{\mathcal{K}_H^G}$, as proven in this subsection.*

3.2.4 Discussion on semantics and semantics for $\mathcal{K}_G$

In this subsection, we explain our choice of semantics for $\mathcal{K}_H^G$ and provide and discuss the semantics for $\mathcal{K}_G$ defined in Definition 40. We also discuss some possible alternative semantics for $\mathcal{K}_G$. Throughout the subsection, we assume we are working in an epistemic group setup $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \|\bullet\|), \mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$.

One natural question about our choice of semantics is why we specify that the members of the subgroup should come to distributively know φ *was the case before executing the events.* Without this specification, the semantics would just become:

$$\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi \quad \textit{iff} \quad \text{there exist } m \in \omega, e \in \mathcal{A}_G^{-fobs} \text{ with } (s,e) \in S \otimes \mathsf{E}^m \text{ such that } (\mathbf{M} \otimes \mathbf{E}^m, \mathbf{E}), (s,e) \vDash D_H\varphi$$

Spelling this out with the meaning of the distributive knowledge operator and writing this in first-order logic, we get:

$$\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi \quad \textit{iff} \quad (\exists m \in \omega)(\exists e \in \mathcal{A}_G^{-fobs})\Big[(s,e) \in S \otimes \mathsf{E}^m \wedge \Big((\forall s' \in S)(\forall e' \in \mathsf{E}^m)[((s',e') \in S \otimes \mathsf{E}^m \wedge (s,e) \sim_H^{\mathbf{M}\otimes\mathbf{E}^m} (s',e')) \rightarrow (\mathbf{M} \otimes \mathbf{E}^m, \mathbf{E}), (s',e') \vDash \varphi]\Big)\Big]$$

The problem with this semantics is that it breaks down for Moore sentences like $\varphi \wedge \neg K_i\varphi$. It would then be impossible for any agent to learn this sentence or others like it, because once it is learnt, it becomes false. An agent can only learn that it used to be true. For this reason, our semantics considers what agents learn to be true before the events occur by evaluating truth in the original model. In first-order logic, making this change gives us:

$$\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi \quad \textit{iff} \quad (\exists m \in \omega)(\exists e \in \mathcal{A}_G^{-fobs}))\Big[(s,e) \in S \otimes \mathsf{E}^m \wedge \Big((\forall s' \in S)(\forall e' \in \mathsf{E}^m)[((s',e') \in S \otimes \mathsf{E}^m \wedge (s,e) \sim_H^{\mathbf{M}\otimes\mathbf{E}^m} (s',e')) \rightarrow \mathbb{S}, s' \vDash \varphi]\Big)\Big]$$

This is the semantics we gave above.

We now move on to the semantics for $\mathcal{K}_G$. Recall that we had defined $\mathcal{K}_G\varphi$ as $\bigwedge_{i\in G}\mathcal{K}_i^G\varphi$, where $\mathcal{K}_i^G\varphi$ was defined as $\mathcal{K}_{\{i\}}^G\varphi$. However, at face value, when we interpret potential collaborative knowledge as a notion of group knowability, this makes the potential collaborative knowledge of the group G such that the chain of events needed to actualise the knowledge in individuals is agent-dependent. In other words, φ is potential collaborative knowledge for G if each agent had some sequence of events they could execute to learn φ. This is not quite in line with the intuitive idea we wanted to model, when viewing potential collaborative knowledge as a notion of group knowability. For this, we would want there to be a potential scenario in which the whole group comes to know that φ was true, rather than several alternative scenarios where combined the agents each come to know φ. To us, potential indicates that there *exists a possible future* where it becomes actualised (recall that we have defined actual group knowledge as "everyone knows"). The semantics for this would be defined as follows: Given a set of agents N, a subgroup $G \subseteq N$, an epistemic group setup $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \| \bullet \|), \mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$, and a state $s \in S$:

$$\mathbb{S}, s \vDash \mathcal{K}^*_G\varphi \quad \textit{iff} \quad \begin{array}{l} \text{there exists } m \in \omega \text{ and } e \in \mathcal{A}_G^{-fobs} \text{ with } (s,e) \in S \otimes \mathsf{E}^m \text{ such that,} \\ \text{for all } i \in G, s' \in S, e' \in \mathsf{E}^m \text{ such that } (s',e') \in S \otimes \mathsf{E}^m, \\ \text{if } (s,e) \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s',e'), \text{ then } \mathbb{S}, s' \vDash \varphi \end{array}$$

Luckily for us, these are equivalent in epistemic group setups, as we defined them!

Theorem 46 (Equivalence of two definitions of group knowledge)**.** *Let* $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \| \bullet \|), \mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$ *be an epistemic group setup. Let* $s \in S$ *be a state and* $G = \{1, ..., n\} \subseteq N$ *a subgroup. Then for all formulas* φ*:*

$$\mathbb{S}, s \vDash \mathcal{K}^*_G\varphi \iff \mathbb{S}, s \vDash \mathcal{K}_G\varphi$$

Proof. $\Rightarrow$: This direction is trivial (if there is one possible future in which all agents come to know φ then there is a future for all agents).

$\Leftarrow$: Assume $\mathbb{S}, s \vDash \mathcal{K}_G\varphi$. Then for each $i \in G$ there exists $m \in \omega$ and $e^i = (e^i_1; ...; e^i_m) \in \mathcal{A}_G^{-fobs}$ with $(s, e^i) \in S \otimes \mathsf{E} \otimes ... \otimes \mathsf{E}$ (m_i times) such that, for all $s' \in S$ satisfying that there exists $e' \in \mathsf{E}$ with $(s', e') \in S \otimes \mathsf{E} \otimes ... \otimes \mathsf{E}$ and $(s', e') \sim_i (s, e^i)$, we have $\mathbb{S}, s' \vDash \varphi$.

Now note that $(s, e^1, ..., e^n) \in S \otimes \mathsf{E} \otimes ... \otimes \mathsf{E}$ by Concatenation of Events (Lemma 32).

Next, note that after executing all these events together, the set of s's we still need to consider shrink or stay the same. Formally, if there are $e^{1*} \in \mathsf{E}^{m_1}, ..., e^{n*} \in \mathsf{E}^{m_n}$ with $(s', e^{1*}, ..., e^{n*}) \in S \otimes \mathsf{E}^{m_1+...+m_n}$ and $(s', e^{1*}, ..., e^{n*}) \sim_i^{\mathbf{M}\otimes\mathbf{E}^{m_1+...+m_n}} (s, e^1, ..., e^n)$, then for any $i \in G$ there is a $e' \in \mathsf{E}^{m_i}$ such that $(s', e') \in S \otimes \mathsf{E}^{m_i}$ and $(s', e') \sim_i (s, e^i)$.

Therefore, if s' satisfies that there are $e^{1*}, ..., e^{n*} \in \mathsf{E}^*$ with $(s', e^{1*}, ..., e^{n*}) \in S \otimes \mathsf{E} \otimes ... \otimes \mathsf{E}$ and $(s', e^{1*}, ..., e^{n*}) \sim_i (s, e^1, ..., e^n)$, then $s' \in S^{\mathbb{S},\{i\}}_{s,e^i}$ for all $i \in G$. This proves $\mathbb{S}, s' \vDash \varphi$.

Therefore, we have $\mathbb{S}, s \vDash \mathcal{K}^*_G\varphi$, as desired.

□

The semantics of the $\mathcal{K}_G$ operator therefore succeeds in defining a notion of group knowability based on what has the potential to become known by the group (in the sense of each member knowing) through collaboration.

Next, we move on to showing how our formalism can accommodate other understandings of actual group knowledge. Here, we are still considering the case where we view potential group knowledge as a notion of group knowability. We consider two alternative understandings of actual group knowledge. This first is only requiring that one agent is able to learn φ for it to be potential collaborative knowledge. Essentially, after adopting this view, a statement would be potential collaborative knowledge if there is a sequence of actions after which it becomes known by someone. If we view group knowledge as consisting of all statements immediately accessible to the members of the group via communication, or as consisting of all statements that have the potential to affect the overall actions of the group, this version may be more appropriate. Fortunately, our formal framework can easily accommodate this perspective as well! To do so, we would need to define $\mathcal{K}_G\varphi$ as $\bigvee_{i\in G}\mathcal{K}_i^G\varphi$ instead of $\bigvee_{i\in G}\mathcal{K}_i^G\varphi$.

The second understanding is the one that results from considering actual group knowledge to be distributive knowledge. Then we would see potential group knowledge as consisting of all statements that have the potential to become distributive knowledge for the group. This can also be accommodated by our framework, this time by defining $\mathcal{K}_G\varphi$ as $\mathcal{K}_G^G\varphi$ instead of $\bigvee_{i\in G}\mathcal{K}_i^G\varphi$.

3.3 Example and link to philosophical motivations

The example in this section serves two purposes. One is to clarify the semantics presented above, and the other is to formally compare what can be achieved together by collaborating to what can be achieved individually.

> **Discovering a planet**: The example involves two agents, Klarise and Alexandru, who want to observe planets and stars with a telescope. Klarise has the skills required to be able to learn how to use the telescope using the available manual. However, she has bad eyesight and will not be able to see fainter/more distant planets or stars through the telescope. Alexandru has good eyesight, but is not able to use the available manual to learn how to use the telescope, because it is in a language he does not understand. They are also both able to communicate knowledge they have about how to use the telescope or what they see. We assume that there is a new planet that would be discovered by Alexandru if he looked through the telescope. At the start, both agents have no knowledge of how to use the telescope or whether or not there is a planet.

Note that in all examples, we make use of Properties of Generated Submodels of Event Models (Proposition 19). Using this allows us to not compute the whole product update when we are only interested in part of it.

Let t and $\neg t$ be propositions summarising the two possible ways the telescope could work. Also, let p be the proposition stating "there is a new planet". We assume that t is the true description of how the telescope works. Let a represent Alexandru and k Klarise.

We first define an epistemic group setup $\mathbb{S}$. The epistemic model in $\mathbb{S}$ is $\mathbf{M} = (S, \sim_a, \sim_k, \|\bullet\|)$, where $S = \{s_1, s_2, s_3, s_4\}$, $\sim_a = \sim_k = S\times S$, and $\|p\| = \{s_1, s_3\}$, and $\|t\| = \{s_1, s_2\}$. Basically, the worlds where there is a new planet are s_1 and s_3, and the worlds where t describes how the telescope works are s_1 and s_2.

$\mathbf{M} :=$

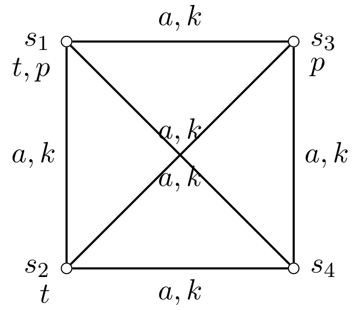

The event model $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i \in N}, \bullet(\bullet), pre)$ in $\mathbb{S}$ has four generated submodels. The first (representing Klarise learning how to use the telescope) is $\mathbf{E}_1 = (\mathsf{E}_1, \sim^1_a, \sim^1_k, \bullet(\bullet)_1 pre_1)$, where $\mathsf{E}_1 = \{(k :_{\{k\}} t), (k :_{\{k\}} \neg t)\}$, $\sim^1_a = \mathsf{E}_1 \times \mathsf{E}_1, \sim^1_k = \{((k :_{\{k\}} t), (k :_{\{k\}} t)),\ ((k :_{\{k\}} \neg t), (k :_{\{k\}} \neg t))\}$, the access map is trivial $(k :_{\{k\}} \neg t)(k) = (k :_{\{k\}} t)(k) = \{k\}$ and $(k :_{\{k\}} \neg t)(a) = (k :_{\{k\}} t)(a) = \{a\}$, $pre_1((k :_{\{k\}} t)) = t$, and $pre_1((k :_{\{k\}} \neg t)) = \neg t$.

$\mathbf{E}_1 :=$

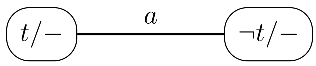

Figure 3.1: Event model with two events: one with precondition t and trivial access map, and another with precondition $\neg t$ and trivial access map.

Note that this is an example of a basic observation event model with $\{k\}$ as the active subgroup, because Klarise is the only one who knows the outcome of the event.

The second (representing Klarise sharing all she knows with Alexandru) is $\mathbf{E}_2 = (\mathsf{E}_2, \sim^2_a, \sim^2_k, \bullet(\bullet)_2, pre_2)$, $\mathsf{E}_2 = \{!(a : k)\}$, $\sim^2_k = \sim^2_a = \{(!(a : k), !(a : k)\}$, $!(a : k)(a) = \{a, k\}$, $!(a : k)(k) = \{k\}$ and $pre_2(!(a : k)) = \top$.

$\mathbf{E}_2 :=$

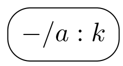

Figure 3.2: Event model with one event that has a trivial precondition and $a : k$ as its access map.

Note that this is a sharing event model in which the agent k is sharing all she knows with the agent a. The active subgroup is then $\{a, k\} = N$.

The third (representing Alexandru looking through the telescope) is $\mathbf{E}_3 = (\mathsf{E}_3, \sim^3_a, \sim^3_k, pre_3)$, where $\mathsf{E}_3 = \{(t/a : p), (t/a : \neg p), (t/a : -_p)\}$, $\sim^3_a = \{((t/a : p), (t/a : p)), ((t/a : \neg p), (t/a : \neg p)), ((t/a : -_p), (t/a : -_p))\}, \sim^3_k = \mathsf{E}_3 \times \mathsf{E}_3$, the access map is trivial (i.e. $e(i) = \{i\}$ for all $e \in \mathsf{E}_3$ and all $i \in N$), and $pre_3((t/a : p)) = K_a t \wedge p$, $pre_3((t/a : \neg p)) = K_a t \wedge \neg p$, and $pre_3((t/a : -_p)) = \neg K_a t$.

$\mathbf{E}_3 :=$

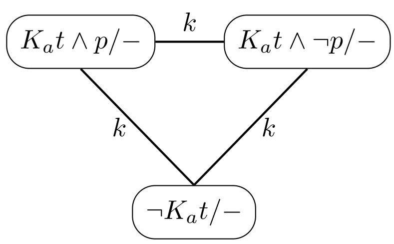

Figure 3.3: Event model with three events: one with precondition $K_a t \wedge p$, another with precondition $K_a t \wedge \neg p$ and the last with precondition $\neg K_a t$, and all with trivial access maps.

Note that this is an informed observation event model, this time with active subgroup $\{a\}$.

The fourth (representing Alexandru sharing all he knows with Klarise) is $\mathbf{E}_4 = (\mathsf{E}_4, \sim_a^4, \sim_k^4, \bullet(\bullet)_4, pre_4)$, $\mathsf{E}_4 = \{!(k : a)\}$, $\sim_k^4 = \sim_a^4 = \{(!(k : a), !(k : a)\}$, $!(k : a)(a) = \{a\}$, $!(k : a)(k) = \{a, k\}$ and $pre_4(!(k : a)) = \top$.

$\mathbf{E}_4 :=$

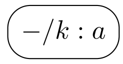

Figure 3.4: Event model with one event that has a trivial precondition and $k : a$ as its access map.

Note that this is again a sharing event model with active subgroup $\{a, k\} = N$.

Now, consider the real world s_1 alone with following sequence of events: $e_1 = (k :_{\{k\}} t), e_2 = !(a : k), e_3 = (t/a : p), e_4 = !(k : a)$. To evaluate which states are still considered possible for each agent after the events, we need to compute the product update of $\mathbf{M}$ with the generated submodels of e_1, e_2, e_3 and e_4 in order. For this the generated submodels are as follows $\mathbf{E}_{e_1} = \mathbf{E}_1$, $\mathbf{E}_{e_2} = \mathbf{E}_2$, $\mathbf{E}_{e_3} = \mathbf{E}_3$, and $\mathbf{E}_{e_4} = \mathbf{E}_4$.

The product update of $\mathbf{M}$ with the basic observation event model $\mathbf{E}_1$ is:

$\mathbf{M} \otimes \mathbf{E}_1 :=$

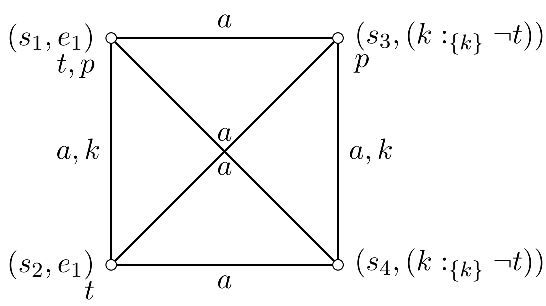

Here, Klarise learns the truth about how to use the telescope.

The product update of $\mathbf{M} \otimes \mathbf{E}_1$ with the sharing event model $\mathbf{E}_2$ is:

$\mathbf{M} \otimes \mathbf{E}_1 \otimes \mathbf{E}_2 :=$

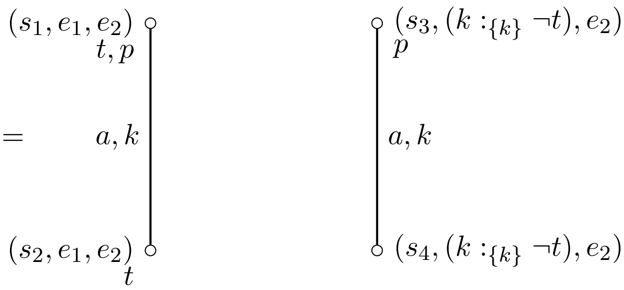

Here, Klarise shares all she knows with Alexandru (so he also learns the truth about how to use the telescope).

Then the product update of $\mathbf{M} \otimes \mathbf{E}_1 \otimes \mathbf{E}_2$ with the informed observation event model $\mathbf{E}_3$ is:

$\mathbf{M} \otimes \mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \mathbf{E}_3 :=$

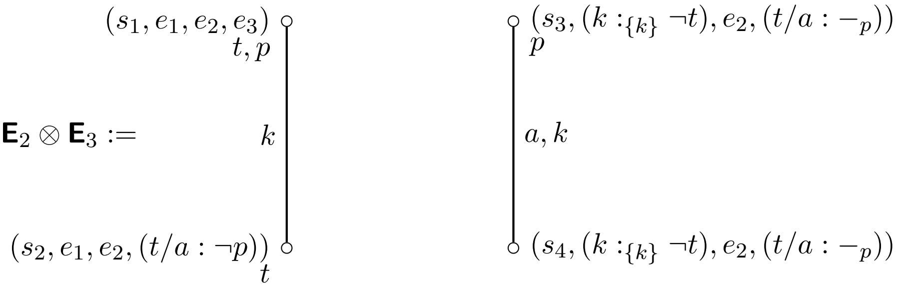

Here, Alexandru makes an observation allowing him to know the real world (in the true case where he knows t).

Lastly, the product update of $\mathbf{M} \otimes \mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \mathbf{E}_3$ with the sharing event model $\mathbf{E}_4$ is:

$\mathbf{M} \otimes \mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \mathbf{E}_3 \otimes \mathbf{E}_4 :=$

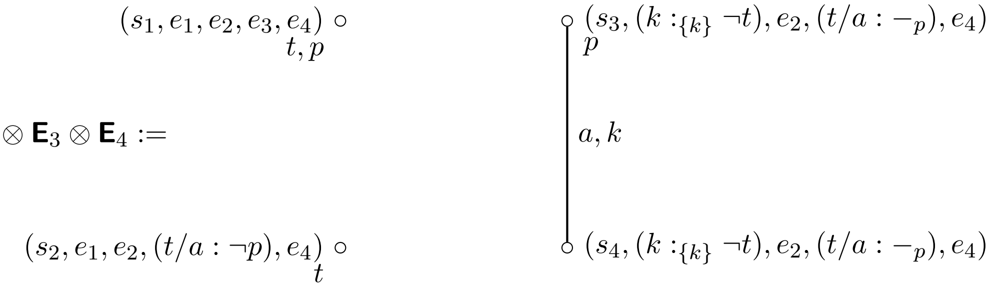

Here, Klarise also learns the real world (in the true case where we have t), after Alexandru shares all he knows.

Now, note that for Alexandru, all worlds are distinguishable from $(s_1, e_1, e_2, e_3, e_4)$ in the final product update. Hence, the only world $(s', e'_1, ..., e'_4)$ in the product update satisfying $(s', e'_1, ..., e'_4) \sim_a (s_1, e_1, e_2, e_3, e_4)$ is $(s_1, e_1, e_2, e_3, e_4)$ itself. Therefore, since $\mathbb{S}, s_1 \vDash p$, we can conclude $\mathbb{S}, s_1 \vDash \mathcal{K}_a^{\{a,k\}} p$.

Similarly, for Klarise, all worlds are distinguishable from $(s_1, e_1, e_2, e_3, e_4)$ in the final product update. So again, the only world $(s', e'_1, ..., e'_4)$ in the product update satisfying $(s', e'_1, ..., e'_4) \sim_a (s_1, e_1, e_2, e_3, e_4)$ is $(s_1, e_1, e_2, e_3, e_4)$ itself. Therefore, since $\mathbb{S}, s_1 \vDash p$, we can again conclude $\mathbb{S}, s_1 \vDash \mathcal{K}_k^{\{a,k\}} p$.

Then by the definition of $\mathcal{K}_G$, we can conclude $\mathbb{S}, s_1 \vDash \mathcal{K}_{\{a,k\}} p$.

However, we do not have $\mathbb{S}, s_1 \vDash \mathcal{K}_{\{a\}} p$ or $\mathbb{S}, s_1 \vDash \mathcal{K}_{\{k\}} p$. To see why $\mathbb{S}, s_1 \nvDash \mathcal{K}_{\{a\}} p$, note that $\mathcal{A}_a = \mathsf{E}_3^*$. This is because $\{a\}$ is not a superset of the active subgroup of any events other than those in E_3. Since Alexandru does not know t in any world in the original model, updating with $\mathbf{E}_3$ just leaves the same model, since at each world a failed observation attempt is performed.

Similarly, to see why $\mathbb{S}, s_1 \nvDash \mathcal{K}_{\{k\}} p$, note that $\mathcal{A}_k = \mathsf{E}_1^*$. This is because $\{k\}$ is not a superset of the active subgroup of any events other than those in E^1. Now, in the product update of $\mathbf{M}$ with $\mathbf{E}_1$ (shown above), we still had that Klarise could not distinguish between p and $\neg p$ worlds. If we perform this product update again to get $\mathbf{M} \otimes \mathbf{E}_1 \otimes \mathbf{E}_1$, the model is not changed from $\mathbf{M} \otimes \mathbf{E}_1$. This shows that Klarise cannot learn p alone. Note also that p is also not a logical consequence of what Klarise and Alexandru can each learn individually.

When we view collaborative group knowledge as a notion of group knowledge, this example shows that potential group knowledge is non-summative, since the group consisting of Alexandru and Klarise, $\{a, k\}$, has more potential group knowledge than the individual members have combined, and closed under logical consequence. When we view collaborative group knowledge as a notion of group knowability, this example demonstrates that the knowledge that can be obtained via collaboration in the group $\{a, k\}$ is more than the logical consequences of the combined knowledge that can be obtained by a or k individually. This therefore supports our claim that via collaboration, the group can come to know more than all the logical consequences of what the members can come to know individually. Either way, this example demonstrates the importance of paying attention to collaboration and social phenomena in epistemology.

3.4 Comparison of potential collaborative knowledge to distributed knowledge

In Section 2.1, we claimed that in the case where all agents can do is share all they know, and all agents can share all they know with all other agents, distributed knowledge and potential collaborative knowledge become the same. We now prove this formally.

Proposition 47. *Let $G = \{a_1, ..., a_k\} \subseteq N$ be a subgroup, $i \in G$ an agent, and E a permissible event set. Assume that $\mathcal{A}_G \cap \mathsf{E} = \{!(i:j) \mid i, j \in G\}$. Then for all epistemic group setups based on E, $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$, $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$, all states $s \in S$, and all formulas φ, we have:*

$$\mathbb{S}, s \vDash \mathcal{K}_G\varphi \iff \mathbb{S}, s \vDash D_G\varphi$$

Proof. For both directions, note that since the precondition for sharing events is always $\top$, the preconditions are always satisfied. In addition, note that events can only be indistinguishable for some agent if they have the same active subgroup. This means that if $e \sim_i f$ for any $i \in N$, $e \in \mathcal{A}_G^{-fobs}$ then $f \in \mathcal{A}_G^{-fobs}$. Given both these facts, the semantics for the $\mathcal{K}_i^G$ modality boils down to:

$$\mathbb{S}, s \vDash \mathcal{K}_i^G\varphi \quad \textit{iff} \quad \text{there exists } m \in \omega \text{ and } e \in \mathcal{A}_G^{-fobs} \text{ such that, for all } s' \in S, e' \in \mathsf{E}^m \text{ if } (s,e) \sim_i (s',e'), \text{ then } \mathbb{S}, s' \vDash \varphi$$

- $\Rightarrow$: Assume $\mathbb{S}, s \vDash \mathcal{K}_G\varphi$. Take any $i \in G$. There is a composite event $e \in \mathcal{A}_G^{-fobs} \cap \mathsf{E}^m$ such that for all $(s', e') \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s, e)$, we have $\mathbb{S}, s' \vDash \varphi$. Now, recall that for each sharing action, the state space stays the same, and the new indistinguishability relation for each agent in the product update becomes an intersection of indistinguishability relations in the original model. This means that $(s', e') \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s, e)$ can be equivalently written as $s' \sim_H s$ for some $H \subseteq G$. Since $H \subseteq G$, we get that all s' satisfying $s' \sim_G s$ also satisfy $s' \sim_H s$. This means that, for all $s' \sim_G s$, we have $\mathbb{S}, s' \vDash \varphi$. Therefore, $\mathbb{S}, s \vDash D_G\varphi$, as desired.
- $\Leftarrow$: Note that $\mathbb{S}, s \vDash D_G\varphi$ means for all s' satisfying $s \sim_G s'$, we have $\mathbf{M}, s' \vDash \varphi$. Recall $G = \{a_1, ..., a_k\}$. Now for agent $i \in G$ we can take as the composite event $e = (!(i : a_1); ...; !(i : a_k))$. Then, in each successive product update, the new indistinguishability relation for i is obtained by intersecting the previous result with the next $\sim_{a_i}$, until all a_is have been included. Therefore, $(s, e) \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s', e')$ is equivalent to $s(\bigcap_{a_i\in G} \sim_{a_i})s'$ which is equivalent to $s \sim_G s'$. Therefore, $\mathbb{S}, s' \vDash \varphi$ for all $s' \in S, e' \in \mathsf{E}^m$ such that $(s, e) \sim_i^{\mathbf{M}\otimes\mathbf{E}^m} (s', e')$, as desired. Since $i \in G$ was arbitrary, we have that

$$\mathbb{S}, s \vDash D_G\varphi \implies \mathbb{S}, s \vDash \mathcal{K}_i^G\varphi$$

for all $i \in G$. Then since $\mathcal{K}_G\varphi = \bigwedge_{i\in G} \mathcal{K}_i^G \varphi$, we have:

$$\mathbb{S}, s \vDash D_G\varphi \implies \mathbb{S}, s \vDash \mathcal{K}_G\varphi$$

As desired.

□

The idea here is: If the available actions are only sharing actions, then potential group knowledge for G implies distributed knowledge for G; If the available actions include all sharing actions, then distributed knowledge for G implies potential group knowledge for G. We can therefore see distributive knowledge of G as the special case of potential collaborative knowledge of G where we have all and only the sharing actions between agents of G as the capabilities of G.

Next, we show that the main axiom for distributive knowledge does not hold in general for $\mathcal{K}_G$. This implies that our notion is not a strict expansion of distributed knowledge; in other words, that formulas can be distributed knowledge without being potential collaborative knowledge. The main axiom for the distributed knowledge modality D_G is:

$$K_i\varphi \to D_G\varphi \qquad\qquad \text{for all } i \in G$$

For the $\mathcal{K}_G$ operator, this formula becomes:

$$\mathcal{K}_{\{i\}}\varphi \to \mathcal{K}_G\varphi \qquad\qquad \text{for all } i \in G$$

To show this is not valid on epistemic setups, take the following counterexample: Let p be a proposition. Consider the case with two agents where agent 1 knows p and agent 2 does not know p. Then $\mathbf{M} = (S, \sim_1, \sim_2, \| \bullet \|)$ where $S = \{s, t\}$, $\sim_1 = \{(s, s), (t, t)\}$, $\sim_2 = \{(s, s), (s, t), (t, s), (t, t)\}$, and $\|p\| = \{s\}$. Now, let the epistemic group setup $\mathbb{S}$ be $\mathbf{M}$ along with no event models (or only the event where nothing happens). Then there are no available actions the group can execute to get agent 2 to know p.

If we want to make potential collaborative knowledge a strict extension of distributive knowledge, we need to add the condition that each agent is able to share all they know with each other agent in G.

Chapter 4

Towards an Axiomatisation

When introducing a semantics, logicians are usually interested in having a sound and complete axiomatisation for that semantics. In this section, we conjecture an axiomatisation for our logic over the language $\mathcal{L}^{static}_{\mathcal{K}^G_H}$. We leave the completeness proof, as well as the axiomatisation for the dynamic extension over the language $\mathcal{L}_{\mathcal{K}^G_H}$ as future work. In Chapter 5, we discuss what reduction laws for the dynamic extension might look like, and the core idea for the completeness proof. In this section, we use both the semantics in Subsection 3.2.2 and in Subsection 3.2.3. Remark 45 explains why we can do this despite axiomatising the logic over $\mathcal{L}^{static}_{\mathcal{K}^G_H}$ only.

4.1 Modal logic validities

In standard epistemic logics, the knowledge modality is a box modality of some normal modal logic. It is then a natural question to ask whether our potential collaborative distributive knowledge operators ($\mathcal{K}^G_H$ for all $G, H \subseteq N$) also are. The following properties are sufficient for this to be the case (for all $G, H \subseteq N$), since they provide an equivalent axiomatisation of the normal modal logic K:

$$\vDash \varphi \to \psi \quad \Longrightarrow \quad \vDash \mathcal{K}^G_H\varphi \to \mathcal{K}^G_H\psi \qquad \text{(Propositional Monotonicity Rules)}$$

$$\vDash \varphi \quad \Longrightarrow \quad \vDash \mathcal{K}^G_H\varphi \qquad \text{(Necessitation)}$$

$$(\mathcal{K}^G_H\varphi \wedge \mathcal{K}^G_H\psi) \to \mathcal{K}^G_H(\varphi \wedge \psi) \qquad \text{(Closure Under Conjunction)}$$

The following validity is both interesting in its own right and useful for proving the other validities considered in this section. It states that for any $G, H \subseteq N$, distributive knowledge for H of φ, $K_i\varphi$ implies potential collaborative distributive knowledge of φ, $\mathcal{K}^G_H\varphi$.

Lemma 48 (Distributive Knowledge Implies Potential Knowledge)**.** *For any subgroups $G, H \subseteq N$, the following is valid:*

$$D_H\varphi \to \mathcal{K}^G_H\varphi$$

Proof. Let $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|), \mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$ be an epistemic setup, and $s \in S$ a state. We can then use the case where there are 0 events in the semantics of the $\mathcal{K}^G_H$ modality to prove $\mathbb{S}, s \vDash D_H\varphi \to \mathcal{K}^G_H\varphi$. Recall the full semantics is:

$$\mathbb{S}, s \vDash \mathcal{K}^G_H\varphi \quad \textit{iff} \quad \text{there exists } m \in \omega \text{ and } e \in \mathcal{A}^{-fobs}_G \text{ with } (s, e) \in S \otimes \mathsf{E}^m$$
$$\text{such that, for all } s' \in S, e' \in \mathsf{E} \text{ such that } (s', e') \in S \otimes \mathsf{E}^m,$$

$$\text{if } (s,e) \sim_H^{\mathbf{M}\otimes\mathsf{E}^m} (s',e'),\ \text{then } \mathbb{S}, s' \vDash \varphi$$

Substituting $m = 0$ into the right gives:

$$\text{For all } s' \in S,\ \text{if } s' \sim_H s,\ \text{then } \mathbb{S}, s' \vDash \varphi$$

Which is just the semantic clause for $D_H\varphi$. Therefore, $\mathbb{S}, s' \vDash D_H\varphi$ implies $\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi$. Since $\mathbb{S}, s$ were arbitrary, $D_H\varphi \to \mathcal{K}_H^G\varphi$ is a valid on all epistemic group setups. □

Proposition 49 ($\mathcal{K}$-Necessitation and Propositional Monotonicity Rules)**.** *The potential collaborative knowledge modality for any $G, H \subseteq N$, $\mathcal{K}_H^G$ satisfies the Necessitation and Propositional Monotonicity Rules.*

Proof. For Necessitation, assume that $\vDash \varphi$ for some formula φ. Let $\vDash \varphi \to \psi$. Let $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$, $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$ be an epistemic setup, and $s \in S$ a state. By necessitation for distributive knowledge, we have $\mathbb{S}, s \vDash D_H\varphi$. Now by Distributive Knowledge Implies Potential Knowledge (Lemma 48), we have $\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi$. Since $\mathbb{S}$ and s were arbitrary, $\vDash \mathcal{K}_H^G\varphi$, as desired.

For the Propositional Monotonicity Rules, let φ and ψ be arbitrary formulas, and assume $\vDash \varphi \to \psi$. Let $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$, $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$ be an epistemic setup, and $s \in S$ a state. Assume $\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi$. Then there is some composite event $e \in \mathcal{A}_G^{-fobs}$ where $(s,e) \in S \otimes \mathsf{E}^m$ and for all $s' \in S, e' \in \mathsf{E}^m$ such that $(s',e') \in S \otimes \mathsf{E}^m$, if $(s,e) \sim_H^{\mathbf{M}\otimes\mathsf{E}^m} (s',e')$ then $\mathbb{S}, s' \vDash \varphi$. But now since $\mathbb{S}, t \vDash \varphi \to \psi$ for all $t \in S$, we have that whenever $\mathbb{S}, t \vDash \varphi$ we also have $\mathbb{S}, t \vDash \psi$. This proves $\mathbb{S}, s \vDash \mathcal{K}_H^G\psi$. Therefore, since $\mathbb{S}$ and s were arbitrary, $\vDash \mathcal{K}_H^G\varphi \to \mathcal{K}_H^G\psi$, as desired. □

The following validity, closure under conjunction, makes use of core properties of our particular event models proven in Subsection 3.1.2. Most notably, we use Concatenation of Events. Note that this means that closure under conjunction only holds because of our particular selection of events, and that with some other choices, we would lose closure under conjunction.

Proposition 50 (Closure under Conjunction)**.** *The potential collaborative distributive knowledge modality for each $G, H \subseteq N$, $\mathcal{K}_H^G$, satisfies Closure under Conjunction.*

Proof. Let $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|)$, $\mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$ be an epistemic setup, and $s \in S$ a state. Assume $\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi \wedge \mathcal{K}_H^G\psi$. Then there exists $m_1 \in \omega$ and $e^1 = (e_1^1; ...; e_{m_1}^1) \in \mathcal{A}_G^{-fobs}$ with $(s, e^1) \in S \otimes \mathsf{E}^{m_1}$ such that, for all $s' \in S_{s,e^1}^{\mathbb{S},H}$, we have $\mathbb{S}, s' \vDash \varphi$.

There also exist $m_2 \in \omega$ and $e^2 = (e_1^2; ...; e_{m_2}^2) \in \mathcal{A}_G^{-fobs}$ with $(s, e^2) \in S \otimes \mathsf{E}^{m_2}$ such that, for all $s' \in S_{s,e^2}^{\mathbb{S},H}$, we have $\mathbb{S}, s' \vDash \psi$.

Then by Concatenation of Events (Lemma 32), we get $(s, e^1, e^2) \in S \otimes \mathsf{E}^{m_1+m_2}$.

We claim that if $s' \in S_{s,e^1,e^2}^{\mathbb{S},H}$ then $s' \in S_{s,e^1}^{\mathbb{S},H}$ and $s' \in S_{s,e^2}^{\mathbb{S},H}$. Note that $s' \in S_{s,e^1}^{\mathbb{S},H}$ follows immediately. Now, note that by Lemma 30, we have $\sim_H^{\mathbf{E}_{e_1^1}\otimes\ldots\otimes\mathbf{E}_{e_{m_1}^1}\otimes\mathbf{E}_{e_1^2}\otimes\ldots\otimes\mathbf{E}_{e_{m_2}^2}} \subseteq \sim_H^{\mathbf{E}_{e_1^2}\otimes\ldots\otimes\mathbf{E}_{e_{m_2}^2}}$. This proves $s' \in S_{s,e^2}^{\mathbb{S},H}$.

Therefore, for each $s' \in S_{s,e^1,e^2}^{\mathbb{S},H}$, we have $\mathbb{S}, s' \vDash \varphi \wedge \psi$. Therefore, $\mathbb{S}, s \vDash \mathcal{K}_H^G(\varphi \wedge \psi)$. □

This proves that for each $G, H \subseteq N$ the operator $\mathcal{K}_H^G$ satisfies all the axioms and inference rules of the box modality of a normal modal logic.

4.2 Common epistemic logic validities

In addition to the axioms of the modal logic K, the knowledge modalities of epistemic logic also satisfy several additional axioms. It is therefore natural to ask which of the standard axioms of actual knowledge the $\mathcal{K}_H^G$ modality satisfies. The most uncontroversial axioms for knowledge, factivity and positive introspection, are satisfied. However, note that this is a consequence of the knowledge modalities of our underlying models satisfying these properties. If we do not want positive introspection for $\mathcal{K}_H^G$, then we could instead consider epistemic models that do not satisfy positive introspection.

Proposition 51 ($\mathcal{K}$-Factivity and $\mathcal{K}$-Positive Introspection)**.** *Potential collaborative knowledge satisfies the following validities of S5:*

$$\mathcal{K}_H^G\varphi \to \varphi \qquad \text{(Factivity)}$$
$$\mathcal{K}_H^G\varphi \to \mathcal{K}_H^G\mathcal{K}_H^G\varphi \qquad \text{(Positive Introspection)}$$

Proof. Let $\mathbb{S} = (\mathbf{M} = (S, (\sim_i)_{i\in N}, \|.\|), \mathbf{E} = (\mathsf{E}, (\sim_i)_{i\in N}, \bullet(\bullet), pre))$ be an epistemic group setup, and $s \in S$ a state.

To prove both Factivity and Positive Introspection, assume $\mathbb{S}, s \vDash \mathcal{K}_H^G\varphi$. Then there is some composite event $e \in \mathcal{A}_G^{-fobs}$ where $(s, e) \in S \otimes \mathsf{E}^m$ and for all $s' \in S$ such that there exists $e' \in \mathsf{E}^m$ with $(s', e') \in S \otimes \mathsf{E}^m$ and $(s) \sim_H (s', e')$, we have $\mathbb{S}, s' \vDash \varphi$. Recall we call the set of states s' that we need to evaluate φ at $S_{s,e}^{\mathbb{S},H}$.

For Factivity, note that by reflexivity of our epistemic models and event models, $s \in S_{s,e}^{\mathbb{S},H}$. Therefore $\mathbb{S}, s \vDash \varphi$, as desired.

For Positive Introspection, let $s' \in S_{s,e}^{\mathbb{S},H}$ be arbitrary. Then there exists e' such that $(s', e') \in S \otimes \mathsf{E}^m$ and $(s, e) \sim_H (s', e')$. We claim that for all $s^* \in S$ where there exists e^* with $(s^*, e^*) \in S \otimes \mathsf{E}^m$ and $(s', e') \sim_H (s^*, e^*)$, we have $s^* \in S_{s,e}^{\mathbb{S},H}$. Consider one such $s^* \in S_{s',e'}^{\mathbb{S},H}$ and its corresponding event e^*. By the transitivity of the $\sim_H^{\mathbf{M}\otimes\mathbf{E}^m}$ relation (transitive since the intersection of transitive relations is transitive), we get $(s, e) \sim_H (s^*, e^*)$. Then since $(s^*, e^*) \in S \otimes \mathsf{E}^m$, we have $s^* \in S_{s,e}^{\mathbb{S},H}$. This establishes our claim, and proves that for all $s' \in S_{s,e}^{\mathbb{S},H}$, we have $\mathbb{S}, s' \vDash \mathcal{K}_H^G\varphi$. Therefore, $\mathbb{S}, s \vDash \mathcal{K}_H^G\mathcal{K}_H^G\varphi$, as desired. □

Negative introspection is the most controversial axiom of standard epistemic logic. Since we have chosen $\mathcal{P}$ to be a countable and not necessarily finite set of propositional variables, our potential collaborative distributive knowledge operator does not satisfy it. We have countably many rather than finitely many propositional variables to be able to represent statements involving numbers, such as "there are n stars in the universe".

Proposition 52 (Failure of Negative Introspection)**.** *Negative introspection is not valid for potential collaborative knowledge:*

$$\nvDash \neg\mathcal{K}_H^G\varphi \to \mathcal{K}_H^G\neg\mathcal{K}_H^G\varphi$$

Proof. Consider the following counterexample: Let $N = \{a\}$ and $\mathbf{M} = \{\omega, \sim_a, \|.\|\}$. Where:

1. $\sim_a = \omega^2$.

2. For each $i \in \omega$, $i \neq 0$, $\|p_i\| = \{i\}$ (p_0 is true nowhere).
   $\|p\| = \{n \in \omega \mid n \text{ is even }\}$.

3. The real world is 0.

Now, let $\mathsf{E} = \{(a :_{\{a\}} p_i) \mid i \in \omega\} \cup \{(a :_{\{a\}} \neg p_i) \mid i \in \omega\}$.

Then, after each finite chain of observations, the agent does not learn p, because the only way for them to learn p is to eliminate all odd worlds, which requires finitely many actions, or to discover the real world (which they cannot check directly). Therefore $\{\omega, \sim_a, \|.\|, \mathsf{E}\}, 0 \vDash \neg\mathcal{K}^{\{a\}}_{\{a\}} p$.

Also, the only world in which a cannot learn whether p is true or false is in the real world, 0. If any other world were the real world, they would be able to eventually observe it. Therefore, the only way for a to learn that they can never learn p is to check all worlds other than 0, one by one. This is not finite, therefore $\{\omega, \sim_a, \|.\|, \mathsf{E}\}, 0 \vDash \neg\mathcal{K}^{\{a\}}_{\{a\}} \neg\mathcal{K}^{\{a\}}_{\{a\}} p$. This shows that negative introspection fails. □

Our notion of potential collaborative distributive knowledge therefore satisfies all the axioms and inference rules of an S4 modality.

4.3 Additional validities

The validities in this subsection are important because they form part of our conjectured axiomatisation.

We have monotonicity in the size of both the collaborating group and the learning group.

Proposition 53 ($\mathcal{K}$-Monotonicity)**.** *If $G \subseteq G'$ and $H \subseteq H'$ then:*

$$\mathcal{K}^G_H \varphi \to \mathcal{K}^{G'}_{H'} \varphi$$

Proof. We show $\vDash \mathcal{K}^G_H \varphi \to \mathcal{K}^{G'}_H \varphi$ and $\vDash \mathcal{K}^{G'}_H \varphi \to \mathcal{K}^{G'}_{H'} \varphi$. For this, let $\mathbb{S} = (\mathbf{M}, \mathbf{E})$ be an epistemic group setup and $s \in \mathbb{S}$ a state. Assume $\mathbb{S}, s \vDash \mathcal{K}^G_H \varphi$. Then we can take the same events to show $\mathbb{S}, s \vDash \mathcal{K}^{G'}_H \varphi$, since $\mathcal{A}^{-fobs}_G \subseteq \mathcal{A}^{-fobs}_{G'}$ (Proposition 35.2). This proves $\mathbb{S}, s \vDash \mathcal{K}^{G'}_H \varphi$. Since $\mathbb{S}$ and s were arbitrary, $\vDash \mathcal{K}^G_H \varphi \to \mathcal{K}^{G'}_H \varphi$.

Now assume $\mathbb{S}, s \vDash \mathcal{K}^{G'}_H \varphi$. Then there is a composite event $e \in \mathcal{A}^{-fobs}_{G'}$ such that for all $s' \in S^{\mathbb{S},H}_{s,e}$, we have $\mathbb{S}, s' \vDash \varphi$. Note that $\sim^{\mathbf{M}\otimes\mathbf{E}^m}_{H'} \subseteq \sim^{\mathbf{M}\otimes\mathbf{E}^m}_{H}$. This means that that $S^{\mathbb{S},H'}_{s,e} \subseteq S^{\mathbb{S},H}_{s,e}$. Since all $s' \in S^{\mathbb{S},H}_{s,e}$ satisfy $\mathbb{S}, s' \vDash \varphi$, it is also the case that all states $s' \in S^{\mathbb{S},H'}_{s,e}$ satisfy $\mathbb{S}, s' \vDash \varphi$. Therefore, $\mathbb{S}, s \vDash \mathcal{K}^{G'}_{H'} \varphi$. Since $\mathbb{S}$ and s were arbitrary, we get $\vDash \mathcal{K}^{G'}_H \varphi \to \mathcal{K}^{G'}_{H'} \varphi$, as desired. □

We also have some validities for formulas of the form $cap(e)$, where e is an event in $\mathcal{L}_e$.

Proposition 54 ($cap(e)$ validities)**.** *The following are valid for all $G \subseteq N$, $i \in G$, $j \in N$, propositional variables p and boolean formulas ψ:*

1. $cap((\psi/j : p)) \leftrightarrow cap((\psi/j : \neg p))$ *and* $cap((\psi/j : \neg p)) \leftrightarrow cap((\psi/j : -_p))$

2. $cap((i :_G p)) \leftrightarrow cap((i :_G \neg p))$

3. $cap(e; e') \leftrightarrow cap(e) \wedge cap(e')$

Proof. Let $\mathbb{S} = (S, (\sim_i)_{i \in N}, \|.\|, \mathsf{E})$ be an epistemic setup and let $s \in \mathbb{S}$ be a state.

1. Assume $\mathbb{S}, s \vDash cap((\psi/j : x))$, where x is $p, \neg p$ or $-_p$. Then $(\psi/j : x) \in \mathsf{E}$. Then, by the definition of permissible event sets E all of $(\psi/j : p), (\psi/j : \neg p)$ and $(\psi/j : -_p)$ are in E (as stated in Remark 27). Therefore $\mathbb{S}, s \vDash cap((\psi/j : x))$ for all $x \in \{p, \neg p, -_p\}$.

2. The proof for this part is essentially identical to the proof for part 1.

3. This follows directly from the definition of E^* and the semantics of $cap(e)$.

□

We now have almost all validities and inference of our proposed axiomsation. We have one important addition left.

4.4 Main axiom schema

This section contains our proposed main axiom schema and its proof, as well as some instances of it that give insight into how our permissible events operate to affect the potential knowledge of the group.

This axiom schema can be understood as a generalisation of the following principle, which applies to propositional variables, to arbitrary formulas:
If executing an event is possible and a proposition would become potential collaborative distributive knowledge after executing that event, then the proposition is potential collaborative distributive knowledge.

Proposition 55 (Main axiom schema)**.** *The following is valid for all $H, G \subseteq N$ and $e \in \mathcal{A}_G^{-fobs}$:*

$$\left[cap(e) \wedge pre_e \wedge \mathcal{K}^G_{e(H)}\Big(\big(\bigvee_{e' \sim_H e} pre_{e'}\big) \rightarrow \varphi\Big)\right] \rightarrow \mathcal{K}^G_H \varphi$$

Proof. Let $\mathbb{S} = (S, (\sim_i)_{i \in N}, \|.\|, \mathsf{E})$ be an epistemic group setup and $s \in S$ a state. Assume $\mathbb{S}, s \vDash cap(e) \wedge pre_e \wedge \mathcal{K}^G_{e(H)}\Big(\big(\bigvee_{e' \sim_H e} pre_{e'}\big) \rightarrow \varphi\Big)$.

Now, since $\mathbb{S}, s \vDash \mathcal{K}^G_{e(H)}\Big(\big(\bigvee_{e' \sim_H e} pre_{e'}\big) \rightarrow \varphi\Big)$, we get that there is some $e'' \in \mathcal{A}_G^{-fobs}$ such that $\mathbb{S}, s \vDash pre_{e''}$ and:

$$\text{for all } s'' \sim_{e''(e(H))} s, \text{ we have } \mathbb{S}, s'' \vDash \big(\bigvee_{e_1 \sim_{e(H)} e''} pre_{e_1}\big) \rightarrow \big(\big(\bigvee_{e' \sim_H e} pre_{e'}\big) \rightarrow \varphi\big). \qquad (*)$$

Now take the composite event $(e''; e)$. Note first that $(e''; e) \in \mathcal{A}_G^{-fobs}$.

Additionally, we have $\mathbb{S}, s \vDash pre_e$ and $\mathbb{S}, s \vDash pre_{e''}$. Therefore, $(s, e) \in S \otimes \mathsf{E}^*$ and $(s, e'') \in S \otimes \mathsf{E}^*$. Then, by Concatenation of Events (Lemma 32), $(s, e'', e) \in S \otimes \mathsf{E}^*$. Therefore, $\mathbb{S}, s \vDash pre_{e''}$ and $(\mathbf{M} \otimes \mathbf{E} \otimes ... \otimes \mathbf{E}, \mathbf{E}), (s, e'') \vDash pre_e$, which by Proposition 38 is equivalent to $\mathbb{S}, s \vDash pre_{(e''; e)}$.

Now, let $s' \in S$ be such that $s' \sim_{(e''; e)(H)} s$ and $\mathbb{S}, s' \vDash \bigvee_{(e_1; e_2) \sim_H (e''; e)} pre_{(e_1; e_2)}$. Then there is some event $(e_1; e_2) \sim_H (e''; e)$ where $\mathbb{S}, s' \vDash pre_{(e_1; e_2)}$. By the definition of the indistinguishability relation for composite events and the definition of preconditions for composite events, this implies $e_1 \sim_{e(H)} e''$ and $\mathbb{S}, s' \vDash pre_{e_1}$. Therefore, $\mathbb{S}, s' \vDash \bigvee_{e_1 \sim_{e(H)} e''} pre_{e_1}$.

In addition $(e_1; e_2) \sim_H (e''; e)$ and $\mathbb{S}, s' \vDash pre_{(e_1; e_2)}$, also by the definition of the indistinguishability relation for composite events and the definition of preconditions for composite events, imply $e_2 \sim_{e(H)} e$

and $(\mathbf{M} \otimes \mathbf{E} \otimes ... \otimes \mathbf{E}, \mathbf{E}), (s', e_1) \vDash pre_{e_2}$. Now, recall that we also have $\mathbb{S}, s \vDash pre_e$, and $s \sim_{e(H)} s'$. Hence we get $\mathbb{S}, s' \vDash (\bigvee_{e' \sim_H e} pre_{e'})$.

Therefore, by $(*)$ above, in which we take $s'' := s'$, we can conclude $s' \vDash \varphi$, as desired. $\square$

Note here the role of $e(H)$ in $\mathcal{K}^G_{e(H)}$, in the antecedent of the validities. Since $H \subseteq e(H)$ and we have monotonicity, the axiom schema would still hold if we replaced $e(H)$ with H. However, taking this as an axiom schema instead would not allow us to prove all validities, and hence would not work as the main axiom schema for a complete axiomatisation.

The validities that follow are instances of our main axiom schema, and they show how each of our permissible events work. The first validity also helps to illustrate the role of $e(H)$ in the axiom schema, and hence justifies our use of $\mathcal{K}^G_H$, rather than just $\mathcal{K}_G$ as the main operator in our formalism.

The first is about sharing events. It says that if an agent j can share all they know with another agent i, and if i and j can take actions with the collaboration of G to come to distributively know φ, then i can come to know φ with the collaboration of G.

Proposition 56. *The following is valid on all epistemic group setups, for $i, j \in G$:*

$$(\mathcal{K}^G_{\{i,j\}}\varphi \wedge cap(!(i : j))) \rightarrow \mathcal{K}^G_i \varphi$$

Proof. Let $\mathbb{S}$ be an epistemic group setup with permissible event set E, and let s be a state in $\mathbb{S}$. Since we have $\mathcal{K}^G_{\{i,j\}}\varphi$, there is an event $e \in \mathcal{A}^{-fobs}_G$ such that for all $s' \in S, e' \in \mathsf{E}^*$ with $(s', e') \in S \otimes \mathsf{E}^*$, if $(s, e) \sim_{\{i,j\}} (s', e')$, then $\mathbb{S}, s' \vDash \varphi$.

Now, we can consider the composite event $(e; !(i : j))$. This event is an element of $\mathcal{A}^{-fobs}_G$. and s satisfies its precondition. Note that for all agents $k \in N$, the only event indistinguishable to $!(i : j)$ is itself, and $!(i : j)(i) = \{i, j\}$. Then all $s' \in S, (e_1; e_2) \in \mathsf{E}^*$ with $(s', e_1, e_2) \in S \otimes \mathsf{E}^*$ and $(s', e_1, e_2) \sim_i (s, e, !(i : j))$ also satisfy $(s', e_1) \in S \otimes \mathsf{E}^*$ and $(s', e_1) \sim_{\{i,j\}} (s, e)$. This shows that if $(s', e_1, e_2) \in S \otimes \mathsf{E}^*$ and $(s', e_1, e_2) \sim_i (s, e, !(i : j))$, then $\mathbb{S}, s' \vDash \varphi$. Therefore, taking our composite event to be $(e; !(i : j))$ allows us to show $\mathbb{S}, s \vDash \mathcal{K}^G_i \varphi$, as desired. $\square$

The above validity is the instance of the main axiom scheme where $e := (!(i : j)$, $H := \{i\}$ and $G := G$. To see this, note first that $pre_{!(i:j)} = \top$. Second, note that $!(i : j) \in \mathcal{A}^{-fobs}_G$. Third, note that for all $k \in N$, $(!(i : j) \sim_k e'$ if and only if $e' =!(i : j)$. Therefore, $(\bigvee_{e' \sim_H e} pre_{e'}) \rightarrow \varphi$ simplifies to φ. Lastly, note that $!(i : j)(\{i\}) = \{i, j\}$.

Now, note that if instead considered the case where we used $\mathcal{K}^G_H$ instead of $\mathcal{K}^G_{e(H)}$ in Proposition 55, the most we could say about the functioning of sharing events would be $(\mathcal{K}^G_i \varphi \wedge cap(!(i : j))) \rightarrow \mathcal{K}^G_i \varphi$ or $(\mathcal{K}^G_{\{i,j\}}\varphi \wedge cap(!(i : j))) \rightarrow \mathcal{K}^G_{\{i,j\}}\varphi$, neither of which capture the functioning of these events. We need to capture how, after sharing events, individuals acquire all the information that was possessed by their sources; in other words, how what was once distributed knowledge before sharing events can become individual knowledge after the events.

The next validity is about informed observations. It says that if l is true and the agent i has the knowledge they require to observe it, i can come to know l.

Proposition 57. *For all $G \subseteq N$, $i \in N$, literals l and boolean formulas ψ:*

$$(l \wedge K_i \psi \wedge cap((\psi / i : l))) \rightarrow \mathcal{K}^{G \cup \{i\}}_i (l)$$

Proof. Let $\mathbb{S}$ be an epistemic group setup with permissible event set E, and let s be a state in $\mathbb{S}$. Since $i \in G \cup \{i\}$ and $\{i\}$ is the active subgroup for $(\psi/i : l)$, we have $(\psi/i : l) \in \mathcal{A}_G^{-fobs}$. Therefore, we can just consider the event $e = (\psi/i : l)$. We have $(s, e) \in S \otimes \mathsf{E}$. Now note that $e(i) = \{i\}$ and the only event $e' \in \mathsf{E}$ satisfying $e' \sim_i e$ is e itself. Hence, for all $s' \sim_{e(i)} s$ and $e' \sim_i e$ with $(s', e') \in S \otimes \mathsf{E}$, we have $\mathbb{S}, s' \vDash l$, as desired. $\square$

This is the instance of the main axiom schema where $e := (\psi/i : l))$, $H := \{i\}$, and $G := G \cup \{i\}$. To see why, note that since e itself is the only event indistinguishable to e for i, $(\bigvee_{e' \sim_H e} pre_{e'}) \to \varphi$ becomes $(K_i\psi \wedge l) \to l$, which is a tautology. Hence, this whole part of the antecedent falls away.

4.5 Summary: Axiomatisation

We have now presented several validities, most of which form part of our proposed axiomatisation. Table 4.1 presents our proposed axiomatisation for our logic over $\mathcal{L}^{static}_{\mathcal{K}^G_H}$.

(I) **Axioms and rules of classical propositional logic**

(II) **Axioms for capability operator**

(Informed Obs.) $cap((\psi/j : p)) \leftrightarrow cap((\psi/j : \neg p)) \leftrightarrow cap((\psi/j : -_p))$

(Basic Obs.) $cap((i :_G p)) \leftrightarrow cap((i :_G \neg p))$

(Comp. Events) $cap(e; e') \leftrightarrow cap(e) \wedge cap(e')$

(III) **Axioms and rules for distributive knowledge:**

(D-Necessitation) From φ, infer $D_G\varphi$

(D-Distribution) $D_G(\varphi \to \psi) \to (D_G\varphi \to D_G\psi)$

(D-Factivity) $D_G\varphi \to \varphi$

(D-Pos. Introspection) $D_G\varphi \to D_G D_G\varphi$

(D-Neg. Introspection) $\neg D_G\varphi \to D_G\neg D_G\varphi$

(D-Monotonicity) $D_G\varphi \to D_H\varphi$, for sets $G \subseteq H$

(IV) **Axioms and rules for potential collaborative distributive knowledge:**

($\mathcal{K}$-Necessitation) From φ, infer $\mathcal{K}^G_H\varphi$

(Prop. Monotonicity rules) From $\varphi \to \psi$, infer $\mathcal{K}^G_H\varphi \to \mathcal{K}^G_H\psi$

(Closure under Conjunction) $(\mathcal{K}^G_H\varphi \wedge \mathcal{K}^G_H\psi) \to \mathcal{K}^G_H(\varphi \wedge \psi)$

($\mathcal{K}$-Factivity) $\mathcal{K}^G_H\varphi \to \varphi$

($\mathcal{K}$-Pos. Introspection) $\mathcal{K}^G_H\varphi \to \mathcal{K}^G_H\mathcal{K}^G_H\varphi$

($\mathcal{K}$-Monotonicity) $\mathcal{K}^G_H\varphi \to \mathcal{K}^{G'}_{H'}\varphi$, for sets $H \subseteq H'$ and $G \subseteq G'$

(Main Axiom Schema) $[cap(e) \wedge pre_e \wedge \mathcal{K}^G_{e(H)}((\bigvee_{e' \sim_H e} pre_{e'}) \to \varphi)] \to \mathcal{K}^G_H\varphi$

Table 4.1: Proposed axiomatisation for our logic over $\mathcal{L}^{static}_{\mathcal{K}^G_H}$

Note that the validities and rules for the distributive knowledge operator $D_G\varphi$ have already been proven [BS24], and since the semantics for $D_G\varphi$ in this thesis matches the standard semantics for $D_G\varphi$, the validity proofs will be identical in our setting.

We have therefore proven that this axiomatisation is sound.

Theorem 58 (Soundness)**.** *The axioms and inference rules in Table 4.1 are sound with respect to epistemic group setups.*

We leave the completeness proof as future work.

Chapter 5

Conclusion and Future Directions

In this thesis, we introduced a notion of group knowledge/knowability that we called potential collaborative knowledge, with the motivation to contribute to the philosophical literature and to create a notion that could have applications in artificial intelligence. We informally defined a statement being potential collaborative knowledge for a group as meaning that the group can collaborate so that each individual member can learn it. We formally defined this notion in two stages. The first stage involves formalising the actions. For this, we first defined special types of data-exchange event models to represent actions taken by groups, and a syntax for the specific elements of these event models. We then defined the notions of active subgroups of actions and capabilities of subgroups to be able to consider the potential collaborative knowledge *of subgroups*. In the second stage, we created a syntax and semantics to model collaboration, and defined potential collaborative knowledge $\mathcal{K}_G$ in this setting.

After defining potential collaborative knowledge, we formalised an example both to demonstrate the functioning of our semantics and to formally show our core philosophical point. Namely, that the group can learn more via collaboration than any individual member could learn alone, and that if we view potential collaborative knowledge as a notion of group knowledge, it is a non-summative one. An implication of this is that some knowledge formation processes are irreducibly social, and that more attention should be paid to social aspects of knowledge formation in epistemology. We also showed formally that potential collaborative knowledge is a generalisation of distributive knowledge. We ended by proving some validities that we believe form an axiomatisation for our static logic. These validities included all the validities of S4 for $\mathcal{K}_H^G$, monotonicity validities for $\mathcal{K}_H^G$, and an axiom schema capturing the essence of the functioning of our potential collaborative distributive knowledge operator.

In terms of future directions, the most clear next steps for this research would be to prove that our conjectured axiomatisation is indeed complete, and to extend our axiomatisation to the logic over $\mathcal{L}_{\mathcal{K}_H^G}$. For the completeness proof, the idea would be to create $S4$ pseudo Kripke models, and to prove completeness with respect to these pseudo models. We would then use the fact that our $\mathcal{K}_H^G$ operators are all $S4$ Kripke operators to create a correspondence between our epistemic group setups and these pseudo models.

For the reduction laws, we would expect that they would be standard for negation, conjunction, and distributive knowledge. The reduction law for the potential collaborative distributive knowledge operator is, however, not as clear. One possibility for it would be:

$$[e]\mathcal{K}_H^G\varphi \leftrightarrow \left((pre_e \wedge cap(e)) \rightarrow \bigwedge_{e' \sim_H e} \mathcal{K}_{e(H)}^G[e']\varphi\right)$$

Checking whether or not these are indeed correct is left as future work.

Another logical next step would be to expand the class of events we consider. For example, we could consider adding events to represent resource sharing, creation, or acquisition. This would require adding post-conditions to our event models. Then, resource acquisitions could be public events in which a subgroup of agents acquire a resource, specified as a postcondition. This would allow us to express examples in which the possession of resources is necessary to make observations, and the creation or acquisition of these resources is part of the collaboration process. Agent i's possession of a given resource r could then be represented as a proposition with additional structure such as $has_i(r)$.

Another option would be to consider which validities we would have if we placed restrictions on which events we allow in our models. We had already shown that distributive knowledge was a special case of potential collaborative knowledge. In Subsection 3.4, we showed that we could make our notion a strict extension of distributive knowledge by always including all sharing events in our epistemic group setups. This suggests considering models with full sharing as a possible future direction of research.

We have already discussed some alternative ways to define potential collaborative knowledge $\mathcal{K}_G$ in terms of $\mathcal{K}_H^G$ that allow us to model other conceptions of group knowledge. A natural future direction from this would be to study the potential collaborative knowledge operators that result from making these changes.

Finally, we could consider different underlying models. We had chosen multi-agent $S5$ epistemic models. We could consider $S4$ epistemic models instead, or even expand to epistemic-doxastic models (to also model beliefs in collaborative settings). A more distant expansion idea along these lines would be to expand to models for evidence, such as topological models.

Bibliography

[Bal+22] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün, and Sonja Smets. "Justified belief, knowledge, and the topology of evidence". In: *Synthese* 200.6 (2022), pages 1–51. https://doi.org/10.1007/s11229-022-03967-6 (cited on page 6).

[Ben70] Jeremy Bentham. *An Introduction to the Principles of Morals and Legislation*. London: Athlone Press, 1970 (cited on page 7).

[Bir14] Alexander Bird. "When Is There a Group that Knows?" In: Nov. 2014, pages 42–63. https://doi.org/10.1093/acprof:oso/9780199665792.003.0003 (cited on page 7).

[BMS16] Alexandru Baltag, Lawrence S. Moss, and Stawomir Solecki. "Dynamic Epistemic Logic". In: *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Winter 2016. Metaphysics Research Lab, Stanford University, 2016. https://plato.stanford.edu/entries/dynamic-epistemic/ (cited on page 14).

[BMS98] Alexandru Baltag, Lawrence S. Moss, and Stawomir Solecki. "The Logic of Public Announcements, Common Knowledge, and Private Suspicions (Extended Abstract)". In: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 1998)*. Evanston, Illinois, USA: TARK, 1998, pages 43–56. http://tark.org/proceedings/tark_jul22_98/p43-baltag.pdf (cited on page 14).

[Bod14] Rachel Boddy. "Epistemic Issues and Group Knowledge". Master's thesis. Institute for Logic, Language and Computation, University of Amsterdam, June 2014. https://eprints.illc.uva.nl/id/eprint/921/1/MoL-2014-03.text.pdf (cited on page 8).

[Bro24] Jessica Brown. *Groups as Epistemic and Moral Agents*. Oxford University Press, 2024. https://doi.org/10.1093/9780191999215.001.0001 (cited on pages 3, 6, 7).

[BRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science 53. Cambridge University Press, 2001 (cited on pages 11, 12, 16).

[BS21] Alexandru Baltag and Sonja Smets. *Learning What Others Know*. 2021. https://arxiv.org/abs/2109.07255 (cited on page 4).

[BS24] Alexandru Baltag and Sonja Smets. "Logics for Data Exchange and Communication". In: *Proceedings of AiML 2024 (Advances in Modal Logic)*. College Publications, 2024. https://www.cs.cas.cz/aiml2024/ (cited on pages 4, 14, 47).

[DHK08] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. 1st edition. Volume 337. Synthese Library. Springer Dordrecht, 2008. https://doi.org/10.1007/978-1-4020-5839-4 (cited on pages 11, 14).

[Fag12] Melinda Fagan. “Collective Scientific Knowledge”. In: *Philosophy Compass* 7.12 (2012), pages 821–831. https://doi.org/10.1111/j.1747-9991.2012.00528.x (cited on page 9).

[Gil87] Margaret Gilbert. “Modelling collective belief”. In: *Synthese* 73 (1987), pages 185–204. https://link.springer.com/article/10.1007/BF00485446 (cited on page 7).

[Han00] Sven Ove Hansson. “Formalization in Philosophy”. In: *The Bulletin of Symbolic Logic* 6.2 (2000), pages 162–175. http://www.jstor.org/stable/421204 (cited on page 5).

[HB20] Lasse D. Hansen and Thomas Bolander. “Implementing Theory of Mind on a Robot Using Dynamic Epistemic Logic”. In: *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20)*. International Joint Conference on Artificial Intelligence Organization, 2020, pages 1615–1621. https://doi.org/10.24963/ijcai.2020/224 (cited on pages 3, 6).

[Hin62] Jaakko Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca, NY: Cornell University Press, 1962 (cited on pages 4, 11, 13).

[Hol16] Wesley H. Holliday. “Epistemic Logic and Epistemology”. In: *Handbook of Formal Philosophy*. Edited by Steen E. Hansson and Vincent F. Hendricks. Springer, 2016, pages 1–18 (cited on page 8).

[HR23] Avram Hiller and Wolfe Randall. “Epistemic Structure in Non-Summative Social Knowledge”. In: *Social Epistemology* 37.1 (2023), pages 30–46. https://doi.org/10.1080/02691728.2022.2121621 (cited on pages 3, 6, 8).

[Lac20] Jennifer Lackey. *The Epistemology of Groups*. New York: Oxford University Press, 2020 (cited on pages 3, 6, 10).

[LP11] Christian List and Philip Pettit. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Edited by Philip Pettit. New York: Oxford University Press, 2011. https://doi.org/10.1093/acprof:oso/9780199591565.001.0001 (cited on pages 3, 6, 7).

[OGG24] Cailin O’Connor, Sanford Goldberg, and Alvin Goldman. “Social Epistemology”. In: *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman. Summer 2024. Metaphysics Research Lab, Stanford University, 2024. https://plato.stanford.edu/archives/sum2024/entries/epistemology-social/ (cited on page 7).

[Qui75] Anthony Quinton. “The Presidential Address: Social Objects”. In: *Proceedings of the Aristotelian Society* 76 (1975), pages 1–viii. http://www.jstor.org/stable/4544878 (cited on page 7).

[Rol08] Kristina Rolin. “Science as collective knowledge”. In: *Cognitive Systems Research* 9.1 (2008), pages 115–124. https://doi.org/10.1016/j.cogsys.2007.07.007 (cited on pages 3, 6).

[Voo93] Frans Voorbraak. “As Far as I Know”. PhD thesis. Utrecht, The Netherlands: Utrecht University, 1993 (cited on page 13).

[Wil00] Timothy Williamson. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000 (cited on page 13).