Epistemic Issues and Group Knowledge

MSc Thesis *(Afstudeerscriptie)*

written by

Rachel Boddy
(born July 29th, 1988 in Amsterdam, The Netherlands)

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

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
June 20, 2014	Dr. Maria Aloni
	Prof. Robert van Rooij
	Prof. Martin Stokhof
	Prof. Johan van Benthem
	Dr. Wesley Holliday
	Prof. Branden Fitelson

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

Formal models for group knowledge can help philosophers gain additional insight into the ramifications of the philosophical concepts that they propose by clarifying the abstract properties of these concepts and their relationship to alternative proposals. To date, however, formal treatments of group knowledge have remained largely disjointed from the related philosophical discussions and are therefore of minimal interest to philosophers. In this thesis, I attempt to bridge this gap by proposing a formal definition of group knowledge that I call *collective knowledge*. *Collective knowledge* is distributed knowledge about common questions and typically lies between common knowledge and full distributed knowledge. It includes two epistemic properties that make it more aligned with philosophical concepts of group knowledge, and that are not modeled by the standard notions from formal epistemology. The first property is that all knowledge is in terms of questions, interpreted as distinctions that define an agent's conceptual framework. The second property is that group knowledge implies an epistemic group, which is a group of agents tied together through mutual interest in each other's knowledge and questions. To model epistemic groups and *collective knowledge*, I introduce new Kripke models that I refer to as *epistemic group models*. I then present an axiomatic system for the logic of *collective knowledge* and prove that it is sound and complete with respect to these new models. As such, I hope to have provided a good first step towards a formal definition of group knowledge that can help advance the philosophical discussion on group knowledge.

Contents

Chapter 1

Introduction

In the formal epistemology literature, notions of group knowledge have long been an important topic of investigation. Several definitions for group knowledge have been proposed and studied, the most important of which are "distributed knowledge" and "common knowledge". These are often used to analyze the information flow within groups of agents, especially in contexts where the agents are assumed to reason about each other's knowledge (Baltag *et al.* 2008; Halpern *et al.* 1995; Lewis, 2006).

In the philosophical literature, discussions of group knowledge have also become increasingly common, with special interest being given to the idea that groups can be treated as collective agents that are capable of knowledge in their own right (Bratman, 1993; Corlett, 1996; Gilbert, 1987; Gilbert, 1989; List and Pettit, 2011; List, 2011; Rolin, 2007). Within this discussion, the term "group knowledge" is used to refer to several distinct concepts, embodying different views about the extent to which knowledge must ultimately be held by an individual, the subjects of group knowledge, and the extent to which group knowledge is accessible to the individual group members.

Epistemic logic attempts to encode the systematic properties of epistemic concepts in order to describe their logical behavior. As such, epistemic logic can help philosophers to gain more insight into the ramifications of epistemic concepts by clarifying the abstract relationship between these concepts (Holliday, 2013; Stalnaker, 2006). In order for epistemic logic to make a real contribution to epistemology, the concepts that it models must coincide with philosophical concepts in terms of pertinent logical properties. With regard to the concept of group knowledge, this is not currently the case as the standard group-epistemic notions from epistemic logic are not well-aligned with the current philosophical discussion of group knowledge as each fails to address some important features of philosophical concepts.

In the philosophical literature, group knowledge is often defined in terms of the types of groups that can be its possible subjects. The underlying assumption is that random sets of individuals cannot have group knowledge simply because they should not be considered possible epistemic subjects (Corlett, 2007; Pettit, 2011). Groups that qualify as epistemic groups must at least be partly defined on the basis of epistemic properties related to knowledge possession, which allows them to behave like (individual) epistemic agents, and explains how it can achieve its knowledge.[1] This suggests that a formal definition of group knowledge should ideally encode not only the way in which group knowledge depends on the knowledge of its constituents, but also the necessary connections between group members that must obtain in order for a group to become

[1] The term "epistemic group" will be further explained below. However, for now it is enough to think of an epistemic group as a group of individuals tied together by an (non-trivial) epistemic property.

an epistemic agent.

Despite the apparent overlap in subject matter, formal treatments of group knowledge have remained largely disjointed from related philosophical discussions. This disconnect is partly attributable to the fact that the former has roots not only in philosophy, but also in mathematics and computer science, and is often driven by other (i.e., non-philosophical) considerations,[2] and that consequently the connection to philosophical discussion is not as directly apparent (Barwise, 1988; Halpern *et al.* 1995).

The objective of this thesis is to present a formal approach to representing group knowledge. More specifically, I propose and defend a formal definition of group knowledge that reflects both how the knowledge of a group depends on the knowledge of its constituent members and how group knowledge is held by an epistemic group. I refer to this conception of group knowledge as *collective knowledge.* To defend this account, I provide both philosophical and logical support. On the philosophical side, this is done by defending two properties that a formal definition of (group) knowledge should reflect, namely, that all knowledge is in terms of questions and that group knowledge implies an epistemic group. To model these properties, I introduce new Kripke models for knowledge, called *epistemic group models.* On the logical side, I support this definition by presenting an epistemic logic for it, and showing that it is sound and complete with respect to the class of *epistemic group models.*

Intuitively, *collective knowledge* is distributed knowledge of a group *about a common question.* As such, like distributed knowledge, it represents group knowledge as the knowledge that the agents in a group would know were they to combine their knowledge. Yet, unlike distributed knowledge, it is based on the further assumptions that group knowledge requires groups of agents that view each other, as well as their group, as knowledge sources, and that group knowledge is limited to answers to common questions.[3] It should be expected that group knowledge behaves like individual knowledge, and similarly, that an epistemic group agent behaves like its individual counterpart. In fact, individual knowledge must be a trivial case of group knowledge. By assuming that questions play an important role in defining group knowledge, I am therefore also assuming that they play an important role in defining individual knowledge.

In defining *epistemic group models* and *collective knowledge*, I draw on resources from formal epistemology and formal approaches to modeling of questions. In particular, following van Benthem and Minică (2009), I add an issue-relation to the standard models for knowledge from epistemic logic and extend the language of epistemic logic with modalities for the key concepts. While logics for questions have been studied in the literature (see e.g. Aloni *et al.* 2013; Groenendijk and Stokhof, 1984; Groenendijk, 1999; Groenendijk and Roelofsen, 2009), to the best of my knowledge the role of questions for group knowledge has not yet been considered.

The thesis is organized as follows: chapter 2 starts with background on epistemology and epistemic logic, which shall henceforth be assumed as common knowledge. Chapter 3

[2]Such as understanding the information flow in a multi-agent system in which the 'agents' are wires rather than humans (Halpern *et al.*, 1995).

[3]Questions are taken to represent the distinctions that define an agent's conceptual framework.

provides a brief discussion of some of the important choices or properties that distinguish competing philosophical notions of group knowledge from each other. I then look at the notions of common knowledge and distributed knowledge and consider to what extent they exhibit these properties. Chapter 4 argues that knowledge presupposes questions. I start with individual knowledge, and explain in what sense an individual's knowledge and questions are tied together. I then propose an appropriate condition to be imposed on the accessibility relation of models for knowledge that captures this tie. In chapter 5, I discuss conditions to be imposed on agents' issue-relations that are needed for them to coherently represent the knowledge and questions of others. I then propose the notion of an *epistemic group* in terms of these conditions and introduce *epistemic group models* that satisfy them. Chapter 6 brings together the ideas expressed in the two preceding chapters. I propose formal definitions for *collective knowledge* and other group-epistemic concepts. In chapter 7, I then present an axiomatic system for a logic with *collective knowledge*. I show that this system, called *epistemic group logic*, is sound and complete with respect to the class of epistemic group models. Chapter 8 provides some concluding remarks.

Note: Some of the work in this thesis is joint with Alexandru Baltag and Sonja Smets, and will appear in a joint paper, called "An interrogative approach to group knowledge" (Baltag *et al.* 2014). In particular, some of the ideas expressed in chapters 6 and 7 draw on this work.

Chapter 2

Background on epistemology and epistemic logic

While I assume that the reader has some basic understanding of philosophical argumentation and the principles of logic, I anticipate that not all readers will possess sufficient background in both epistemology and epistemic logic in order to follow the subsequent argumentation. This chapter attempts to at least partially fill this gap by providing a summary of a number of key concepts from both areas. In so doing, it also introduces the main formal notation that I use in this thesis.

2.1 Epistemology

One of the oldest problems in philosophy and a central problem of epistemology is to elucidate the concept of knowledge by proposing conditions that are necessary and jointly sufficient for its possession. In the past, philosophers have mainly been concerned with propositional knowledge that is held by individuals.[1] Apart from propositional knowledge, however, there are several other types of knowledge, including knowledge *how, who, what, where* and *why* (etc.). We routinely ascribe knowledge to agents not only when they know that a proposition is true, but also when they know *how* to grill beef, *what* they are wearing, *why* they are wearing it (etc.). Yet, philosophers have focused on propositional knowledge. Thus understood, knowledge is a binary relation between a subject s and a proposition p: s knows that p (Schaffer, 2007). Other propositional attitudes (such as belief and hope) are similarly analyzed as binary relations between a subject and a proposition. In the quest to define the necessary and sufficient conditions for s to know that p, a definition should capture all the knowledge ascriptions that philosophers want to legitimize, and at the same time offer a response to skepticism. Contemporary mainstream epistemology is still largely concerned with this problem (Hendricks, 2006).[2]

2.1.1 Knowledge as justified true belief

Although disagreement among philosophers about the definition of knowledge continues, the view that knowledge is some version of justified true belief (JTB) is still prevalent in contemporary epistemology. This conception of knowledge is traced as far back as

[1] The term "proposition" refers to a factual statement or, in modal terms, a set of possible worlds. Further, it is enough to assume an uncontroversial understanding of truth, *viz.* that a proposition is true whenever it describes an aspect of the world that in fact obtains.

[2] Here, the term 'mainstream epistemology' refers to proposals that analyze knowledge in terms of the conditions that are necessary and jointly sufficient for its possession (Ibid.: 14).

Plato's writings, and has set the standards for an adequate analysis of knowledge. The three components – truth, belief and justification – lead to the standard definition of knowledge as follows:

Definition 1 (Knowledge as *JTB*)**.** An agent a knows that p, if and only if: (1) p is true, (2) a believes that p and (3) a is justified in believing that p.

It is hardly controversial that knowledge is factive: a knows that p only if p is true. We can only know true propositions. While there are other types of knowledge as well, epistemology is primarily concerned with knowledge of fact–hence the truth component. Similarly, an agent cannot know what she does not believe to be true: an agent knows that p only if she believes that p is true. Beliefs are mental states–more precisely, they are attitudes towards a proposition. The other two conditions (truth and justification) are meant to distinguish knowledge from other propositional attitudes, such as (mere) belief, hope and guessing. Finally, there is the justification condition: an agent knows that p only if she is justified in believing that p. This condition is meant to ensure that knowledge is not truthful by accident–as the result of wishful thinking or lucky guessing–but rather that a believes that p *because* p *is true*. The justification component is most controversial. Philosophers disagree on what makes a belief justified,[3] but also on whether knowledge should be justified belief or not at all. Competing theories of knowledge tend to disagree on issues regarding the sources and justification of knowledge, and in particular, on the matter of what distinguishes knowledge from mere justified true belief. This latter problem has been a prominent topic in epistemology ever since Edmund Gettier raised convincing counterexamples against the *JTB* definition of knowledge (Gettier, 1963).

Pre-Gettier it seemed that truth, belief and justification were not only necessary but also jointly sufficient for knowledge. In his 1963 paper, however, Gettier raised two examples in order to show that justified true belief is not sufficient for knowledge. In these examples, an agent has justified true belief in some proposition p on the basis of a false, though justified, statement q that entails p. So even though the agent is justified in believing p (since it is entailed by her justified belief that q) and p is true, intuitively this is still not sufficient for knowledge (since q is false). The agent's belief in p is justified, but this justification is somehow not related to the truth of p in the right way.

Given the validity of the Gettier examples, knowledge apparently involves more than justified true belief. These examples are counter-examples to the *JTB* definition only given the assumption that 'knowledge that p' requires that an agent believes that p (is true) *for the right reasons*. One can simply reject the counter-example by denying the latter assumption. The history of mainstream epistemology, however, shows that philosophers have not been willing to take this route. Rather than discrediting the *JTB* definition altogether, the Gettier examples have induced an extensive body of literature attempting to overcome the problems illustrated by them, thereby further clarifying the concept of knowledge. There are multiple ways to evade this type of counterexample

[3] Justification can be analyzed in multiple ways–e.g. as having sufficient evidence, as coherence with other beliefs, as a reliable process, etc. (Hendricks, 2006).

while still maintaining that knowledge is some form of justified true belief. The typical strategy is to either add a fourth knowledge condition or make the justification condition more fine-grained (Hendricks, 2006).[4]

2.1.2 Fallibilism

An implicit assumption of epistemology is that we are in pursuit of knowledge (Hendricks, 2006). In order to effectively pursue this goal, we must be able to distinguish truth from falsehood and use this ability when forming beliefs and engaging in inquiry. As such, the possibility of knowledge appears to be dependent upon our ability to distinguish truth from falsehood *without error*. Yet, given that our beliefs are undeniably fallible and given that we are hardly ever (perhaps never) able to fully eliminate the possibility of error, the skeptic argues that we therefore cannot acquire knowledge: knowledge is an unattainable goal. Thus we are forced to conclude that knowledge is not possible. Given that *all* our beliefs may still be false (for we can never fully exclude the possibility of error), belief can never meet the epistemic standards of justification needed for knowledge.

For any analysis of knowledge that a philosopher can come up with, there is a skeptical response starting with "but..." - filling in the dots with some possibility of error: but, what if there are hallucinogens in the water, do you really know that there is a cat on the mat in front of you? But, have you excluded the possibility that you are a brain in a vat? But, have you excluded the possibility that the cat is a hologram? Or that it is a cleverly disguised dog? To quote Lewis,

> "Let your paranoid fantasies rip – CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself – and soon you find that uneliminated possibilities of error everywhere. Those possibilities of error are far-fetched, of course, but possibilities all the same. They bite into even our most everyday knowledge. We never have infallible knowledge." (Lewis, 1996: 549)

Lewis, together with many other philosophers, maintains that it is a Moorean fact that we know all sorts of things. Moorean facts, named after the philosopher G. E. Moore, are facts that anyone in her right mind simply cannot deny.[5] To use Lewis' words again, it "is one of those things that we know better than we know the premises of any philosophical argument to the contrary". It is a Moorean fact that I know that I have two hands. Similarly, it is a Moorean fact that I know what I am wearing at the moment, and what day of the week it is. To doubt these facts, let alone deny

[4]Notice that the Gettier-examples are based on two common assumptions about epistemic justification: namely, that an agent can be justified in believing a false proposition (e.g. Smith's belief that q) and that justification is preserved by deductive inference (e.g. by Smith's inference from q to p via the introduction rule for the existential quantifier) (Gettier, 1963).

[5]Moore famously argued against skepticism by appealing to common sense. For instance, according to his reasoning, I can now prove that two human hands exist. How? By holding up my two hands and saying, as I make a gesture with the right hand, "Here is one hand," and adding as I make a gesture with the left, "and here is another." (Moore, 1959: 145-6).

them, “in any serious and lasting way would be absurd” (Lewis, 1996: 549). Such doubt goes against common sense and our pre-analytic intuitions. Note that the data of analytical philosophy, *viz.* the facts or intuitions that any satisfactory theory must be able to accommodate, consists of common sense and scientific assumptions about our knowledge, including Moorean facts. These assumed-to-be-known facts are not limited to mundane facts, though knowledge of the mundane is perhaps most intuitive or in any case least controversial. If knowledge that p required that all possible scenarios in which p is false are excluded and thereby any uncertainty about the truth of p removed (as the skeptic maintains), then we would indeed know nothing. But we know all sorts of things. So knowledge is possible without the need to rule out all possibilities of error. This view of knowledge is called *fallibilism*.

2.1.3 Modal conditions on knowledge

In the contemporary epistemology literature, knowledge is often defined in modal terms, that is, with respect to other possible worlds, scenarios or states. Intuitively, a possible world, counterfactual scenario or state represents a way the world might have been.[6] We are in the actual world, and as epistemic agents, we seek to know what facts obtain here. Put differently, for an agent to know that p, p must be true in the actual world. Additionally, in order to be justified in her belief, she must have excluded the possibility that she is in error about the truth of p. The (actual) world might have been different in many different ways, some more farfetched than others. There may be possible scenarios in which the agent has the exact same evidence as in the actual world (say, a scenario in which she is deceived by an evil demon), and which possibilities of error she therefore cannot eliminate.

So-called “relevant alternatives” theories of knowledge approach this matter by claiming that ‘knowledge that p’ only requires that all *relevant* possibilities of error are excluded by the agent. An agent need not be infallible with respect to all possible worlds, but rather only with respect to a restricted set of such worlds, namely, the epistemically relevant alternatives. For an agent to know that p, she must have eliminated all possibilities in which not-p, except for those possibilities that she may properly ignore because they are not epistemically relevant (Lewis, 1996).[7] Competing theories draw the distinction between epistemically relevant and irrelevant worlds differently. Such theories include epistemic contextualism as defended by e.g. Lewis (1996), Dretske (1970) and DeRose (1995); counterfactual epistemology as defended by Nozick (1981) and Sosa (2004). Epistemic contextualism, for instance, is the view that the set of epistemically relevant possibilities (those possibilities that may not be properly ignored) is dependent upon the context of knowledge ascription: this set is determined by the needs of the

[6]Thus, the terms “possible worlds”, “scenarios” and “states” can be interpreted as referring to alternate realities, scenarios, contexts or simply counterfactual circumstances at the actual world. It does not matter. For the present purpose, we can stick with the intuitive characterization: a possible world is a way the world might have been.

[7]Note that the actual world is by definition epistemically relevant: singling it out from all other possibilities is the goal of epistemic inquiry.

situation in which the knowledge ascription is made. As such, different contexts come with different standards for knowledge, depending on many factors–e.g. the cost of error (Lewis, 1996).

2.1.4 Social epistemology

The traditional subjects of epistemology are individuals – in particular, individuals considered in isolation. The JTB analysis and its post-Gettier descendants provide necessary and sufficient conditions for an agent to *possess* knowledge that p. It has long been recognized that knowledge, especially knowledge *aquisition*, has important social aspects, most notably because testimony provides agents with social evidence for their beliefs (Goldman, 2001). In recent years, the idea that groups of agents, and even social systems, can also be proper subjects of epistemology has gained increasing attention. Philosophical investigations of the social aspects of knowledge are generally categorized as "social epistemology" (Goldman, 2009). Part of this discussion addresses the possibility of "epistemic group agents", in addition to the individual knowers from traditional epistemology. Various authors argue that groups should be considered as epistemic subjects in their own right (including Gilbert, 1987; Pettit, 2011; Rolin, 2008; Wray 2007). Often this is motivated by the observation that important epistemic phenoma cannot be properly explained in terms of individual knowledge, such as collaborative knowledge that cannot be attributed to any individual agent.

2.2 Epistemic logic

Epistemic logic is generally said to have started with Hintikka's *Knowledge and Belief* (Hintikka, 1962). In his book, Hintikka proposed to treat knowledge as a modal operator, similar to the modal operator for necessity, which can be given an interpretation in terms of standard Kripke semantics. An approach to formal epistemology in terms of Kripke semantics (also called "possible worlds semantics") is especially well-aligned with the proposals from mainstream epistemology that define knowledge in modal terms. Both these approaches define knowledge as truth in all epistemically possible worlds, *viz.* as truth in all worlds that an agent cannot distinguish from the actual world (and from each other) on the basis of her knowledge.

2.2.1 Kripke semantics for epistemic logic

The language of basic epistemic logic is obtained by adding an operator for knowledge, K_a, to the language of propositional logic. This operator has the following intended meaning:

- $K_a\varphi$ (agent a knows that φ)

Language and syntax The language of basic epistemic logic L_K has a countable set of propositions p, Boolean operators $\neg$ and $\wedge$ and a modal operator K_a. It has the following syntax:

$$\varphi \;:=\; p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi$$

The other Boolean connectives get their standard abbreviations in terms of $\neg$ and $\wedge$: thus, $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$ and $\varphi \to \psi := \neg(\varphi \wedge \neg\psi)$.

Semantics The language is interpreted on Kripke models (here called "epistemic models").

Definition 2 (Multi-agent epistemic model)**.** A multi-agent epistemic model is a tuple $\mathbf{S} = (S, \to_{a(a \in A)}, \| \bullet \|)$ consisting of a finite set S of states (or "possible worlds"); a set A of agents; for each agent a, a (binary) reflexive, transitive relation ("epistemic indistinguishability" relation) $\to_a \subseteq S \times S$; and a valuation which maps the atomic sentences $p \in \Phi$ to sets of worlds $\|p\| \subseteq S$.

A proposition is a set $P \subseteq S$ of worlds. Further standard notation includes: $\neg P := S \setminus P$ is the negation (complement) of proposition P, $P \wedge Q := P \cap Q$ is the conjunction (intersection) of P and Q, $P \vee Q := P \cup Q$ is their disjunction (union), $\top := S$ is the tautologically true proposition and $\bot := \emptyset$ is the inconsistent proposition, etc. Regarding the standard operations on relations, $R_1, R_2 \subseteq S \times S$: union of relations is $R_1 \cup R_2$, intersection of relations is $R_1 \cap R_2$, relational composition is $R_1 R_2 = \{(s,t) \in S \times S : \exists w \in S\, (s,w) \in R_1 \wedge (w,t) \in R_2\}$, the n^{th} iteration of a relation is R^n (which is defined recursively by putting $R^0 = id := \{(s,s) : s \in S\}$ to be the identity relation and $R^{n+1} = R^n R$) and the reflexive-transitive closure of a relation is $R^* := \bigcup_{n \in N} R^n = id \cup R \cup R^2 \cup R^3 \cup \cdots$.

The epistemic indistinguishability (or possibility or accessibility) relation represents agent a's epistemic uncertainty: two states s and t are related by $\to_a$ (and thus $s \to_a t$) whenever at s agent a cannot distinguish s from t on the basis of her knowledge (at s). Since s and t are both compatible with her knowledge, a has no other epistemic means by which to distinguish them from each other. That is, all states t' such that $s \to_a t'$ are epistemic alternatives (or possibilities) for a, because from the perspective of her knowledge, she cannot distinguish them from s. Note that multi-agent epistemic models include for each agent $a \in A$ her own accessibility relation. On the syntactic side, each agent $a \in A$ gets her own knowledge operator K_a.[8]

The semantics is given by an interpretation map that associates each formula of the language with a proposition $\|\varphi\|_{\mathbf{S}} \subseteq S$ in models $\mathbf{S}$. Intuitively, $\|\varphi\|_{\mathbf{S}}$ is the set of all worlds in S satisfying φ. The definition is by induction, in terms of the obvious compositional clauses (using the operators defined above):

Definition 3 (Standard Kripke semantics)**.** Given a model $\mathbf{S}$ and a world s,
$s \models_{\boldsymbol{S}} p$ iff $s \in \|p\|_{\mathbf{S}}$
$s \models_{\boldsymbol{S}} \neg\varphi$ iff $s \notin \|\varphi\|_{\mathbf{S}}$
$s \models_{\boldsymbol{S}} \varphi \wedge \psi$ iff $s \models_{\boldsymbol{S}} \varphi$ and $s \models_{\boldsymbol{S}} \psi$
$s \models_{\boldsymbol{S}} K_a\varphi$ iff $\forall s' : s \to_a s' \Rightarrow s' \models_{\boldsymbol{S}} \varphi$

[8]More precisely, an epistemic logic for multiple agents is obtained from single-agent logics, which are then combined into a fusion of logics, i.e., one big logic.

In words, formulas $K_a\varphi$ get the following semantic interpretation: φ is true at all worlds s' that a cannot distinguish from s on the basis of her knowledge.

For any binary relation $R \subseteq S \times S$ on the set S of all possible worlds, the corresponding Kripke modality $[R]$ can be introduced as follows:

$$[R]P = \{s : \forall t \in S(sRt \to t \in P)\}$$

The knowledge operator corresponds to the Kripke modality for the epistemic indistinguishability relation:

$$K_a P = [\to_a]P$$

The knowledge relation is a primitive component of epistemic models. This means that knowledge is not defined in terms of some other, more basic, notion such as belief.

Kripke semantics has the feature that conditions on the accessibility relation of models correspond to formulas that are valid with respect to these models. These formulas describe properties of the modal operator that correspond to the Kripke modality for the relation. For epistemic models this means that constraints on $\to_a$ correspond to axioms that describe properties of knowledge. The conditions imposed on the accessibility relation $\to_a$ of the epistemic models assumed in this thesis are reflexivity and transitivity.

Reflexivity corresponds to the **T** axiom: $K_a\varphi \to \varphi$. This axiom states that knowledge is factive: knowledge implies truth.[9] It is commonly assumed that this condition is a minimal requirement for a logic of knowledge, as factivity is supposed to be one of the defining properties of knowledge (*cf.* chapter 2.1.1). Transitivity corresponds to the **4** axiom: $K_a\varphi \to K_aK_a\varphi$. This axiom states that epistemic agents are positively introspective: if a knows that φ, then she knows that she knows that φ. Together with the **T** axiom, it leads to a notion of knowledge that is veracious and that is only held by positively introspective agents. Transitivity thus corresponds to the assumption that epistemic agents have introspective access to their knowledge. While this assumption is not as uncontroversial as the assumption that knowledge is factive, many philosophers seem to accept it.

Tautologies are valid on all Kripke frames, irrespective of any conditions imposed on the accessibility relation.[10] Epistemic logic thus models agents that know all tautologies. Similarly, for the **K** axiom: $K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$. This axiom states that agents know the consequences of their knowledge. The agents thus modeled are perfect reasoners or *logically omniscient* (Baltag *et al.* 2008: 28). The system that validates axioms **K**, **T** and **4** is known as $S4$. These axioms are valid on the epistemic models that I assume in this thesis.

A more dubious knowledge property, especially from a philosophical perspective, is negative introspection: if a does not know that φ, then she knows that she does not know

[9]It is only valid when interpreted on models in which the epistemic relation is reflexive. So if reflexivity is not imposed as a condition on the epistemic relation, then knowledge is not by definition factive.

[10]A *frame* is the same type of structure as a model minus the valuation function.

that φ. This property is described by the **5** axiom: $\neg K\varphi \rightarrow K\neg K\varphi$. It effectively states that agents have knowledge of their lack of knowledge. This axiom correponds to the condition that $\rightarrow_a$ is Euclidean. Note that in the context of reflexivity and transitivity, the semantic condition that is necessary to validate axiom **5** results in $\rightarrow_a$ being an equivalence relation– hence reflexive, transitive and symmetric. In this context, axiom **5** really encodes the assumption that knowledge is *fully* introspective: epistemic agents have full introspective access to their knowledge and the lack thereof.

As is often noted, this property does not appear to characterize the concept of knowledge that philosophers have in mind. If epistemic agents were fully introspective, then they would know of every proposition whether they know it or not. In other words, they would be certain about the extent of their knowledge. However, it seems reasonable to assume that epistemic agents are able to consistently believe that they know p, even when in fact they do not (because p is false or because their justification is insufficient). Moreover, fallible knowledge is typically not fully introspective, for it does not require agents to exclude all possibilities of error (*cf.* 2.1.2). For example, if agent a cannot distinguish state s from states s', then her knowledge at s is given by the propositions that s and s' agree on, and her lack of knowledge is given by the propositions that s and s' do not agree on. Full introspection implies that if at s she knows p, then at all s' she also knows p. Similarly, if at s she does not know p, then she also does not know p at s', which means that she knows that she does not know p, and therefore she cannot consistently believe that she knows p.

This said, the most common logic for knowledge is the modal system S5. S5 models validate axioms **K**, **T**, **4** and **5**. Formal approaches to epistemology – such as game theory and computer science – typically assume the S5 conditions for knowledge, which is (partly) explained by the convenient formal properties of the logic. Philosophers typically opt for a weaker notion. Hintikka (1962), for instance, argues that the proper logic for knowledge is the modal system S4. In this thesis, I follow Hintikka and assume that knowledge satisfies the S4 conditions, in order to avoid the assumption that knowledge is fully introspective. In the remainder of this thesis, I therefore only consider epistemic models that are positively introspective. It should be noted that the S5 conditions can be obtained as a special case by interpreting the language on fully introspective models:

Definition 4 (Special case: fully introspective agents)**.** An epistemic model is *fully introspective* iff all epistemic relations $\rightarrow_a$ are equivalence relations.

2.2.2 Common knowledge and distributed knowledge

Given multi-agent epistemic models, notions of group knowledge can be defined in terms of the knowledge relations of the individual agents by constructing group knowledge relations from the individual knowledge relations. Here, I shall only introduce the formal definitions for "common knowledge" and "distributed knowledge". The notion of common knowledge captures the knowledge that a group of agents has whenever all group members know that p and know of each other that they all know that p, and know of each other that they know of each other that they all know that p, *ad infinitum*.[11] It has proven particularly relevant for analyzing scenarios and puzzles that involve coordination amongst individuals within groups, as in the well-known 'muddy children' and 'cheating wives' puzzles. The notion of distributed knowledge is typically taken to represent the knowledge that the agents in a group would know were they to combine their knowledge (Halpern and Moses, 1992). That is to say, a group has distributed knowledge that p, whenever p is entailed by the (combined) knowledge of the group members. This means that p can be distributed knowledge in G without it being the case that any member of G knows that p (Baltag *et al*, 2008).

"Common knowledge" and "distributed knowledge" typically get their own modal operators, which are added to the standard language of epistemic logic: Ck_G and Dk_G, respectively, where $G \subseteq A$ are groups of agents. These modalities are then given the following intended meaning: $Ck_G\varphi$ is read as 'it is common knowledge in G that φ' and $Dk_G\varphi$ is read as 'it is distributed knowledge in G that φ'. Given a multi-agent epistemic model, the common knowledge Ck_G of a group $G \subseteq A$ of agents corresponds to the following Kripke modality:

$$Ck_G P = [(\bigcup_{a \in G} \rightarrow_a)^*]P.$$

Here, R^* is the reflexive-transitive closure of relation R. Common knowledge can be alternatively expressed as an infinite conjunction of iterated knowledge (about others' knowledge) within the group G:

$$Ck_G P \iff \bigwedge_{a \in G} K_a P \wedge \bigwedge_{a,b \in G} K_a K_b P \wedge \ldots$$

The distributed knowledge Dk_G of a group $G \subseteq A$ of agents is given by the Kripke modality for the relation $\rightarrow_G$, where $\rightarrow_G := \bigcap_{a \in G} \rightarrow_a$ (for any group $G \subseteq A$) :

$$Dk_G P = [\rightarrow_G]P.$$

The semantics for these modalities is as expected, thus:

Definition 5. Given a model $\mathbf{S}$ and a world s.
$s \models_{\mathbf{S}} Ck_a\varphi$ iff $\forall s' : s(\bigcup_{a \in G} \rightarrow_a)^* s' \Rightarrow s' \models_{\mathbf{S}} \varphi$
$s \models_{\mathbf{S}} Dk_a\varphi$ iff $\forall s' : s(\rightarrow_G) s' \Rightarrow s' \models_{\mathbf{S}} \varphi$.

[11] The notion of 'common knowledge' traces back to the work of David Lewis, who investigated it in the context of analysing the notion of a convention (Lewis, 2008).

Chapter 3

Conceptions of group knowledge

Formal epistemology intends to represent philosophical concepts. By encoding and modeling the formal properties of such concepts, it hopes to provide insight into the implications of the various concepts individually and their relation to each other. The standard group-epistemic notions from epistemic logic, however, are not well-aligned with the current philosophical discussion on group knowledge, as they fail to address some important features of philosophical concepts. This is probably partly attributable to the fact that the formal notions have broader roots, such as in computer science, and are often driven by other (non-philosophical) considerations.

This chapter provides a brief discussion of some of the important choices or properties that distinguish competing philosophical notions of group knowledge from each other.[1] In the final section, I look at the notions of common knowledge and distributed knowledge and consider to what extent they exhibit these properties, and highlight properties that should ideally be represented by formal notions. In a later chapter, these observations are used to motivate an alternative formal definition of group knowledge, called *collective knowledge.*

3.1 Group knowledge properties

In the philosophical literature, the term "group knowledge" is used to refer to several concepts, embodying different views of e.g. the extent to which knowledge must ultimately be held by an individual, the subjects of group knowledge, and the extent to which group knowledge is accessible to the individual group members.

3.1.1 Summativism and non-summativism

Interpretations of "group knowledge" are often classified based on whether group knowledge is seen to be held by individual group members or by the group in its own right. Probably the most conservative interpretation of the term is as shorthand for claims such as "everybody knows" or "someone in the group knows". Such an interpretation reduces ascriptions of group knowledge to ascriptions of knowledge to individual group members. This view of group knowledge is known in the literature as *summativism*, introduced by Anthony Quinton (1976). He gives the following explanation:

> Groups are said to have beliefs, emotions and attitudes and to take decisions and make promises. But these ways of speaking are plainly metaphorical. To

[1] Several authors actually use the term "collective knowledge" rather than "group knowledge". I shall use the term "group knowledge" as the umbrella term, and reserve "collective knowledge" for the notion proposed in this thesis.

ascribe mental predicates to a group is always an indirect way of ascribing such predicates to its members. (Quinton, 1976: 19)

If all knowledge can be analyzed as individual knowledge, then there does not seem to be a *prima facie* reason to assume that groups can have knowledge, for 'their' knowledge can already be explained at the individual level. Moreover, as Christian List puts it, if groups are considered capable of knowledge "one has to be prepared to consider groups as epistemic agents over and above their individual members" (List, 2011: 223). Philosophers that do not want to commit themselves to the metaphysical baggage associated with group agency opt for a summative understanding of group knowledge.[2]

This said, in recent years the idea that groups can be treated as collective agents that are capable of knowledge has gained increasing attention (Bratman, 1993; Corlett, 1996; Gilbert, 1987; Gilbert, 1992; List and Pettit, 2011; List, 2011; Rolin, 2008). Within this discussion, group knowledge is typically argued to be non-summative, meaning *not* fully reducible to the knowledge of the group members. This presupposes that groups can somehow have knowledge and "minds of their own" (Pettit, 2010). Different conditions have been considered that reflect structural properties of groups that are deemed necessary (and sufficient) for collective epistemic agency (Pettit, 2010; Wray, 2001). A common objective of non-summative analyses has been to explain the status of certain instances of "collaborative" knowledge that cannot adequately be explained as (or reduced to) individual knowledge. The prime example of such irreducible knowledge is scientific knowledge. Modern science is undoubtedly highly collaborative. Indeed, even for most of their individual knowledge scientists are dependent upon the knowledge of others. This type of dependence is often referred to as *epistemic dependence* (Hardwig, 1985; 1991). Given such dependence, it seems natural to assume that the knowledge obtained through scientific practice should be attributed to groups in their own right rather than to individuals – more so since oftentimes no individual scientist appears to meet the conditions necessary for this assumed-to-be knowledge, thus necessitating a non-summative concept (de Ridder, 2014; Rolin, 2007; Wray, 2007). For example, K. Brad Wray (2007) gives the following argument in favor of non-summative group knowledge:

> Collective knowing may be the only way to get at some of the knowledge we now take for granted. Indeed, now that such discoveries have been made they can in principle be known by individuals as well. But, some of the knowledge we now take for granted could never have been discovered without the efforts of plural subjects, agents formed by people working in groups with intentions that are irreducibly the intentions of the group. (Wray, 2007: 345)

Note that summativism need not deny that groups are needed to acquire (scientific) knowledge, but what it denies is that consequently groups have (scientific) knowledge as distinct from their individual members (Fagan, 2012).

[2]These include Giere, 2002; Goldman, 1999; Kitcher, 1994.

To many philosophers, the possibility of "proper" group knowledge stands or falls with the possibility of group beliefs. Summativism is characterized by a denial of this latter possibility,[3] and so also of group knowledge–at least of the proper kind, *viz.* as distinct from the knowledge distributed in, or shared by, random (unstructured) sets of individuals. Yet, summativism need not assume that any group of agents has group knowledge. J. Angelo Corlett (2007), for example, argues for so-called "sophisticated summativism". According to his analysis, group knowledge requires that the individual agents' beliefs are sufficiently similar in content and supported by shared motives. So for those philosophers who do not want to attribute epistemic states to groups directly, there are still concepts of group knowledge that are richer, i.e. philosophically more interesting, than plain summativism.

3.1.2 Epistemic groups

Group knowledge is often defined in terms of the types of groups that are its possible subjects. The underlying assumption is that random sets of individuals cannot have group knowledge simply because they should not be considered possible epistemic subjects (Corlett, 2007; Pettit, 2011).[4] Groups that qualify as epistemic groups or *plural subjects* must at least be partly defined on the basis of epistemic properties related to knowledge possession. This means that the concept of group knowledge is based on some concept of shared epistemic properties – such as shared belief or shared justification – that ties the group members together from an epistemic perspective. For example, the individuals that are currently in New York City can clearly be distinguished as a set based on their location and there is undoubtedly knowledge distributed in this set. Yet, location is not an epistemic property and so, from an epistemic perspective, it appears that this set should not be referred to as a group that can possess knowledge.

Philosophers have considered several conditions (or epistemic group properties) in terms of which to define epistemic groups. Here, I briefly explain four prominent proposals that address this matter. To begin, Christian List argues that epistemic groups must be characterized by "an institutional structure (formal or informal) that allows the group to endorse certain beliefs or judgments as collective ones" (List, 2011: 223). Such an institutional structure (e.g. an electoral system) is needed in order for the group members to support certain information as collective information. He proposes to represent such institutional structures by aggregation procedures (List and Pettit, 2002). Second, Corlett (2007) appears to support a similar view. He argues that epistemic groups consist of members who identify relationally with each other (as group members) based on the fact that they have shared epistemic motives and decision-making capacitities that enable them to form beliefs (Corlett, 2007: 232, 235). Corlett explains,

> But this group [of television watchers] is so amorphous that its putative beliefs "held" by various segments of the group – that, for example, "Frontline",

[3] Though the converse does not hold.

[4] Epistemic subjects exhibit behavior that can be evaluated from an epistemic perspective (Goldman, 2004).

"Nova," and certain other programs are qualitatively superior to others – are held in common more accidentally than as a group or a sub-group intention. (Corlett, 2007: 233)

In other words, epistemic groups must be more than merely sets of individuals.

Third, Wray (2007) argues that this is not enough: only groups with so-called "organic solidarity" can have group knowledge. Such groups are characterized by functional interdependence of the group members, which in turn requires both a common goal and an agreed upon division of labor needed in order to achieve this goal.[5] "Organic solidarity" is contrasted with "mechanical solidarity", which only requires shared beliefs and motives. Groups of the latter type can be said to share knowledge (in the summative sense), though they are incapable of group knowledge because their knowledge cannot go beyond the sum of their individual knowledge (Wray, 2007: 342). Wray provides three conditions for epistemic groups: (1) epistemic dependence, (2) an agreed upon division of labor and (3) the ability to adopt views that are not necessarily identical to the views of their members.

As a final example, Rolin (2008) argues that epistemic groups consist of individuals that are jointly committed to defend background assumptions – so-called "default entitlements" – that constitute a context of epistemic justification within which the group members work. While every group member is thus committed, the burden of proof (or epistemic responsibility) for these default entitlements is distributed amongst them. This allows members to acquire further knowledge based on entitlements (i.e., group knowledge) even when they are not themselves able to defend them. Moreover, the joint commitment ensures that the group members are aware of their group's entitlements (Rolin, 2008: 121-2). As such, the value of epistemic groups is that they allow their members to share the epistemic responsibility for each other's individual knowledge. Note that Wray and Rolin both defend non-summative concepts of group knowledge, though not the same concept: Wray is defending a type of group knowledge that is only acquired by groups, and that thus necessitates epistemic groups, while Rolin focuses on the background assumptions that agents within the group can rely upon when acquiring additional *individual* knowledge.

3.1.3 Accessibility of group knowledge

The final important aspect of the concept is the extent to which group knowledge is accessible to the individual members. There seems to be an implicit assumption in the philosophical literature that the group members should be able to come to know their group's knowledge (De Ridder, 2014; Gilbert, 1987; Goldman, 2004; List, 2011; Wray, 2007). To "non-summative" philosophers group knowledge is valuable exactly because it offers a path to individual knowledge. Recall, for example, the quote from Wray (2007) in which he says that "[c]ollective knowing may be the only way to get at some of the knowledge we now take for granted. Indeed, now that such discoveries have been

[5]Wray follows Emile Durkheim (1893) in making this distinction between types of groups.

made they can in principle be known by individuals as well".[6] Similar considerations apply for summative concepts: in order for group knowledge to be valuable to individual group members, they must have access to it (at least in principle) (Goldman, 2004). In the formal-epistemic literature, this assumption is known as the *principle of full communication* (Van der Hoek *et al.* 1999).[7]

3.2 Common knowledge and distributed knowledge

In order to be pertinent to philosophy, the notions offered by formal epistemology should represent philosophical concepts. For the concept of group knowledge this means that a formal definition should encode *how* group knowledge depends on the knowledge of its constituents, and also *how* the group members are tied together through a shared epistemic property. In this section, I assess how the properties discussed above are reflected in the notions of common knowledge and distributed knowledge.

Starting with common knowledge, its standard formal definition allows random sets of individual agents to have common knowledge. Given any set of agents with knowledge, its common knowledge is computed from the individual knowledge relations. Still, arguably such knowledge presupposes an epistemic group since from a conceptual perspective the group members are not tied together directly by the knowledge that they share, but indirectly via their interest in each other's knowledge. Recall that a group is said to have common knowledge that p whenever every group member knows that p, every group member knows that every group member knows that p, every group member knows that every group member knows that every group member knows that p, etc. This means that the group members have iterated higher-order knowledge of each other's knowledge. As such, on a conceptual level, common knowledge presupposes a group of agents that is at least tied together by agents' mutual interest in each other's knowledge.[8] This mutual interest is typically explained with reference to a common goal, such as coordinated action or reaching agreement within a group (Fagin *et al.* 1995; Lewis, 2008). In order for common knowledge to facilitate coordinated action, clearly all group members must know that it is common knowledge, and know that all others' know that it is common knowledge.[9]

Distributed knowledge is a well-studied notion from epistemic logic that has remained largely disconnected from the philosophical literature. It is typically taken to represent the knowledge that the agents in a group would possess were they to combine their

[6] Similarly, Rolin motivates her notion of group knowledge by drawing attention to its value for individual group members: they need to be aware of their group's knowledge in order to acquire (individual) knowledge based upon it.

[7] Van der Hoek *et al.* claim that group knowledge should have this property, stating that: "It is questionable whether group knowledge [i.e., distributed knowledge] is of any use if it cannot somehow be upgraded to explicit knowledge by a suitable combination of the agents' individual knowledge sets, probably brought together through communication" (Ibid.: 226).

[8] This is not represented in basic epistemic models.

[9] Common knowledge should thus be positively introspective – which it is. Any infinite conjunction of iterated knowledge will be positively introspective, due to the assumption that knowledge is veracious.

knowledge (Halpern and Moses, 1992).[10] Distributed knowledge is often referred to as *implicit* knowledge, as it need not be held by any of the group members (Fagin and Halpern, 1988; Levesque, 1984).[11] The notion of implicit knowledge plays a similar role to that of non-summative group knowledge. As said, "non-summative" philosophers typically have some particular piece of (presumed) knowledge in mind that is not attributable to any individual. In order to defend the status of this presumed knowledge, they propose to view (particular types of) groups as epistemic agents (in addition to the individual agents from traditional epistemology). Similarly, knowledge that is not held by any individual can be explained as implicit knowledge. Distributed knowledge thus need not be interpreted non-summatively.[12]

Any set of individuals has distributed knowledge. As such, the notion appears to lack an important group knowledge property: its application is not restricted to epistemic groups, however construed. As a concept of group knowledge, distributed knowledge is based on the (often) implicit assumption that it is accessible, at least in principle, to the individual group members (Van der Hoek *et al.* 1999). It appears to be supposed that this is enough for distributed knowledge to be useful to the group members. Arguably, this type of knowledge is of use to random sets of individuals as these individuals may still acquire knowledge through each other's testimonies. Nonetheless, groups that qualify as epistemic groups (from a philosophical perspective) must at least be partly defined on the basis of epistemic properties related to knowledge possession. This is what makes its knowledge *group* knowledge, as opposed to e.g. knowledge that just happens to be distributed within a group.

In providing an analysis of group knowledge, two important aspects of notions of group knowledge merrit clarification: the types of groups that are subjects of group knowledge, and given such groups, the extent of their knowledge. The philosophical discussion has tended to focus on the former aspect, whereas the standard formal definitions of group knowledge address only the latter aspect. In this thesis, I propose a formal definition of group knowledge that addresses both of these important aspects. In particular, I present a formal definition of group knowledge, called *collective knowledge*, that is based on the notion of distributed knowledge, and that has the following additional property: group knowledge is about some common issue. Common issues, I argue, are only held by epistemic groups, which are groups of individuals that are tied together through mutual interest in each other's knowledge and questions. As such, *collective knowledge* is based on the assumption that the epistemic property that is pertinent to "epistemic group agency" is mutual interest of the group members in each others' knowledge and questions. *Collective knowledge* is based on an additional knowledge property that is not captured by the standard semantics for knowledge, namely, that all knowledge of an agent is an answer to her question(s). The first task, then, is to explain in what sense knowledge and questions are tied together. This is the topic of the next chapter.

[10] Here, the term "group" is used loosely. Groups of agents are not distinguised from sets of agents.

[11] Fagin and Halpern originally named the notion "implicit knowledge" rather than "distributed knowledge" (Fagin and Halpern, 1988).

[12] Virtually no formal epistemologist does this anyway.

Chapter 4

Epistemic issues and knowledge

It is generally recognized that questions are relevant to epistemology. In both the philosophy and formal logic literature there are proposals that address the role of questions in epistemology. These proposals typically stress the role of questions as cognitive goals that motivate and regulate inquiry. Hintikka's "Interrogative model of inquiry" and his so-called "Socratic epistemology" are prime examples of such an approach to epistemology (Hintikka, 1981; 2007). Yet, a few exceptions notwithstanding, this recognition has not motivated philosophers to propose (formal) accounts of knowledge that include questions as properties of knowledge. Certainly, questions are not considered a necessary condition for knowledge by the standard definitions of knowledge from mainstream epistemology. This is perhaps explained by philosophers' focus on knowledge possession, rather than acquisition, as exemplified by the traditional *JTB* analysis, and their focus on demarcating knowledge possession from "mere" belief possession. The *JTB* conditions do not require that known propositions are answers to the knower's questions. In fact, known propositions need not be relevant to anything whatsoever.[1]

In this chapter, I argue that all knowledge implies a question. In order to model this property, I propose a condition to be imposed on the accessibility relation of models for knowledge. The *epistemic issue models* of van Benthem and Minică (2009) are used as the starting point. In the final section of this chapter, I introduce additional terminology that shall be used in later chapters.

4.1 Questions

A number of (formal) epistemologists have explicitly linked questions to knowledge. They have interpreted the term e.g. (1) as specifying the relevant alternatives that an agent must rule out in order to have knowledge (Schaffer, 2007); (2) as requests for information (Groenendijk, 1999; Groenendijk and Stokhof, 1984), and; (3) as epistemic goals that together comprise the agent's so-called *research agenda* (Olssen and Westlund, 2006). In this thesis, I adopt a different interpretation. I interpret the term as representing the distinctions that define an agent's conceptual framework. A question consists of a family of answers, which represent states that the agent can (conceptually) distinguish

[1] A second assumption ingrained in mainstream epistemology is that in its most fundamental form, knowledge is propositional knowledge, which is taken to express a binary relation: s knows that p. As Schaffer (2007) explains, since propositional knowledge need not explicitly refer to a question, there appears to be no need to include questions in the knowledge definition. It should be noted that Schaffer emphasizes that he has not found any support for this assumption (*viz.*, that knowledge expresses a binary relation), other than the fact that 's knows that p' seems to express a binary relation (Schaffer, 2007: 385, 400).

from each other. Thus understood, having a question means making distinctions. For a discussion of these other treatments of questions, see appendix A.

Following Groenendijk and Stokhof (1984), questions are taken to denote partitions of the state space into cells, such that each cell corresponds to a possible answer (proposition) to the question. Questions are not, however, explicitly part of the semantic models that I propose. Rather, I propose to model questions indirectly via an epistemic relation. This relation is meant to encode roughly the learnable answers to an agent's questions. In order to do so, I start from the epistemic issues models of Van Benthem and Minică (2009). Van Benthem and Minică identify questions with their corresponding equivalence relations, called "issue relations", which they then add to the standard epistemic models (*cf.* section 2.2). In their models, both the knowledge-relation and the issue-relation are equivalence relations. In this thesis, however, I only require that these relations are reflexive and transitive.

Definition 6 (Epistemic issue model)**.** Given a set A of agents and a set Φ of atomic sentences, an epistemic issue model over (A, Φ) is a tuple $\mathbf{S} = (S, \rightarrow_{a(a \in A)}, \approx_{a(a \in A)}, \| \bullet \|)$ consisting of a finite set S of states; for every agent a, a reflexive, transitive relation $\rightarrow_a \subseteq S \times S$; for every agent a, a reflexive, transitive relation $\approx_a \subseteq S \times S$; and a valuation which maps the atomic sentences $p \in \Phi$ to $\|p\| \subseteq S$.

It should be emphasized that, unlike Van Benthem and Minică, I interpret the issue-relation $\approx_a$ as an additional epistemic relation, rather than as a question relation. As an epistemic relation, it encodes the knowledge that an agent could acquire based on answers to her questions.[2] As such, $\approx_a$ can be viewed as embodying "question-based potential" knowledge. While questions (as partitions) correspond to an equivalence relation, the knowledge that agents can acquire based on their questions need not satisfy the S5 axioms. Indeed, I only assume that $\approx_a$ is reflexive and transitive. This is based on the assumption that knowledge, including the knowledge that agents would possess had they answered all their questions, has the $S4$ properties but that it is not necessarily negatively introspective.

A related point is that the family of answerable questions (those questions that the agent could actually find an answer to) is not necessarily closed under negation. Some questions can only ever be answered if the answer is "yes" but not if the answer is "no" (and vice-versa, if the same question is stated in negated form). For instance, the question "am I deceived by an evil demon?" can only be decisively answered if the answer is "yes", and the evil demon chooses to reveal itself. If the answer is "no" and there is no evil demon, then the question will never be decisively answered, as no amount of evidence will ever be enough to exclude the positive answer. Thus, the knowledge that agents can acquire based on their questions is not necessarily partitional because not all answers can be known (that is, are learnable).

The semantics is as specified in chapter 2 (for epistemic models). Given an epistemic issue model, to make use of the issue-relation, a new modal operator is added to the language, Q_a, with the following semantics:

[2] The notation $\approx_a$ should not be taken to denote an equivalence relation.

Definition 7 (Semantics for Q_a)**.** Given a model $\mathbf{S}$ and a world s,
$s \models_{\mathbf{S}} Q_a\varphi$ iff $\forall s' : s \approx_a s' \Rightarrow s' \models_{\mathbf{S}} \varphi$.

Here '$Q_a\varphi$' is read as "φ can be known solely based on learnable answers to a's questions".[3] In other words, "the true answers to a's questions entail φ". Given this reading, Q_a is an epistemic modality, in addition to K_a. More precisely, Q_a is interpreted as the Kripke modality for the issue-relation:

$$Q_a P = [\approx_a]P.$$

The Q_a modality thus represents the maximum knowledge that an agent can acquire given her questions and the answers that are learnable for her.

4.2 Epistemic issues and knowledge acquisition

Questions (as partitions) are not part of the semantic models, but they nonetheless play an important role in the background. The term "questions" can mean many things. In this thesis, I interpret the term as encoding the relevant (conceptual) distinctions between possible states (or worlds) that an agent can (or is willing to) make. These distinctions define her conceptual framework. A question consists of a family of answers (propositions). States in the same answer are *conceptually indistinguishable* for the agent, and she therefore represents them as the same state. Thus, an agent's questions are determined by her ability to distinguish possible states from one another.

The set of distinctions that an agent can make, *viz.* the set of all her questions, is her most refined question or *issue*.[4] As said in the previous section, however, questions are *not* identified with the issue-relation ($\approx_a$) from the above models. The issue-relation is interpreted as an indistinguishability relation, in parallel to the epistemic indistinghuishability relation for knowledge, such that $s \approx_a t$ holds if at state s agent a cannot conceptually distinguish s from t, and thus represents them as the same state. As such, there is no answer that she can learn (at s) that will allow her to distinguish s from t (so she cannot learn that t is not the actual state). Thus understood, questions include answered questions. The underlying assumption is that the fact that an agent can distinghuish a state t from the actual state s based on her knowledge does not negate her ability to conceptually distinghuish them. If she knows at s that t is not the actual state, then she can distinghuish s from t, both conceptually and epistemically.[5] Thus, an answered question still remains a question in this sense. Further, answers are

[3] This interpretation differs from the interpretation adopted by Van Benthem and Minică, which is as follows: "φ holds in all issue-equivalent worlds" (Van Benthem and Minică, 2009: 4).

[4] So issues are just 'conjunctive' questions.

[5] In its common usage, "questions" refers to *open* questions. The interpretations of "questions" considered by Olsson and Westlund (2006) and the Inquisitive Semantics approaches focus on currently open questions at state s. Within the setting of epistemic issue models, locally open questions can be represented by the issue-relation as follows. For a given state s, the set of all states that are epistemically indistinguishable for a at s is given by $s(a) := \{s' \in S : s \rightarrow_a s'\}$. The restriction $\approx_a |s(a) := \approx_a \cap s(a)$ of the issue-relation $\approx_a$ to agent a's current epistemic state $s(a)$ represents agent a's 'question-based potential knowledge' at s.

closed under finite non-empty intersection, but they are not closed under subsets, as such subsets go beyond the conceptual distinctions made by the agent.

Based on this interpretation of "question", I claim that questions are a necessary condition for knowledge acquisition and thus that knowledge acquisition presupposes them. Knowledge is preceded by uncertainty (and thus an open question): an agent knows that p whenever she has eliminated her uncertainty regarding the truth of p. Epistemic uncertainty is possible only for agents that make distinctions between possible states: an agent is epistemically uncertain whenever she can distinguish state s from another state t, but does not know at s which of them is actuality. Indeed, without conceptual distinctions there is nothing to be uncertain about, and hence there is nothing about which to inquire. Epistemic agents aim to locate the actual world (or state) within their conceptual framework. An agent who makes no distinctions only considers one possible world, namely, the single world that she can conceive of. From her subjective perspective, she knows all there is to know and so she will inquire no further. Thus understood, questions are fundamental to epistemology, for without questions we cannot acquire any knowledge at all.

To be sure, the knowledge as possessed by an agent is the same in nature as the knowledge as acquired by her. There are not two different concepts at work here. So if knowledge as acquired presupposes a question, then knowledge as possessed must similarly presuppose a question. Since individual knowledge is a trivial case of group knowledge, i.e., of the singleton group, it must also hold for group knowledge.

I propose that knowledge has the following property:

Condition 1. If a knows that p, then a can know p solely based on answers to her question(s).

All knowledge is in terms of questions. In the language of epistemic logic, this can be expressed by the following formula:

$$(P1)\ K_a P \to Q_a P.$$

$P1$ states that all knowledge is based on answers to agents' questions: a knows that p only if p is entailed by an answer to her issue. Conversely, if p is not entailed by (or based on) an answer to a's question(s), then she does not know p.

In order to assume $P1$ as a knowledge property, an appropriate condition must be imposed on the epistemic issue models. $P1$ corresponds to the following semantic condition:[6]

$$\approx_a \subseteq \to_a .$$

In words, conceptual indistinguishability implies epistemic indistinguishability: if a does not conceptually distinghuish p from other possibilities, then she does not epistemically distinghuish p from them either–hence she cannot know p.

[6]By standard modal correspondence theory (Blackburn *et al.* 2001).

Van Benthem and Minică do not impose this condition as a constraint on their models.[7] Note that, given that questions define an agent's conceptual framework, distinctions that go beyond her issue are not conceptualizable for her. As a consequence, the knowledge that an (individual) agent can acquire is limited by her issue.

Observation 1. Agent a can acquire the knowledge that p if and only if p is based on an answer to her issue.

In other words, agents only acquire knowledge that is relevant to their questions. This is the dynamic formulation of condition 1.

4.3 Further support for $P1$

If $P1$ did not describe a knowledge property, this would imply that knowledge need not be based on an answer (for the agent), and thus that an agent can know propositions that are more refined than her representation of the world (and therefore presuppose distinctions that she does not herself make). Moreover, it would imply that knowledge need not be the result of successful, goal-directed inquiry. Such a concept of knowledge is far removed from the traditional concept of knowledge from philosophy, if not outright inconsistent.

To illustrate, suppose that Jane knows that petrels are tubenosed. Let p denote "petrels are tubenosed". Furthermore, suppose that p is not entailed by an answer to her issue. This means that she does not distinguish p from a particular set of other possibilities e.g. because she cannot discriminate a tubenose from other bird traits. This scenario can be expressed in logic by the following formula: $K_j p \wedge \neg Q_j p$. Given that Jane knows that petrels are tubenosed, it follows that she has the required type of justified true belief. Jane knows that petrels are tubenosed if she believes that they are tubenosed for *the right reasons*. As the Gettier-examples and its responses exemplify, knowledge is supposed to be the result of successful, goal-directed activity (that is, inquiry aimed at truth), as opposed to lucky coincidence or guesswork. Yet, given that p is not solely based on answers to Jane's questions, how did she acquire her knowledge that p? Intuitively, Jane cannot have acquired her knowledge that p if she lacks the conceptual resources to represent p. If Jane cannot distinguish a tubenose from other bird traits, then it seems that she does not know *what* a tubenose is and that she should not be able to know that petrels are tubenosed. Similarly, if she does not distinghuish petrels from other seabirds, then she does not know that petrels are tubenosed.[8]

In this example, it is unclear how Jane could have acquired her knowledge if not in response to some question. The only alternative is that she must have come to

[7] They introduce a so-called "resolution modality", $R_a P$, that is an intersection modality for the knowledge relations ($\rightarrow_a$) and issue relations ($\approx_a$). Given $P1$, however, the issue relation is required to be included in the knowledge relation, and so the following holds: $R_a P = Q_a P$.

[8] As a second example, consider the following situation. Chris is standing in front of a pack of snow. Suppose that he knows that the snowpack is firn, and suppose that he cannot distinguish firn from other types of snow (e.g. powder snow and perennial snow). If Chris cannot distinguish firn from other types of snow, then it seems that he does not know what firn is and that he therefore should not be able to know that the snowpack is firn.

believe p in response to no question. Yet, as argued by Olsson and Westlund (2006), only beliefs that are held in response to an agent's questions can be considered rational (*cf.* appendix A). The assumption that knowledge need not be based on an answer is inconsistent with a common intuition about the concept of knowledge, namely, that knowledge can only be acquired rationally. Moreover, agents' knowledge is based on their representation of the world, which is in terms of properties (or distinctions): states that are indistinguishable to an agent satisfy the same properties. This means that (from her subjective perspective) these states correspond to the *same world*. So for Jane to have acquired knowledge that p (in the above example), she must have eliminated possibilities that were never in her "model". This is another counter-intuitive consequence. Given the counter-intuitive consequences of the denial of $P1$, it seems warranted to claim that $P1$ indeed describes a property of knowledge.

The above argument is further reinforced when expressed in more formal terms. As noted in section 4.2, $P1$ corresponds to the semantic condition that $\approx_a \subseteq \rightarrow_a$. This condition requires that $s \approx_a s'$ implies $s \rightarrow_a s'$. That is, if an agent cannot conceptually distinguish s from s', then she also cannot distinguish s from s' on the basis of her knowledge. This condition is also justifiable by counterposition: Suppose that $s \not\rightarrow_a s'$. This means that at (actual) world s, the agent knows that s' is *not* the actual world. If so, then intuitively there should exist some property by which she distinguishes s' from the actual world. As such, s should differ from s' with respect to at least some of the answers to a's questions (hence $s \not\approx_a s'$).

4.4 Additional terminology

Some further terminology can be introduced that makes it easier to refer to the propositions that are knowable (or learnable) to an agent given her issue. The propositions that an agent can come to know given her issue are called *agent-relevant propositions*. The underlying assumption is that a proposition is relevant to an agent, if it is a learnable answer to one of her questions. The set of agent-relevant propositions of an agent is called her *question-based potential knowledge* or *potential knowledge* for short. As such, $\approx_a$ can be viewed as the potential knowledge relation.

Definition 8 (Agent-relevant proposition)**.** Proposition P is an *agent-relevant proposition* for agent a if and only if for all states s and s' the following holds: if P is true at s and a cannot distinguish s from s', then P is true at s'.

Definition 9 (Agent-relevant truth)**.** Proposition P is an *agent-relevant truth* for a at state s if and only if P is an agent-relevant proposition and P is true at s ($s \in P$).

Definition 10. Q_aP is true at state s if and only if P is entailed by some agent-relevant truth at s.

Following these definitions, p is an agent-relevant proposition if there exists a proposition P' such that P is of the form Q_aP'. Note that agent-relevant truths are the local versions of agent-relevant propositions.

Chapter 5

Testimonial knowledge and epistemic groups

In the previous chapter, I argued that knowledge is limited to answers to agents' questions. An agent cannot come to know p if she cannot conceptually distinguish states in which p is true from states in which p is false. This means that, irrespective of the knowledge source, the knowledge that an agent can acquire is limited to the (true) answers to her current questions.

In this chapter, I propose two conditions to be imposed on agents' issue-relations. These conditions are needed, I argue, for agents to coherently represent the knowledge of others. Based on these conditions, I present a notion of an *epistemic group* and then introduce *epistemic groups models* to represent such groups. As the proposed conditions on epistemic group models reflect the assumption that testimony is a legitimate source of knowledge, I will begin this chapter with a brief introduction to testimonial knowledge.

5.1 Testimonial knowledge

It has long been recognized that much of our individual knowledge is dependent upon the knowledge of others, most notably because testimony provides agents with social evidence for their beliefs. In recent years, philosophical discussions of testimony have focused primarily on the nature and evidential status of testimony (see e.g. Coady, 1992; Hardwig, 1991; Fricker, 2006; Goldman, 2001; Lackey and Sosa, 2006; Lackey, 2008). Testimony is taken to be a justifiable source of knowledge in addition to e.g. perception, introspection and reasoning (Goldman, 2001; Goldman and Whitecombe, 2011). It is typically assumed, however, that knowledge is only *transmitted* through testimony (as opposed to produced), and thus that an agent can acquire knowledge that p based on the testimony of another agent only if the latter knows that p (Lackey, 2008). In this context, the central problem is why testimony should be considered a justified source of knowledge. To address this conundrum, several philosophers have attempted to specify the conditions under which testimonial knowledge can be considered reliable.[1] The 'conditions' referred to here primarily concern the reliability of the testifier – more precisely, her status as a reliable source of knowledge.

Needless to say, irrespective of the reliability of the testifier, testimony can only be a source of knowledge if agents can coherently represent the knowledge of others. Stronger still, an agent can only acquire knowledge that p based on the testimony of another agent if she believes that the testifier knows that p. Yet, if as was claimed in the previous chapter, agents represent the world in terms of their own questions, then to what extent can they represent the knowledge and questions of others? The knowledge

[1] Testimony is typically taken to be a source of knowledge, evidence or information.

and questions of others, after all, may go beyond an agent's issue. So what higher-order knowledge can agents achieve? The brief answer is that agents represent this information in terms of their own questions. It follows from $P1$ that testimonial knowledge is possible for agents only if the knowledge of others is relevant to them (i.e., potential knowledge). In order to make use of the testimony of others as a source of knowledge, their issues must therefore also include something like "what do others know about my questions?".

5.2 Conditions needed to represent knowledge of others

Recall that in order for an agent to be able to represent (any) knowledge, that knowledge must at least be relevant to her in the sense that she can distinguish it as a possible answer to one of her questions. Further, she must recognize the source of that knowledge as legitimate. In order to be able to recognize knowledge sources as such, their status (as knowledge sources) must be relevant to her. Agents are therefore only able to coherently represent the knowledge of others, if the fact that they (the others) possess this knowledge, or could possess it, is relevant to them. The knowledge that others do possess offers a more direct route to knowledge to the agent, as they can share it with her. The knowledge that others could possess (which is limited to their potential knowledge) indicates their range as knowledge sources and also indicates to what extent the agents' conceptual frameworks coincide.

The above requirements can be expressed in terms of the following two conditions:

Condition 2. If, based on their current knowledge, others can provide answers to an agent's questions, then this fact is itself relevant to her.

Condition 3. If, based on their potential knowledge, others can provide answers to an agent's questions, then this fact is itself relevant to her.

In terms of logic, condition 2 is expressed by the following conditional:

$$(P2)\ K_b Q_a P \rightarrow Q_a K_b P.$$

In words, if b knows something that is relevant to a, then it is relevant to a that b knows this. So both $K_b P$ and P must be potential knowledge for her. This condition can be seen as a social rationality condition on epistemic agency, for it states that the knowledge of others is relevant to epistemic agents–at least, in as far as it answers their questions. Only given this condition can they acquire testimonial knowledge.[2] This corresponds to the following semantic condition:

$$\approx_a \rightarrow_b \subseteq \rightarrow_b \approx_a .$$

The composite relation $\rightarrow_b\approx_a$ encodes the knowledge of b about (the answers to) a's question(s): all b's knowledge that a can know solely based on true answers to her issue

[2]Similar considerations apply to other sources of knowledge, such as perception and memory.

(which answers b can provide). This knowledge is given by the worlds t that are $\approx_a$-indistinguishable from any worlds that are considered by b to be epistemically possible alternatives of the actual world s.[3]

Condition 3 extends condition 2 to include b's potential knowledge. In terms of logic, this condition is expressed as follows:

$$(P3)\ Q_bQ_aP \to Q_aQ_bP.$$

$P3$ states that if b can know (given her issue) anything that is relevant to a, then this fact (that b's potential knowledge includes potential knowledge of a) is itself relevant to a. This condition is based on the assumption that epistemic agents represent others as future knowledge sources. Epistemic agents are interested in the issues (and thus potential knowledge) of others, for these issues allow them (the others) to become knowledge sources for the agents. $P3$ corresponds to the following semantic condition:

$$\approx_a\approx_b \subseteq \approx_b\approx_a .$$

The composite relation $\approx_b\approx_a$ encodes the potential knowledge of b about (the answers to) a's question(s): all b's potential knowledge that a can know solely based on true answers to her issue (which answers b could provide had she answered her own questions).

The assumption that underlies these two conditions is that testimony is a legitimate source of knowledge. In as far as an agent acquires testimonial knowledge (or can acquire such knowledge), both of these conditions must be met. The ability to acquire such knowledge is an epistemically-desirable property, because it allows an agent to make use of (and built upon) the knowledge of others.

5.3 Epistemic groups and epistemic group models

Thus far, I have only considered how agents represent the knowledge of other individuals. Yet, as was discussed in chapter three, groups can also be seen to possess knowledge that may go beyond the sum of the knowledge possessed by their individual members. Following the argument of the previous section, agents are only able to coherently represent the knowledge of a group *as a whole*, if the fact that the group possesses this knowledge is relevant to them. It was further mentioned that epistemic groups are groups of individuals tied together by some pertinent epistemic property. In characterizing epistemic groups, I take agents' mutual interest in each others' knowledge and questions (i.e., conceptual distinctions) to be the epistemic property that ties them together.

[3] The need for this condition can also be explained via its counterpositive. Assume that it does *not* hold that $s \to_b \approx_a s'$. This means that b *knows* that the answer to a's questions at s' is *not* the correct answer (at the actual world s); but then this fact (that b knows this) should itself be relevant for a. It is an epistemic fact (about b's knowledge) that settles some of a's own questions, and hence it should be a part of a's issue. Any possible world $t \to_b s'$ is one in which b *wouldn't know* the above relevant fact a, so any such world should be conceptually distinguishable for a from the actual world, because it differs from s with respect to a fact that should be relevant to a: $\forall t(t \to_b s' \Rightarrow s \not\approx_a t)$. So $s \approx_a \to_b s'$ should *not* obtain.

Agents within the same epistemic group view each other, as well as their group as a whole, as knowledge sources. Conversely, groups that consist of agents that do not view each other as knowledge sources are not subjects of group knowledge. Here, the idea is that at the minimum, the agents of an epistemic group should view each other as knowledge sources–otherwise it is not clear why it is valuable to them to be part of an epistemic group (from an epistemic perspective). Being part of an epistemic group must allow the individual members to achieve something (e.g. additional knowledge) that they could not otherwise achieve. For example, suppose that Jane and Chris are both interested in Twin Peaks. Chris knows several things about Twin Peaks that are of interest to Jane, as he has read a biography of the director. Jane, however, does not trust Chris' testimony on the topic. She is therefore not interested in his knowledge about it. Intuitively, the "group" that consists of Jane and Chris should not be referred to as a group that can possess knowledge. It is unclear how their "group knowledge" is valuable to Jane and Chris, that is, it is unclear what it allows them to achieve individually and as a group.

Furthermore, agents that do not represent the larger group *as a whole* as a knowledge source do not take full advantage of the epistemic benefit of being part of a larger group (that can be associated with more knowledge). In its current formulation, condition $P2$ does not require this, as it only ties an agent's issue to the knowledge of one other agent (at a time)–and similarly for $P3$.[4] As said in chapter 3, the knowledge of a group can go beyond the sum of the individual members' knowledge, and this "additional" knowledge may be relevant to an agent. In fact, it is this "additional" knowledge that makes group knowledge a concept that cannot be expressed simply in terms of individual knowledge alone. In order for group knowledge (of some sort) to be useful to individual agents,[5] they must represent their group *as a whole* as an epistemic subject. This requirement can be expressed by the following condition:

Condition 4. If, based on their distributed knowledge, groups can provide answers to an agent's questions, then this fact is itself relevant to her.

In terms of logic, this condition is expressed by the following conditional (where $a \in G$):

$$(P2')\ Dk_G Q_a P \rightarrow Q_a Dk_G P$$

In words, if the group has distributed knowledge that P is relevant to a, then this fact is itself relevant to a. This corresponds to the following semantic condition:

$$\approx_a \rightarrow_G \subseteq \rightarrow_G \approx_a .$$

[4]Formally, the 'problem' is that $(\bigcap_{b\in G} \rightarrow_b); \approx_a \neq \bigcap_{b\in G}(\rightarrow_b; \approx_a)$, where $\rightarrow_G$ is obtained from the intersection of the group members' knowledge relations: $\rightarrow_G := \bigcap_{a\in G} \rightarrow_a$. The first (composite) relation is obtained from $\rightarrow_G$, which is then restricted to a's issue, whereas the latter is obtained in a step by step manner from all $\rightarrow_b$ (for all $b \in G$, including a), which are restricted separately to a's issue. The former relation may thus contain more potential knowledge than the latter.

[5]That is, in addition to individual knowledge.

The composite relation $\to_G \approx_a$ encodes the group's distributed knowledge about a's question(s).[6] As before, condition $P2'$ can be seen as a social rationality condition on epistemic agency. It states that epistemic agents are interested in the combined knowledge of themselves with others.

Given conditions $(P2')$ and $(P3)$, I propose the following definition for the notion of an *epistemic group*:

Definition 11 (Epistemic group)**.** A group G of agents A is called an *epistemic group* whenever each group member's issue-relation meets conditions $P2'$ and $P3$.

Only epistemic groups are capable of some type of group knowledge. The epistemic property that ties the group members together is mutual interest in each others' (potential) knowledge. As such, the group members self-identify as group members on the basis of epistemic properties that (conceptually) precede any group knowledge, namely, $P2'$ and $P3$.

To model epistemic groups and their knowledge, I propose new models, called *epistemic group models*, that are subject to conditions $P1$, $P2'$ and $P3$. An epistemic group model is defined as follows:

Definition 12 (Epistemic group model)**.** Given a set A of agents and a set Φ of atomic sentences, an *epistemic group model* over (A, Φ) is a tuple $\mathbf{S} = (S, \to_{a(a\in A)}, \approx_{a(a\in A)}, \|\bullet\|)$, such that $(S, \to_{a(a\in A)}, \|\bullet\|)$ is an epistemic model over (A, Φ) and $\approx_a$ is a map associating each agent $a \in A$ with a reflexive, transitive relation $\approx_a \subseteq S \times S$, satisfying the following three conditions[7]: $(*)\ \approx_a\ \subseteq\ \to_a$; $(**)\ \approx_a \to_G\ \subseteq\ \to_G \approx_a$; $(***)\ \approx_a \approx_b\ \subseteq\ \approx_b \approx_a$.

Recall that '$\to_a$' is assumed to be reflexive and transitive (*cf.* def. 2). These models are just a special class of epistemic-issue models, as first introduced by van Benthem and Minică (2009). The differences with their models are that (1) the current models are required to satisfy three additional conditions that are meant to capture the specific interpretation of the issue-relation proposed in this thesis, and (2) neither the knowledge-relation nor the issue-relation is assumed to be an equivalence relation.[8] Finally, the three conditions $(*)$, $(**)$ and $(***)$ correspond to the following universally true conditionals, *viz.* to $P1$, $P2'$ and $P3$, respectively:

$$(*) :\ K_aP \Rightarrow Q_aP\ (P1)$$
$$(**) :\ Dk_GQ_aP \Rightarrow Q_aDk_GP\ (P2')$$
$$(***) :\ Q_bQ_aP \Rightarrow Q_aQ_bP\ (P3)$$

In the remainder of this thesis, I shall use only epistemic group models. Any set of agents $G \subseteq A$ in such models is an epistemic group, including $\{a\}$. Further, I shall also

[6] Note that $\to_G ; \approx_a = (\bigcap_{b\in G} \to_b); \approx_a$.

[7] These conditions can be stated more explicitly by saying that, for all worlds $s, t, w \in S$ we have: $(*)$ if $s \approx_a w$ then $s \to_a w$; $(**)$ if $s \approx_a t \to_b w$ then there exists some $t' \in S$ such that $s \to_b t' \approx_a w$; $(***)$ if $s \approx_a t \approx_b w$ then there exists some $t' \in S$ such that $s \approx_b t' \approx_a w$.

[8] Knowledge is thus not assumed to meet the S5 conditions, but only S4.

use the names $P1$, $P2'$ and $P3$ to refer to the semantic conditions (instead of $(*)$, $(**)$ and $(***)$).

For some purposes, condition $P2$ is enough. To model groups of agents that satisfy conditions $P1$, $P2$ and $P3$, the notion of *testimonial model* can be defined. The agents in such models are assumed to represent the knowledge of their fellow (individual) group members, but they are not assumed to represent the knowledge of their group as a whole. This is where the name "testimonial model" comes from.

Definition 13 (Testimonial model)**.** Given a set A of agents and a set Φ of atomic sentences, a *testimonial model* over (A, Φ) is a tuple $\mathbf{S} = (S, \to_{a(a\in A)}, \approx_{a(a\in A)}, \| \bullet \|)$, such that $(S, \to_{a(a\in A)}, \| \bullet \|)$ is an epistemic model over (A, Φ) and $\approx_a$ is a map associating each agent $a \in A$ with a reflexive, transitive relation $\approx_a \subseteq S \times S$, satisfying the following three conditions: $(*)$ $\approx_a \subseteq \to_a$; $(**)$ $\approx_a \to_b \subseteq \to_b \approx_a$; $(***)$ $\approx_a \approx_b \subseteq \approx_b \approx_a$.

Chapter 6

Collective knowledge

This chapter brings together the ideas expressed in the foregoing chapters. I propose (formal) definitions for group epistemic notions that are interpreted on epistemic group models. In particular, I introduce a notion that expresses the potential knowledge of an individual within a group (*potential individual knowledge*) and extend this notion to that of *collective knowledge*. I further introduce the notion of a *common issue* and then characterize *collective knowledge* as distributed knowledge about a *common issue*. The formal definition of *collective knowledge* encodes how it depends on the knowledge of the group members, as well as how the group members are tied together through shared epistemic properties. As such, this notion is well-aligned with philosophical concepts of group knowledge. *Collective knowledge* can also be viewed as potential common knowledge, in parallel to *potential individual knowledge*. As further support for the characterizations of *potential individual knowledge* and *collective knowledge* as types of potential knowledge, I briefly consider knowledge dynamics. Finally, I look at the philosophical implications of *collective knowledge*. In this chapter, individual knowledge, common knowledge and distributed knowledge are defined as usual (*cf.* chapter 2.2).

6.1 Potential individual knowledge within a group

The epistemic benefit of being part of a larger group is that it gives agents access to more knowledge. The potential knowledge of an individual within a group, here referred to as *potential individual knowledge*, is the knowledge that an agent can in principle acquire from the combined testimony of her group members. *Potential individual knowledge* can be expressed by the following informal definition:

> An agent a, belonging to a group G, potentially knows that P within G if and only if:
>
> (1) P is entailed by some proposition P' that is (question-based) potential knowledge for her[1];
>
> (2) P' is distributed knowledge in G.

In other words, an agent's potential knowledge within a group consists of the set of propositions that is entailed by the distributed knowledge of the group, restricted to her

[1] My use of the term "potential" may be somewhat confusing here. Note that an agent's *potential individual knowledge within a group* is included in her *question-based potential knowledge*. The latter notion captures the maximum knowledge that an agent can acquire given her questions and the answers that are learnable for her, whereas the former notion captures the knowledge that an agent can acquire through the testimony of her group.

(question-based) potential knowledge.[2] According to this characterization, agents come to know only propositions that are relevant to them. As such, the value to an agent of being part of an epistemic group is determined not only by the extent of her group members' knowledge, but also by the extent to which their knowledge coincides with her potential knowledge. The distributed knowledge of a group is of little use to its members if it is not relevant to them.

Formally, *potential individual knowledge* within a group $G \subseteq A$ is denoted by K_a^G, which is interpreted as the Kripke modality for the composite relation $\rightarrow_G \approx_a$. As stated in chapter 5, '$\rightarrow_G \approx_a$' denotes the distributed-knowledge-about-a's-question(s)-relation.

Definition 14 (Potential individual knowledge within a group)**.**

$$K_a^G P := [\rightarrow_G \approx_a] P \quad \text{(assuming that } a \in G\text{)}.$$

Agent a potentially knows proposition P within group G if and only if P is true at all states t' that she cannot conceptually distinguish from the states s' that are epistemically possible based on the group's distributed knowledge at s. Thus, a has potential knowledge that P (at s), if P is true at all states t' such that $s \rightarrow_G s' \approx_a t'$. This means that there is distributed knowledge in G that P is an agent-relevant truth:

Proposition 1. For any fixed $G = \{a_1, a_2, \ ... \ a_n\}$ the following holds*:*

$$K_a^G P \Leftrightarrow Dk_G Q_a P.$$

Proof. Suppose $w \models_{\mathbf{S}} K_a^G P$. This means that $\forall w', w'' \in \mathbf{S} : (w \rightarrow_G w' \approx_a w'' \Rightarrow w'' \models_{\mathbf{S}} P) \Leftrightarrow \forall w'(w \rightarrow_G w' \Rightarrow \forall w''(w' \approx_a w'' \Rightarrow w'' \models_{\mathbf{S}} P)) \Leftrightarrow \forall w'(w \rightarrow_G w' \Rightarrow w' \models_{\mathbf{S}} Q_a P)$, and hence $w \models_{\mathbf{S}} Dk_G Q_a P$. □

Further, within groups, K_a^G commutes:

Proposition 2. For any $a, b \in G$ the following holds*:*

$$K_a^G K_b^G P \Leftrightarrow K_b^G K_a^G P$$

Proof. ($\Rightarrow$) Suppose $K_a^G K_b^G P$. From $K_a^G P \Leftrightarrow Dk_G Q_a P$, it follows that $K_a^G K_b^G P \Leftrightarrow Dk_G Q_a Dk_G Q_b P$. By veracity of Dk, it follows that $Dk_G Q_a Q_b P$. Further, by condition $P3$, it then follows that $Dk_G Q_b Q_a P$. Since Dk is positively introspective, it follows that $Dk_G Dk_G Q_b Q_a P$, which implies, by condition $P2'$, that $Dk_G Q_b Dk_G Q_a P$, where $Dk_G Q_b Dk_G Q_a P \Leftrightarrow K_b^G K_a^G P$.

($\Leftarrow$) Because of the symmetry of $(K_a^G K_b^G P \Leftrightarrow K_b^G K_a^G P)$, the proof for this direction will match the above. □

Note that this equivalence only holds within *the same group.* K_a^G commutes with respect to the same knowledge source, and thus not with respect to different groups. While K_a is indeed a special case of K_a^G, it presupposes that $G = \{a\}$: that is, $K_a = K_a^{\{a\}}$. Conceptually, this equivalence holds because the group is viewed by its members as a knowledge source in its own right to which they all members have equal access, and because the group members are mutually interested in each other's potential knowledge.

[2]Note that this also includes her 'normal' individual knowledge.

6.2 Collective knowledge

Potential individual knowledge is the knowledge that an agent could know based on the testimony of her group. As such, it is still a type of individual knowledge. In chapter 4, I claimed that $P1$ is a property of knowledge, whether held by individuals or groups. This makes it natural to view group knowledge as the distributed knowledge about a *common issue*. Here, as well as in the remainder of this thesis, I abuse my own terminology and use the term "common issue" to refer to (question-based) potential common knowledge, which is the knowledge that could be common knowledge based on the answers to the group members' questions. In order to have *collective knowledge*, an epistemic group must therefore have a *common issue*.

I propose the following informal definition of *collective knowledge*:

> Group G has *collective knowledge* that P if and only if the following holds:
>
> (1) P is entailed by some proposition P' that is (question-based) potential common knowledge in G, and;
>
> (2) P' is distributed knowledge in G.

Thus, if a group collectively knows that P, then it is potential common knowledge within the group that P, which means that P can be known solely based on true answers to the group members' common questions.

Formally, *collective knowledge* of group $G \subseteq A$ is denoted by K_G, which gets the following definition:

Definition 15 (Collective knowledge)**.**

$$K_G P \ := \ [(\bigcup_{a \in G} \rightarrow_G \approx_a)^*]P$$

Or stated otherwise, $K_G P \ := \ \bigwedge_{a \in G} K_a^G P \wedge \bigwedge_{a,b \in G} K_a^G K_b^G P \wedge \ldots$

Recall that R^* is the reflexive-transitive closure of relation R. As such, K_G can be viewed as the "common" version of *potential individual knowledge*, as it mirrors the definition of common knowledge, only replacing the instances of K_c with K_c^G (for all $c \in G$).

Thus, K_G encodes the knowledge that could become common knowledge through the groups' testimony. It entails that group knowledge is limited to propositions that are relevant to all group members.[3] This assumption is supported, for example, by Rolin's characterization of an epistemic group as a group of agents that are jointly committed

[3]This appears to be a reasonable assumption. To see this, suppose that Chris is an ornithologist and an expert on the anatomy of seabirds. He knows that puffins are not tubenosed. Further, suppose that Jane cannot distinguish different sea birds from each other–they are all gulls to her. Indeed, she has no interest in the topic. Intuitively, the fact that puffins are not tubenosed should not be considered group knowledge of the "group" that consists of Jane and Chris, for it is unclear what the value of this knowledge is to Jane (since it is not potential knowledge for her), and what it allows their group to achieve.

to the propositions (i.e., 'default entitlements') on which their group knowledge is based (Rolin, 2008). Corlett appears to make a similar assumption when he argues that epistemic groups consist of members who identify relationally with each other based on the fact that they have shared epistemic motives and decision-making capacities (Corlett, 2007).

Formally, the notion of *common issue* is denoted by Q_G, and interpreted as follows:

Definition 16 (Common issue)**.** For any group $G \subseteq A$*:*

$$Q_G P \;:=\; [(\bigcup_{a \in G} \approx_a)^*]P$$

Or stated otherwise, $Q_G P \;:=\; P \wedge \bigwedge_{a \in G} Q_a P \wedge \bigwedge_{a,b \in G} Q_a Q_b P \wedge \dots$

Within groups, the conjunctions in the characterization of both $Q_G P$ and $K_G P$ can be reduced to a single string with no repititions:

Proposition 3. For any fixed $G = \{a_1, a_2, \ldots a_n\}$ the following holds*:*

$$Q_G P \Leftrightarrow Q_{a_1} Q_{a_2} \ldots Q_{a_n} P$$

Proof. ($\Rightarrow$) Suppose $Q_G P$. By definition $Q_G P = (\bigwedge_{a_1 \in G} Q_{a_1} P \wedge \bigwedge_{a_1, a_2 \in G} Q_{a_1} Q_{a_2} P \wedge \ldots)$. This infinite conjunction of strings implies any one of its strings. The string $(Q_{a_1} Q_{a_2} \ldots Q_{a_n} P)$ is one of the conjunctions in the characterization of $Q_G P$.

($\Leftarrow$) Suppose $(Q_{a_1} Q_{a_2} \ldots Q_{a_n} P)$. It is enough to show that if $(b_1, \ldots b_k)$ is a string of any length (possibly with repetitions) in G, then $(Q_{a_1} Q_{a_2} \ldots Q_{a_n} P)$ implies $(Q_{b_1} \ldots Q_{b_k} P)$. Since $b_1, \ldots b_k \in G$, they come from $G = \{a_1, a_2, \ldots, a_n\}$. By positive introspection of Q_a, Q_a's can be repeated. Thus $(Q_{a_1} Q_{a_2} \ldots Q_{a_n} P)$ implies $(Q_{a_1} Q_{a_1} Q_{a_2} \ldots Q_{a_n} Q_{a_1} Q_{a_2} \ldots P)$ for however many repetitions. This can be done any number of times in order to match the Q_a's in $(Q_{b_1} \ldots Q_{b_k} P)$ (though they might be in a different order). Similarly, by veracity of Q_a, any number of Q_a's can be deleted from the string in order to match the number in $(b_1, \ldots b_k)$. By condition $P3$, Q_a permutes, and so in some finite number of steps, $(Q_{a_1} Q_{a_1} Q_{a_2} \ldots Q_{a_n} Q_{a_1} Q_{a_2} \ldots P)$ can be reordered so as to obtain $(Q_{b_1} \ldots Q_{b_k} P)$. Hence $(Q_{a_1} Q_{a_2} \ldots Q_{a_n} P)$ implies $(Q_{b_1} \ldots Q_{b_k} P)$. □

Although the formal definition of a *common issue* can be simplified for epistemic groups, the intended meaning does not change. Thus, $Q_{a_1} Q_{a_2} \ldots Q_{a_n} P$ should not be read as only stating that P is a mutual issue (i.e., that $Q_{a_1} P$ and $Q_{a_2} P$ and $Q_{a_3} P$ etc). Indeed, proposition P is intended common knowledge in G, and so P is not just potential knowledge for all agents.

Similar considerations apply for the characterization of *collective knowledge*:

Proposition 4. For any fixed $G = \{a_1, a_2 \ ... \ a_n\}$ the following holds*:*

$$K_G P \Leftrightarrow K^G_{a_1} K^G_{a_2} \dots K^G_{a_n} P$$

Proof. ($\Rightarrow$) Suppose $K_G P$. By definition $K_G = (\bigwedge_{a \in G} K^G_a P \wedge \bigwedge_{a,b \in G} K^G_a K^G_b P \wedge \dots)$. This infinite conjunction of strings implies any one of them. The string $(K^G_{a_1} K^G_{a_2} \dots K^G_{a_n} P)$ is one of the conjunctions in the characterization of $K_G P$.

($\Leftarrow$) Suppose $(K^G_{a_1} K^G_{a_2} \dots K^G_{a_n} P)$. It is enough to show that if $(b_1, \dots b_k)$ is a string of any length (possibly with repititions) in G, then $(K^G_{a_1} K^G_{a_2}, \dots K^G_{a_n} P)$ implies $(K^G_{b_1} K^G_{b_2}, \dots K^G_{b_k} P)$. The proof mirrors the proof of proposition 3, because, like Q_a, K^G_a commutes (K^G_a commutes only *within groups.* This means that $K^G_a K^G_b = K^G_b K^G_a$ *only if* $a, b \in G$.) and is positively introspective. □

Collective knowledge is distributed knowledge about a *common issue*. Moreover, *collective knowledge* behaves like individual knowledge: the properties of the (individual) knowledge relation are preserved. As such, it presupposes that epistemic groups agents have the same properties as individual agents (from the perspective of logic). Propositions $5-8$ make these claims precise:

Proposition 5. For any fixed $G = \{a_1, a_2, \ ..., \ a_n\}$ the following holds*:*

$$K_G P \Leftrightarrow Dk_G Q_G P$$

Proof. ($\Rightarrow$) Suppose $K_G P$. By definition $K_G P = (\bigwedge_{a \in G} K^G_a P \wedge \bigwedge_{a,b \in G} K^G_a K^G_b P \wedge \dots)$. This is an infinite conjunction of all possible strings. By unfolding the K^G_a's, each of the terms in these strings can be expressed in the following form: $(Dk_G Q_a Dk_G Q_a, \dots P)$. By factivity of Dk_G and Q_a, it follows that $(Q_{a_1} \dots . Q_{a_n} P)$. Then by normality of Dk_G, it follows that $(Dk_G Q_{a_1} \dots Q_{a_n} P)$. This is the same as writing $Dk_G Q_G$.

($\Leftarrow$) Suppose $Dk_G Q_G P$. It immediately follows that $(Dk_G Q_{a_1} Q_{a_2} \dots Q_{a_n} P)$ (*cf.* prop. 3). By positive introspection of Dk_G, any number of Dk_G's can be added in front. Thus, $(Dk_G Q_{a_1} Q_{a_2} \dots Q_{a_n} P)$ implies $(Dk_{G_1} \dots Dk_{G_n} Q_{a_1} Q_{a_2} \dots Q_{a_n} P)$. Next, condition $P2'$ allows Dk_G's to pass one Q_a at a time. Hence, in some finite number of steps, $(Dk_{G_1} \dots Dk_{G_n} Q_{a_1} Q_{a_2} \dots Q_{a_n} P)$ can be reordered to the following: $(Dk_G Q_{a_1} DK_G Q_{a_2} \dots Dk_G Q_{a_n} P)$.
This shows that $(Dk_G Q_{a_1} DK_G Q_{a_2} \dots Dk_G Q_{a_n} P) = (Dk_G Q_{a_1} Q_{a_2} \dots Q_{a_n} P)$.
Finally, $(Dk_G Q_{a_1} DK_G Q_{a_2} \dots Dk_G Q_{a_n} P) = (K^G_{a_1} K^G_{a_2} \dots K^G_{a_n} P) = K_G P$. □

Proposition 6. For any fixed $G = \{a_1, a_2, \ ... \ a_n\}$ the following holds*:*

$$K_G P \Rightarrow Q_G P$$

In words, *collective knowledge* implies a *common issue*.

Proof. Suppose that $K_G P$. By proposition 5, this is equivalent to $Dk_G Q_G P$. By veracity of Dk_G it follows that $Q_G P$. □

Proposition 7. $K_G P \Rightarrow Dk_G P$.

In words, *collective knowledge* implies distributed knowledge.

Proof. Suppose $K_G P$. By definition $K_G P \Leftrightarrow (K^G_{a_1} K^G_{a_2}, ... K^G_{a_n} P)$ $\Leftrightarrow (Dk_G Q_{a_1} DK_G Q_{a_2} ... Dk_G Q_{a_n} P)$. By condition $P2'$ it follows that $(Q_{a_1} Dk_G Q_{a_2} Dk_G ... Q_{a_n} Dk_G P)$. Since Q_a and Dk_G are both factive, it follows that $Dk_G P$. □

Proposition 8. $K^{\{a\}}_a P \Leftrightarrow Dk_{\{a\}} P \Leftrightarrow K_a P \Leftrightarrow K_{\{a\}} P$

Proof. To show: $\rightarrow_a ; \approx_a = \rightarrow_a = \rightarrow_{\{a\}} = (\rightarrow_a ; \approx_a)^*$.

Clearly $\rightarrow_a = \rightarrow_{\{a\}}$. For $\rightarrow_a ; \approx_a = \rightarrow_a$, both inclusions have to be shown: ($\Rightarrow$) Suppose $\rightarrow_a ; \approx_a$. By $P1$ (ie, $\approx_a \subseteq \rightarrow_a$) and monotonicity of relational composition, it follows that $\rightarrow_a ; \approx_a \subseteq \rightarrow_a \rightarrow_a$, and by transitivity of $\rightarrow_a$, this means that $\rightarrow_a ; \approx_a \subseteq \rightarrow_a \rightarrow_a = \rightarrow_a$. ($\Leftarrow$). Suppose $\rightarrow_a$. By veracity of $\approx_a$ (i.e., $id \subseteq \approx_a$) and monotonicity of relational composition, it follows that $id; \rightarrow_a \subseteq \rightarrow_a ; \approx_a$ and hence that $\rightarrow_a \subseteq \rightarrow_a ; \approx_a$. By transitivity of $\rightarrow_a$, it follows that $(\rightarrow_a ; \approx_a)^* = (\rightarrow_a)^* = \rightarrow_a$.

□

6.3 Further support: knowledge dynamics

Collective knowledge and *potential individual knowledge* are types of implicit knowledge. They are meant to capture the knowledge that (group) agents would know, were the members of the group to share all their knowledge with them. In this section, I consider the action of truthful public announcement through which agents share their knowledge with their group members in an epistemic model. Learning from truthful public announcement is certainly an idealized form of learning, though the possibility of testimonial knowledge also presupposes it. By showing that *collective knowledge* can indeed become individual knowledge–even common knowledge–through public announcement, and similarly for *potential individual knowledge*, considering such actions can provide further justification for their characterizations.

Condition $P1$ is a static condition: it states that if agent a knows that p, then it follows that p is entailed by answers to her questions. From the perspective of knowledge acquisition, it must also hold that agents only acquire knowledge that is relevant to their issues. This condition can be viewed as the dynamic version of $P1$, as it describes the corresponding property of knowledge acquisition.

The above condition, however, does not specify how to handle the case where an agent is confronted with testimony that is partially relevant to her issue. Suppose that an agent is informed that some proposition P is true, but that P distinguishes between states that are conceptually indistinguishable for her. The above condition can be interpreted as implying either that a will ignore P altogether or that she will come to know only the agent-relevant truths from P (even though P contains information that she cannot conceptualize). I adopt the latter interpretation. This is based on the assumption that epistemic agents are able to extract relevant information. I formalize this as P_a, where $P \subseteq S$ and $a \in A$:

$$P_a := \{s \in S : s' \approx_a s \text{ for some } s' \in P\} = \bigcap\{P' \subseteq S : P' \text{ is agent relevant and } P \subseteq P'\}$$

In words, P_a is the strongest agent-relevant consequence of P.

In Dynamic Epistemic Logic, the action of publicly (and truthfully) announcing P is usually denoted by $!P$ – where $!P$ takes any given model $\mathbf{S}$ to a new model $\mathbf{S^P}$ (Baltag *et al.* 2008). However, in order to ensure that agents only update their model with P_a (rather than P), the standard semantics for this action has to be modified. Moreover, the action $!P$ concerns a single proposition. Yet, both *collective knowledge* and *potential individual knowledge* are characterized as the knowledge that would become common knowledge and individual knowledge (respectively) were all group members to share *all* their knowledge. This latter action is denoted by $!G$. That is, $!G$ denotes the action by which all agents in G publicly announce all their knowledge. At state s, this corresponds to them announcing (together) $P = \{w : s \rightarrow_G w\}$–that is, the conjunction $\bigcap_{b \in G}\{w : s \rightarrow_b w\}$ of the knowledge of each $b \in G$ at s. Thus $!G$ is the public announcement of P. From this announcement, each agent will only learn its strongest agent-relevant consequence: $P_a = \{t : s \rightarrow_G \approx_a t\}$. As such, a's new knowledge $\rightarrow_a^{!G}$ after action $!G$ is given by the combination of her old knowledge with P_a: $s \rightarrow_a^{!G} t$ iff $s \rightarrow_a t$ and $s \rightarrow_G \approx_a t$. The action $!G$ modifies the models as follows:

Definition 17 (Sharing all knowledge in epistemic group models)**.** Given an epistemic group model $\mathbf{S} = (S, \rightarrow, \approx, \| \bullet \|)$ over (A, Φ) and a group $G \subseteq A$, a new epistemic model $\mathbf{S}^{!G} = (S^{!G}, \rightarrow^{!G}, \approx^{!G}, \| \bullet \|^{!G})$, is obtained as follows:

1. the set of possible worlds stays the same: $S^{!G} = S$;
2. $s \rightarrow_a^{!G} t$ iff both $s \rightarrow_a t$ and $s \rightarrow_G \approx_a t$.
3. the issue-relations stay the same: $s \approx_a^{!G} t$ iff $s \approx_a t$.
4. the valuation stays the same: $\|p\|^{!G} = \|p\|$, for $p \in \Phi$.

When $P \subseteq S$ is a set of states, the following two propositions describe the dynamics of K_G and K_a^G:

Proposition 9. $K_G p \Leftrightarrow [!_G] Ck_G p$ (for atomic sentences p).

Proof. To show: $w \models_{\mathbf{S}} K_G p$ if and only if $w \models_{\mathbf{S}} [!_G] Ck_G p$.
By definition $w \models_{\mathbf{S}} K_G p \Leftrightarrow \forall w' \in \mathbf{S} : (w\, (\bigcup_{a \in G} \rightarrow_G ; \approx_a)^* \, w' \Rightarrow w' \models_{\mathbf{S}} p \Leftrightarrow \forall w' \in \mathbf{S}^{!G} : (w\, (\bigcup_{a \in G} \rightarrow_a^{!G})^* \, w' \Rightarrow w' \models_{\mathbf{S}^{!G}} p) \Leftrightarrow w \models_{\mathbf{S}^{!G}} Ck_a p \Leftrightarrow w \models_{\mathbf{S}} [!_G] K_G p$. ☐

Proposition 10. $K_a^G p \Leftrightarrow [!_G] K_a p$ (for atomic sentences p).

Proof. To show: $w \models_{\mathbf{S}} K_a^G p$ if and only if $w \models_{\mathbf{S}} [!_G]K_a p$.
By definition $w \models_{\mathbf{S}} K_a^G p \Leftrightarrow \forall w' \in \mathbf{S} : (w(\to_G; \approx_a)w' \Rightarrow w' \models_{\mathbf{S}} p) \Leftrightarrow \forall w' \in \mathbf{S}^{!_G} : (w \to_a^{!_G} w' \Rightarrow w' \models_{\mathbf{S}^{!_G}} p) \Leftrightarrow w \models_{\mathbf{S}^{!_G}} K_a p \Leftrightarrow w \models_{\mathbf{S}} [!_G]K_a p$. □

It should be noted that the new models $\mathbf{S}^{!G}$ are not necessarily epistemic group models, as condition $P2'$ is not preserved by the transformation. Still, they do satisfy all the other semantic conditions of epistemic group models. In fact, after action $!G$ the new model is a testimonial model, thus satisfying condition $P2$ (*cf.* chapter 5).

6.4 Philosophical implications of collective knowledge

How does the notion of *collective knowledge* relate to other philosophical concepts of group knowledge? *Collective knowledge* suggests a non-summative interpretation, for it allows the group to hold knowledge that is not held by any of its constituents. It inherits this property from the distributed knowledge on which it is based. It should however be recalled that distributed knowledge is generally *not* interpreted non-summatively, but rather as implicit knowledge. *Collective knowledge* could also be interpreted as non-summative group knowledge and/or as implicit knowledge.

Regarding the possible subjects of group knowledge, *collective knowledge* is held only by epistemic groups. Such groups consist of members that are tied together by mutual interest in each others' (potential) knowledge. Intuitively, this means that the agents in an epistemic group view each other and their group as a whole as sources of knowledge. Their common issue can be viewed as a common goal, keeping in mind that it need not be actively pursued by all group members. In as far as the group has a common issue, it can have collective knowledge. Put differently, collective knowledge has the property that it is issue-based in the sense that all collective knowledge is based on answers to a common issue ($K_G P \to Q_G P$). This is well-aligned with Wray's characterization of epistemic groups, though it does not require an agreed upon division of epistemic labor among the group members. In fact, it does not explicitly make any requirements with respect to epistemic labor. It also comes close to Rolin's characterization of epistemic groups as groups of individuals tied together by a joint commitment to certain pieces of (individual) knowledge, which moreover share the epistemic responsibility for each others' individual knowledge. Further, it is also well-aligned with Corlett's view of epistemic groups as consisting of agents that identify with each other based on shared epistemic motives. Common among these views is the assumption that group knowledge is held only by groups that have common interests ($K_G P \to Q_G P$).

It is an explicit part of the definition of *collective knowledge* that it implies *potential individual knowledge*, and therefore is accessible to its members. This is made explicit by the dynamics of K_G and K_a^G. This 'accessibility' assumption appears to be common in both the philosophical and the formal literature on group knowledge. I have not come across any definition of group knowledge that is not based on this assumption.

Chapter 7

The logic of collective knowledge

This chapter introduces an axiomatic system for a logic with *collective knowledge*. This system, called *epistemic group logic* (EGL), is a sound and complete axiomatization with respect to the class of epistemic group models.

7.1 Axiomatization of EGL

Epistemic group logic takes Q_a and Dk_G as basic operators. It is given the following syntax where ($G \subseteq A$):

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \varphi \mid Q_a\varphi \mid Dk_G\varphi$$

Here, $p \in PROP$ and $a \in A$, where $PROP$ and A are sets of propositions and agents, respectively. Existenial modalities are abbreviated as usual, thus $\neg Dk_G\neg\varphi$ is written as $< Dk_G\varphi >$ and $\neg Q_a\neg\varphi$ is written as $< Q_a\varphi >$. I assume as given a finite set A of agents, together with a fixed enumeration of A without repetitions, i.e., a linear order $<$ on A. The other modalities are then expressed via the following abbreviations:

- $K_a^G\varphi := Dk_G Q_a\varphi$
- $K_a\varphi := Dk_{\{a\}}\varphi$
- $Q_G\varphi := Q_a Q_b \ \ldots \ Q_n\varphi$ (where $a < b < \ldots < n$ is the enumeration of the agents in G according to the order $<$ given above.)
- $K_G\varphi := Dk_G Q_G\varphi$

In addition to the standard Boolean operators, the six modalities are given the following intended meaning:

- $Q_a\varphi$ (agent a has (question-based) potential knowledge that φ),
- $Dk_G\varphi$ (it is distributed knowledge in G that φ),
- $K_G\varphi$ (group G collectively knows φ),
- $K_a\varphi$ (agent a knows that φ),
- $K_a^G\varphi$ (agent a potentially knows within G that φ),
- $Q_G\varphi$ (group G has potential common knowledge that φ),

Semantics The semantics is given by an interpretation map that associates with each sentence φ of EGL a proposition $\|\varphi\|_{\mathbf{S}} \subseteq S$ in any given epistemic group model $\mathbf{S}$, as specified in chapter 2. Intuitively, $\|\varphi\|_{\mathbf{S}}$ is the set of all worlds in S satisfying φ. The definition is by induction, in terms of the obvious compositional clauses (using the operators defined above).

Proof system. The axiomatic system of EGL includes the following:

1. The rules and axioms of propositional logic
2. Kripke's axioms for the two modalities Q_a and Dk_G:
   - $Q_a(\varphi \Rightarrow \psi) \Rightarrow (Q_a\varphi \Rightarrow Q_a\psi)$
   - $Dk_G(\varphi \Rightarrow \psi) \Rightarrow (Dk_G\varphi \Rightarrow Dk_G\psi)$
3. Necessitation Rules for the two modalities Q_a and Dk_G:
   - from φ infer $Q_a\varphi$
   - from φ infer $Dk_G\varphi$
4. $Dk_G\varphi \Rightarrow \varphi \wedge Dk_G Dk_G\varphi$
5. $Q_a\varphi \Rightarrow \varphi \wedge Q_a Q_a\varphi$
6. $K_a\varphi \Rightarrow Q_a\varphi$
7. $Dk_G Q_b\varphi \Rightarrow Q_b Dk_G\varphi$
8. $Q_b Q_a\varphi \Rightarrow Q_a Q_b\varphi$
9. $Dk_G\varphi \Rightarrow Dk_{G'}\varphi$ (for $G \subseteq G'$)

7.2 Completeness for EGL

Proposition 11. *EGL is a sound and complete axiomatization with respect to the class of epistemic group models.*

7.2.1 Main ideas of the proof

The completeness proof is based on a more sophisticated version of the method used by Fagin *et al.* (1992) in their original completeness proof for the standard logic of distributed knowledge. The current proof requires additional steps in order to accommodate conditions $P1, P2'$ and $P3$. Here, I briefly explain the main steps of the proof by Fagin *et al.* as it helps clarify the ideas behind the current proof. I then explain the main ideas (and steps) of the proof for soundness and completeness for EGL.

The proof method used by Fagin *et al.* (1992) involves three steps. First, a more general notion of model is considered (called "pseudo-model"), in which each distributed knowledge operator Dk_G receives its own accessibility relation $\to_G$, which is *not* necessarily equal to the intersection of the basic accessibility relations in the group. Apart from this, pseudo-models satisfy all the other semantic conditions required of a model (e.g. reflexivity, transitivity etc.). Completeness for pseudo-models is proved via the standard canonical (pseudo-)model construction. Then, in Step 2, this pseudo-model is unraveled into a tree, made of all possible finite "histories" (i.e. finite paths from a given fixed world w_0 to any arbitrary world w_n in the pseudo-model). The tree is naturally endowed with accessibility relations R_G going one-step on the branches of the tree, i.e. from any path of the form $(w_0, \to_{G_0}, w_1, \ldots, w_n)$ to extended paths of the form $(w_0, \to_{G_0}, w_1, \ldots, w_n, \to_{G_n}, w_{n+1})$ with $w_n \to_G w_{n+1}$. This "tree model" no longer satisfies any of the original semantic conditions (so it is not even a pseudo-model), but it preserves the truth of modal formulas, including (all the axioms and) the formulas that were satisfied at w_0 (since there exists a bounded morphism from the tree model to the pseudo-model, defined by mapping any path $(w_0, \to_{G_0}, w_1, \ldots, w_n)$ to its end-point w_n). Moreover, the tree model satisfies a new condition, called "the tree property": between any two worlds there is at most one path. In Step 3, the tree model is modified by changing only the accessibility relations in such a way that they satisfy all the original semantic conditions (thus obtaining a full-fledged model in the original sense). More precisely, new individual knowledge relations $\to_a$ are defined by taking the reflexive-transitive closure of the union of all R_G relations with $a \in G$. The tree property plays a key role in proving that this construction preserves the truth of sentences in the language of epistemic logic with distributed knowledge. The completeness result is thus "transferred" from pseudo-models to the intended models via the intermediary step provided by tree models.

The first step of the proof of completeness for EGL is the same as in Fagin *et al.* (1992), except that there are now more relations to consider since the pseudo-models have both independent relations $\to_G$ for Dk_G and relations $\approx_a$ for Q_a. These relations satisfy the same semantic constraints as epistemic group models, except for the fact that $\to_G$ is not necessarily equal to $\bigcap_{a \in G} \to_a$. Completeness for pseudo-models is proved by a canonical model construction. However, the tree construction has to be modified in order to satisfy conditions $P2'$ and $P3$. This is done as follows. After unraveling the pseudo-model, the tree is immediately "collapsed back" by identifying histories that have the same end-point w_n and the same "path-type". In particular, such histories have the same number of arrows of each kind, as well as their last world (from the pseudo-model). This structure still has accessibility relations $R_G^{\to}$ and $R_a^{\approx}$ going from each history type to its corresponding one-step extensions. This "model" no longer has the tree property, but it instead has a "tree-like" property: between any two worlds there is at most one "type" of path (i.e. all paths between any w and w' have the same number of arrows of each kind). For brevity, I do not explicitly present the intermediary tree models. Instead, step 2 provides preliminary definitions and results on paths (corresponding to arrows) and "path types" in pseudo-models. These definitions and results are used in step 3. In this step, an epistemic group model is constructed from a pseudo-model in one blow. New issue relations $\approx_a$ are defined as the reflexive-transitive closure of the corresponding

$R_a^{\approx}$, and new knowledge relations $\rightarrow_a$ are defined as the reflexive-transitive closure of the union of $\approx_a$ and all $R_G^{\rightarrow}$ relations for $a \in G$. The "worlds" of this epistemic group model consist of pairs (γ, w), where w is any world in the original pseudo-model and γ is a "path-type" of some path p from a fixed designated world w_0 to w. (This construction corresponds exactly to identifying two histories in the tree model iff they have the same end-point and the same path type.) The issue relations $\approx_a$ and knowledge relations $\rightarrow_a$ are defined directly as binary relations between such pairs (γ, w) and (γ', w'). The preliminary results from step 2 are then used to prove that this construction preserves the truth of EGL formulas.

7.2.2 The proof

Step 1: soundness and completeness for pseudo-models

The first step is to define "pseudo-model" in which each modality get its 'own' model (thus $A+1$ single-agent logics are combined).

Definition 18 (pseudo-model)**.** A pseudo-model is a structure $M = (S, \rightarrow_G, \approx_a \| \bullet \|)$ consisting of a finite set of states S; a finite set of agents A, where $G \subseteq A$; binary accessibility relations $\rightarrow_G$ and $\approx_a$ and a valuation function $\| \bullet \|$, that satisfies the following constraints:

- $\rightarrow_G$ and $\approx_a$ are reflexive and transitive relations;
- $\approx_a \subseteq \rightarrow_{\{a\}}$;
- $\approx_a ; \rightarrow_G \subseteq \rightarrow_G ; \approx_a$;
- $\approx_a ; \approx_b \subseteq \approx_b ; \approx_a$.
- $\rightarrow_G \subseteq \rightarrow_{G'}$ for $G' \subseteq G$.

I adopt the following additional notation. Given a pseudo-model $M = (S, \rightarrow_G, \approx_a \| \bullet \|)$, I put $\rightarrow_a := \rightarrow_{\{a\}}$, and $K_a := Dk_{\{a\}}$.

Definition 19 (Pseudo-satisfaction)**.** The pair (M, s) pseudo-satisfies a formula φ, written as $s \models_M \varphi$, whenever it satisfies φ in the usual way, except that $s \models_M Dk_G\varphi$ iff $s' \models_M \varphi$ for all s' such that $s \rightarrow_G s'$.

Proposition 12. Every epistemic group model "is" a pseudo-model. More precisely: let $M = (S, \rightarrow_a, \approx_a \| \bullet \|)$ be an epistemic group model. Then $M' = (S, \rightarrow_G, \approx_a \| \bullet \|)$, where $\rightarrow_G := \bigcap_{a \in G} \rightarrow_a$ for any $G \subseteq A$, is called the *pseudo-model associated with* M, such that $\forall s \in S$ and all sentences φ:

$$s \models_M \varphi \text{ iff } s \models_{M'} \varphi.$$

Proof. The fact that M' is a pseudo-model follows directly from the definition of epistemic group models. The semantic clauses for each connective are essentially the same in M and M'. Hence it follows by induction that the truth of modal formulas is preserved when going from M' to M. $\square$

As such, in order to prove soundness for epistemic group models, it suffices to show that pseudo models are sound.

Proposition 13. All axioms of the proof system of EGL are valid on pseudo-models.

Proof. The axioms will be treated individually:

($Dk_G\varphi \rightarrow \varphi$) Suppose $s \models_M Dk_G\varphi$. By reflexivity of $\rightarrow_G$ it follows that $s \models_M \varphi$ and thus $s \models_M Dk_G\varphi \rightarrow \varphi$.

($Dk_G\varphi \rightarrow Dk_G Dk_G\varphi$) Suppose $s \models_M Dk_G\varphi$. Let t be an arbitrary state such that $s \rightarrow_G t$. It is enough to show that $t \models_M Dk_G\varphi$. For this, let w be an arbitrary state such that $t \rightarrow_G w$. It is enough to show that $w \models_M \varphi$. But, this follows from $s \rightarrow_G t \rightarrow_G w$ by transitivity of $\rightarrow_G$.

($Q_a\varphi \rightarrow \varphi$) Proof is similar to the Dk_G case.

($Q_a\varphi \rightarrow Q_aQ_a\varphi$) Proof is similar to the Dk_G case.

($K_a\varphi \rightarrow Q_a\varphi$) Suppose $s \models_M Dk_G\varphi$. Take any arbitrary world t such that $s \approx_a t$. By P1 it follows that $s \rightarrow_{\{a\}} t$. Hence we get $t \models_M \varphi$ for all t such that $s \approx_a t$, and so: $s \models_M Q_a\varphi$.

($Dk_GQ_a\varphi \rightarrow Q_aDk_G\varphi$) Suppose $s \models_M Dk_GQ_a\varphi$. Let t be an arbitrary state such that $s \approx_a t$. To show that $s \models_M Q_aDk_G\varphi$, it is enough to show that $t \models_M Dk_G\varphi$. To show this, let w be any arbitrary state such that $t \rightarrow_G w$. It is enough to show that $w \models_M \varphi$. By $P2$ there exists some t' such that $s \rightarrow_G t' \approx_a w$. Given the assumption that $s \models_M Dk_GQ_a\varphi$, it follows that $t' \models_M Q_a\varphi$ and $w \models_M \varphi$.

($Q_bQ_a\varphi \rightarrow Q_aQ_b\varphi$) Suppose $s \models_M Q_aQ_b\varphi$. Let t be an arbitrary state such that $s \approx_a t$. To show that $s \models_M Q_bQ_a\varphi$, it is enough to show that $t \models_M Q_b\varphi$. To show this, let w be any arbitrary state such that $t \approx_b w$. It is enough to show that $w \models_M \varphi$. By $P3$ there exists some t' such that $s \approx_b t' \approx_a w$. Given the assumption that $s \models_M Q_aQ_b\varphi$, it follows that $t' \models_M Q_b\varphi$ and $w \models_M \varphi$.

($Dk_G\varphi \rightarrow Dk_{G'}\varphi$ (for all $a \in G \subseteq G'$)) Suppose $s \models_M Dk_G\varphi$. Let t be such that $s \rightarrow'_G t$. It follows that $s \rightarrow_a t$ for all $a \in G'$. Together with $G \subseteq G'$, it follows that $s \rightarrow_a$ for all $a \in G$. Hence: $s \rightarrow_G t$. By the semantics of Dk_G, it follows from $s \models_M Dk_G\varphi$ that $t \models_M \varphi$. For any arbitrary t such that $s \rightarrow'_G t$ it follows that $s \models_M Dk'_G\varphi$.

□

To prove completness with respect to pseudo-models, I construct a "canonical pseudo-model", M^C, which is the canonical Kripke structure for EGL. In this model, every EGL-consistent formula is satisfiable. States s^c in M^C are maximally consistent sets of EGL formulas. In order to show completeness with respect to this structure (from which completeness of pseudo-models follows), some preliminary definitions are needed:

Definition 20. A formula φ is *inconsistent* if its negation $\neg\varphi$ can be proven in EGL. Otherwise, φ is *consistent.*

Definition 21. A set of formulas is *maximally consistent* if it is a consistent set of formulas Φ such that whenever a formula ψ is not in Φ, then $\{\psi\} \cup \Phi$ is inconsistent.

Lemma 1 (Lindenbaum's lemma)**.** If Φ is a consistent set of formulas, then there exists a maximally consistent set of formulas Φ^+ that contains Φ, i.e., $\Phi \subseteq \Phi^+$.[1]

Definition 22 ("Canonical pseudo-model")**.** The "canonical pseudo-model" is a structure $M^C = (W, \to_G, \approx_a \| \bullet \|)$ such that:

- $W = \{w : w \text{ is a maximally consistent set of } EGL \text{ formulas}\}$
- $w \to_G w'$ iff $\forall\varphi(Dk_G\varphi \in w \Rightarrow \varphi \in w')$
- $w \approx_a w'$ iff $\forall\varphi(Q_a\varphi \in w \Rightarrow \varphi \in w')$
- $\|p\| = \{w : w \in M^C : p \in w\}$

Note that the canonical relations can be alternatively characterized as follows (Blackburn *et al.* 2001: 200):

- $w \to_G w'$ iff $\forall\theta(\theta \in w' \Rightarrow < Dk_G > \theta \in w)$
- $w \approx_a w'$ iff $\forall\theta(\theta \in w' \Rightarrow < Q_a > \theta \in w)$

Proposition 14. The "canonical pseudo-model" is a pseudo-model.

Proof. The semantic properties are treated individually:

To show that $\to_G$ is reflexive, suppose $Dk_G\varphi \in w$. Then by our axioms, it follows that $\varphi \in w$. By the definition of $\to_G$ (in M^C), it follows from this that $id \subseteq \to_G$, i.e., that $\to_G$ is reflexive.

To show that $\to_G$ is transitive, suppose that $s \to_G w \to_G t$. Suppose $Dk_G\varphi \in s$. Then, it follows by axiom 4 that $Dk_G Dk_G\varphi \in s$, and moreover, that $Dk_G\varphi \in w$ and $\varphi \in t$. By the definition of $\to_G$ (in M^C) it follows that $s \to_G t$.

The proofs for reflexivity and transitivity of $\approx_a$ are similar to the $\to_G$ case.

To show that $\approx_a \subseteq \to_{\{a\}}$, suppose that $s \approx_a t$. Let $Dk_{\{a\}}\varphi \in s$. Since $(Dk_{\{a\}}\varphi \to_{\{a\}} Q_a\varphi)$ is one of the axioms and t is a maximally consistent set, if $Dk_{\{a\}}\varphi \in s$ then so is $Q_a\varphi \in s$. Since $s \approx_a t$, it follows from $Q_a\varphi \in s$ that $\varphi \in t$.

To show that $\approx_a; \to_G \subseteq \to_G; \approx_a$, suppose that $w \approx_a s \to_G t$. Then, the first thing to show is the existence of some $s' \in W$ such that $w \to_G s' \approx_a t$. Let

[1] For the proof, see (Blackburn *et al.*, 2001).

$\Phi = \{\varphi \,|\, Dk_G\varphi \in w\} \cup \{< Q_a > \psi \,|\, \psi \in t\}$. The next step is to show that Φ is consistent. Suppose not: $\Phi \vdash \bot$. This means that $\exists(\varphi_1, ...\varphi_n, \psi_1, ...\psi_n)$ such that: $(Dk_G\varphi_1, ...Dk_G\varphi_n) \in w, (< Q_a > \psi_1, ..., < Q_a > \psi_n) \in t$, and $\vdash \varphi_1 \wedge ... \wedge \varphi_n \wedge < Q_a > \psi_1 \wedge ... < Q_a > \psi_n \to \bot$. Further, let $\varphi = \varphi_1 \wedge ... \wedge \varphi_n$ and let $< Q_a > \psi =< Q_a > \psi_1 \wedge ... \wedge < Q_a > \psi_n$. We thus have that $Dk_G\varphi \in w$ and $\psi \in t$, and that $\varphi \wedge < Q_a > \psi \to \bot$. It follows that $\varphi \to Q_a\neg\psi$. Since w is maximally consistent, it follows that $Dk_G\varphi \to Dk_GQ_a\neg\psi$, where $Dk_G\varphi \in w$ and $Dk_GQ_a\neg\psi \in w$. It follows by $P2$ that $Q_aDk_G\neg\psi \in w$, and thus that $Dk_G\neg\psi \in s$ (since $w \approx_a s \to_G t$. From this it follows that $\neg\psi \in t$. However, this contradicts the assumption that $\psi \in t$, and so Φ must be consistent. Next, by Lindenbaum's lemma, since Φ is consistent, it follows that $\exists s' \in W$ such that $\Phi \subseteq s'$. But, by the choice of Φ, $\Phi \subseteq s'$ implies both of the following: $\forall\varphi(Dk_G\varphi \in w \Rightarrow \varphi \in s')$ and $\forall\psi(Q_a\psi \in t \Rightarrow < Q_a > \psi \in s')$. Hence, $w \to_G s' \to_G t$.

The proof for $\approx_a; \approx_b \subseteq \approx_b; \approx_a$ is similar to that of $\approx_a; \to_G \subseteq \to_G; \approx_a$. ☐

Corollary 1. The EGL axioms are valid on the canonical pseudo-model.

Lemma 2 (Truth lemma)**.** $w \models_{M^C} \varphi$ iff $\varphi \in w$.[2]

Proposition 15. If φ is consistent, then φ is pseudo-satisfiable.

Proposition 16. The logic EGL is sound and complete with respect to pseudo-models.

Proof. Soundness follows from propsitions 12 and 13. For completeness, suppose Φ is a consistent set of formulas from L_{EGL}. It must be shown that Φ is pseudo-satisfiable. By Lindenbaum's lemma, it follows that $\exists w \in W : \Phi \subseteq w$. By the truth lemma, it follows from $\Phi \in w$ that $w \models_{M^C} \Phi$. This means that Φ is true in some world in some model, namely, the canonical model. Hence Φ is also satisfiable in pseudo-models (since the canonical model is a special case of pseudo-models). ☐

Step 2: paths and "path types" in pseudo-models

Definition 23 (Paths in pseudo-models)**.** Given a pseudo-model $M = (S, \to_G, \approx_a \|\bullet\|)$, and two worlds $w, v \in S$, a *path from w to v* is a sequence $p = (w_0, R_0, w_1, ...R_{n-1}, w_n)$, of length $n \geq 0$, where

- $w_0 = w$
- $w_n = v$
- all $R_i \in \{\to_G : G \subseteq A\} \cup \{\approx_a : a \in A\}$

I adopt the following additional notation. For any path $p = (w_0, R_0, w_1, ...R_{n-1}, w_n)$: $\text{last}(p) := w_n$ and $\text{first}(p) := w_0$.

[2]For the proof, see (Blackburn *et al.*, 2001).

Definition 24 (Path composition)**.** Given a path $p = (w_0, R_0, w_1, ...R_{n-1}, w_n)$ from $w = w_0$ to $v = w_n$, and another path $p' = (w'_0, R'_0, w'_1, ...R'_{n-1}, w'_n)$ from $v = w'_0$ to $s = w'_n$, the *composed path* $p \cdot p' = (w_0, R_0, w_1, ..., w_n = w'_0, R'_0, ...R'_{n-1}, w'_n)$ is a path from w to s.

Note that for two paths p and q to be composable, p and q must satisfy $\text{last}(p) = \text{first}(q)$.

Definition 25 ("Path types")**.** A *"path type"* is a map $\gamma : A \cup \mathcal{P}(A) \to \mathbb{N}$.

Definition 26 (Path types)**.** Given a path $p = (w_0, R_0, w_1, ...R_{n-1}, w_n)$ from $w = w_0$ to $v = w_n$, the *type of* p is a path type $\gamma^p : A \cup \mathcal{P}(A) \to \mathbb{N}$, given by recursion:

- $\gamma^p(x) = 0$ for all x, if $p = (w, w)$ is a path of lenght 0;
- If either $x \subseteq A$ and $R = \to_x$, or $x \in A$ and $R = \approx_x$, then $\gamma^{p\cdot(v,R,s)}(x) = \gamma^p(x) + 1$.
- Else*:* $\gamma^{p\cdot(v,R,s)}(x) = \gamma^p(x)$.

So, γ^p counts the number of relations of each kind ($\approx_a$ or $\to_G$) in the path p. It only keeps track of the number of arrows for each agent and the number of arrows for each group in the path, not the order.

Definition 27 (Operations with path types)**.** For $\gamma, \gamma' : A \cup \mathcal{P}(A) \to \mathbb{N}$, path types $\gamma + \gamma'$ and $\gamma - \gamma'$ are such that for all x:

- $\gamma(x) + \gamma'(x) := \gamma(x) + \gamma'(x)$
- $\gamma(x) - \gamma'(x) := \gamma(x) - \gamma'(x)$

The operation $\gamma - \gamma'(x)$ is not always well-defined, but only when $\gamma(x) \geq \gamma'(x)$ for all x. Moreover, the following holds:

- $(\gamma + \gamma') - \gamma' = \gamma(x)$
- $(\gamma + \gamma') = \gamma' + \gamma$
- $(\gamma + \gamma') + \gamma'' = \gamma + (\gamma' + \gamma'')$
- $\gamma^{p\cdot p'} = \gamma^p + \gamma^{p'}$
- $\gamma^{p'} = \gamma^{p\cdot p'} - \gamma^p$ (whenever $p \cdot p'$ is defined).

Given these definitions, three kinds of path are defined, which correspond to the arrows in the new models (to be constructed in the final step).

Definition 28 ($\overset{\approx}{a}$-paths)**.** An $\overset{\approx}{a}$-path from w to v is any path p (from w to v) such that for all x: $\gamma^p(x) \neq 0 \Rightarrow x = a$.

The count of γ^p of an $\overset{\approx}{a}$-path is non-zero only for $\approx_a$: an $\overset{\approx}{a}$-path only contains $\approx_a$.

Definition 29 ($\vec{a}$-paths)**.** An $\vec{a}$-path from w to v is any path p (from w to v) such that for all x: $\gamma^p(x) \neq 0 \Rightarrow (x = a$ or $x \in \{G \subseteq A : a \in G\}$.

Intuitively, an $\vec{a}$-path is a sequence of $\approx_a$'s and $\rightarrow_G$'s such that $a \in G$. Note that any $\overset{\approx}{a}$-path is also an $\vec{a}$-path, as the former is a special case of the latter.

Definition 30 ($\vec{G}$-paths)**.** A $\vec{G}$-path from w to v is any path p (from w to v) such that it is simultaneously an $\vec{a}$-path for all $a \in G$.

Note that a $\vec{G}$ does not necessarily consist of $\rightarrow_G$ arrows. Moreover, the following observations can be made:

- If p is an $\overset{\approx}{a}$-path and $\gamma^p = \gamma^{p'}$, then p' is an $\overset{\approx}{a}$-path.
- If p is an $\vec{a}$-path and $\gamma^p = \gamma^{p'}$, then p' is an $\vec{a}$-path.

I adopt the following notation: The cardinality of a set P is denoted by '$|P|$'.

Proposition 17. Let G be a group with $|G| \geq 2$. A path p is a $\vec{G}$-path iff it is of the form $p = (w = w_0, \rightarrow G_0, w_1, \rightarrow G_1 ... \rightarrow G_{n-1}, w_n = v)$ for some groups G_i such that $G \subseteq G_i$ for all i.

As a consequence, when $|G| \geq 2$, no $\vec{G}$-path can contain any $\approx_a$'s.

Proof. ($\Rightarrow$) Choose $a, b \in G$ such that $a \neq b$. The following two claims must be shown: (1) $\gamma^p(c) = 0$ for every $c \in A$ (i.e., no $\approx_a$), and (2) $\gamma^p(H) = 0$ for every $H \subseteq A$ such that $G \not\subseteq H$ (i.e., all G_i contain G).

(1). Let $c \in A$. This means that either $c \neq a$ or $c \neq b$. In the first case, since p is an $\vec{a}$-path (because $a \in G$ and p is a $\vec{G}$-path), it follows that $\gamma^p(c) = 0$. In the second case ($c \neq b$), since p is also a $\vec{b}$-path, it also follows that $\gamma^p(c) = 0$.

(2). Let $G \not\subseteq H$. Choose some $b \in G$ such that $b \notin H$. But p is a $\vec{G}$-path and $b \in G$, so p is a $\vec{b}$-path. Hence: $\gamma^p(H) = 0$. Putting (1) and (2) together, it follows that for every i between 0 and $n - 1$ there exists some G_i such that $S_i = \rightarrow_{G_i}$ and $G \subseteq G_i$

($\Leftarrow$) Let $p = (w = w_0, \rightarrow G_0, w_1, \rightarrow G_1 ... \rightarrow G_{n-1}, w_n = v)$ for some groups G_i such that $G \subseteq G_i$ for all i. To show that p is a $\vec{G}$-path, it is enough to show that p is an $\vec{a}$-path for every $a \in G$. Now, if $a \in G$ then $a \in G_i$ for all i (since $G \subseteq G_i$). Hence p is by definition an $\vec{a}$-path. □

Proposition 18. If there is an $\overset{\approx}{a}$-path from w to v, then $w \approx_a v$.

Proof. Let $p = (w = w_0, S_0, w_1, S_1 ... S_{n-1}, w_n = v)$ be an $\overset{\approx}{a}$-path from w to v. Then, by definition 28, $S_i = \approx_a$ for all i. This means that $p = (w = w_0 \approx_a w_1, \approx_a ... \approx_a w_n = v)$. Then, by transitivity of $\approx_a$ it follows that $w \approx_a v$. □

Proposition 19. If there is an $\vec{a}$-path from w to v, then $w \rightarrow_a v$.

Proof. Let $p = (w = w_0, S_0, w_1, S_1 ... S_{n-1}, w_n = v)$ be an $\vec{a}$-path from w to v. Then, by definition 29, $S_i \in \{\approx_a\} \cup \{\rightarrow_G : a \in G\}$. Together with w_{i-1}, S_{i-1}, w_i for all i, this means that $w_{i-1} \rightarrow_a w_i$ for all i (since $\approx_a \subseteq \rightarrow_a$ and $\rightarrow_G \subseteq \rightarrow_{\{a\}} = \rightarrow_a$ for $a \in G$). Then, by transitivity of $\rightarrow_a$, it follows that $w \rightarrow_a v$. □

Proposition 20. If there exists some $\vec{G}$-path from w to v, then $w \rightarrow_G v$ (in the pseudo-model M).

Proof. Let $p = (w = w_0, S_0, w_1, S_1 ... S_{n-1}, w_n = v)$ be some $\vec{G}$-path from w to v. There are two cases to consider, namely, $|G| = 1$ and $|G| \geq 2$.

($|G| = 1$). Let $G = \{a\}$. Then p is an $\vec{a}$-path from w to v. By proposition 19, this means that $w \rightarrow_a v$. This is the same as $w \rightarrow_{\{a\}} v$, hence also as $w \rightarrow_G v$.

($|G| \geq 2$). By proposition 17, p is of the form $p = (w = w_0, G_0, w_1, ... G_{n-1}, w_n = v)$ for some group G_i such that $G \subseteq G_i$. This means that $w_{i-1} \rightarrow_{G_{i-1}} w_i$ (for all i). Since $G \subseteq G_i$, it follows that $w_{i-1} \rightarrow_G w_i$. And thus: $p = (w = w_0, \rightarrow_G, w_1, \rightarrow_G ... \rightarrow_G, w_n = v)$. Then, by transitivity of $\rightarrow_G$, it follows that $w \rightarrow_G v$. □

Lemma 3. Let $w, v \in S$, $G \subseteq A$ and $\{p_a : a \in G\}$ be some family of paths such that for all $a \in G$:

- p_a is an $\vec{a}$-path from w to v
- $\gamma^{p_a} = \gamma$

Then there exists some $\vec{G}$-path p from w to v such that $\gamma^p = \gamma$.

Proof. Given some fixed $a \in G$, it must be shown that p_a is a $\vec{G}$-path from w to v. Indeed, by assumption, $\gamma^{p_a} = \gamma^{p_b}$ for all $b \in G$, and hence by previous observations, γ^{p_a} is a $\vec{b}$-path for any $b \in G$. So (by definition) γ^{p_a} is a $\vec{G}$-path. □

Lemma 4. If $w \approx_a v$, then for every $\overset{\approx}{a}$-path p there exists some path p' from w to v such that $\gamma^p = \gamma^{p'}$.

Proof. Let $p = (w, \approx_a, w_1, \approx_a ... \approx_a, w_n = v)$ be any $\overset{\approx}{a}$-path p (of some length n). Then the path $p' = (w, \approx_a, w_1, \approx_a ... \approx_a, w_n = v)$ (also of length n) is clearly a path from w to v, and $\gamma^p = \gamma^{p'}$. □

Given the preliminary definitions and propositions thus far, the next (and final preliminary) proposition generalizes conditions $P2'$ and $P3$ to paths. It shows that $P2'$ and $P3$ preserve path types.

Proposition 21. Suppose that p is an $\widetilde{\widetilde{a}}$-path from w to v and q is some path from v to s. Then there exists some world $v' \in S$ and some paths q', p' such that:

- q' is a path from w to v' that has the same type as q (i.e. $\gamma^q = \gamma^{q'}$),
- p' is a path from v' to s that has the same type as p (i.e. $\gamma^p = \gamma^{p'}$).

Proof. The proof is by induction on the length of q.

(Base case). Let length$(q) = 0$. Then $p = (w, \approx_a, w_1, \approx_a ... \approx_a, w_n = v)$, $q = (v) = (s)$. Take $q' = (w = w_0)$, then $p' = (w, \approx_a, w_1, \approx_a ... \approx_a, v = s) = p$.

(Inductive step). Let length$(q) = n + 1$. Then $q = \bar{q} \cdot (t, R_n s)$ for some path $\bar{q}$ of length n from v to t (for some new t). By applying the induction hypothesis to p and $\bar{q}$, it follows that there exists some v'', some $\bar{q'}$ and some p'', such that:

- $\bar{q'}$ is a path from w to v'';
- p'' is an $\widetilde{\widetilde{a}}$-path from v'' to t;
- $\gamma^{\bar{q'}} = \gamma^{\bar{q}}$;
- $\gamma^{p''} = \gamma^p$.

By proposition 18, $v'' \approx_a t$ (since there exists some $\widetilde{\widetilde{a}}$-path from v'' to t). Thus: $v'' \approx_a tR_n s$. By $P2'$ and $P3$, it follows that there exists some v' such that $v'' R_n v' \approx_a s$. Take $q' := \bar{q'} \cdot (v'', R_n v')$ and $p' :=$ some $\widetilde{\widetilde{a}}$-path from v' to s of the same type as p. By lemma 4 such a p' exists, because $v' \approx_a s$ and p are both $\widetilde{\widetilde{a}}$-paths. Then, q' is indeed a path from w to v' and $\gamma^{q'} = \gamma^{\bar{q'}} + \gamma^{(v'', R_n, v')} = \gamma^{\bar{q}} + \gamma^{(v, R_n, s)} = \gamma^{\bar{q} \cdot (v, R_n, s)} = \gamma^q$. Also, p' is an $\widetilde{\widetilde{a}}$-path from v' to s with $\gamma^{p'} = \gamma^p$. $\square$

Step 3: the construction of the epistemic group model

Definition 31 (Epistemic group model)**.** Given any pseudo-model $M = (S, \rightarrow_G, \approx_a, \| \bullet \|)$ and a fixed world $w_0 \in S$, its associated epistemic group model is defined as $\bar{M} = (\bar{S}, \rightarrow_a, \approx_a, \| \bullet \|)$, with:

- $\bar{S} = \{(\gamma, w) : w \in S, \gamma = \gamma^p$ for some path p from w_0 to $w\}$.
- $(\gamma, w) \rightarrow_a (\gamma', w')$ iff $\gamma' = (\gamma + \gamma^p)$ for some $\vec{a}$-path p from w to w'.
- $(\gamma, w) \approx_a (\gamma', w')$ iff $\gamma' = (\gamma + \gamma^p)$ for some $\widetilde{\widetilde{a}}$-path p from w to w'.
- $\|p\|_{\bar{M}} = \{(\gamma, w) \in \bar{S} : w \in \|p\|_M\}$.

Thus, M and $\bar{M}$ have a different state space. Worlds in $\bar{S}$ are pairs (γ, w) of path types γ and end-worlds w (from M).

Definition 32 (Bounded morphism)**.** Given pseudo-models $M = (S, \rightarrow_G, \approx_a, \| \bullet \|)$ and $M' = (S', \rightarrow_G, \approx_a, \| \bullet \|)$, a *bounded morphism* is a function $F : S \rightarrow S'$ that satisfies the following three properties:

- Preservation of valuation: $w \in |p\|_M$ iff $F(w) \in \|p\|_{M'}$,
- "Forth" condition for relations $\rightarrow_G$ and $\approx_a$: for all groups $G \subseteq A$ and agents $a \in A$ it holds that $w \rightarrow_G s$ implies $F(w) \rightarrow_G F(s)$ and $w \approx_a s$ implies $F(w) \approx_a F(s)$;
- "Back" condition for relations $\rightarrow_G$ and $\approx_a$: for all groups $G \subseteq A$ and agents $a \in A$ it holds that $F(w) \rightarrow_G s'$ implies that there exists some $s \in S$ such that $w \rightarrow_G s$ and $F(s) = s'$ and $F(w) \approx_a s'$ implies that there exists some $s \in S$ such that $w \approx_a s$ and $F(s) = s'$.

Proposition 22. Bounded morphisms preserve the truth of EGL formulas. More precisely, if $F : S \rightarrow S'$ is a bounded morphism between pseudo-models M and M', $w \in S$ is a world in the first model and φ is a sentence in the syntax of EGL, then the following holds:

$$w \models_M \varphi \,\text{iff}\, F(w) \models_{M'} \varphi.$$

Proof. This is an instance of the general result that modal formulas are preserved by bounded morphisms between Kripke models (Blackburn *et al.* 2001: 79). The proof is by induction on the complexity of φ, using the preservation of valuation property to deal with atomic sentences, and the back-and-forth properties to deal with the modalities (here, Dk_G and Q_a). For more details of this proof, see (Blackburn *et al.* 2001). □

To prove completeness for epistemic group models, it suffices to show that (1) $\bar{M}$ is an epistemic group model (i.e., that it has the required semantic properties), and that (2) there exists a bounded morphism from $\bar{M}$ to M having w_0 in its range – hence, the pre-image of w_0 in $\bar{M}$ satisfies the same formulas as w_0 in M. I start by showing (1), i.e., that $\bar{M}$ has the semantic properties of epistemic group models.

Proposition 23. $\bar{M}$ is an epistemic group model, with $\rightarrow_G = \bigcap_{a\in G} \rightarrow_a$.

Proof. The semantic properties are treated individually:

(Reflexivity of $\approx_a$ and $\rightarrow_a$). For $(\gamma, w) \in S$, take the 0-length path $p = (w, w)$.

(Transitivity of $\approx_a$ and $\rightarrow_a$). Suppose that $(\gamma, w) \rightarrow_a (\gamma', w') \rightarrow_a (\gamma'', w'')$. Then $\gamma' = \gamma + \gamma^p$ for some $\vec{a}$-path p from w to w' and $\gamma'' = \gamma' + \gamma^q$ for some $\vec{a}$-path q from w' to w''. It follows that $p \cdot q$ is an $\vec{a}$-path from w to w'', with $\gamma + \gamma^{p\cdot q} = \gamma + (\gamma^p + \gamma^q) = (\gamma + \gamma^p) + \gamma^q = \gamma' + \gamma_q = \gamma''$. Hence $(\gamma, w) \rightarrow (\gamma'', w'')$. The proof for $\approx_a$ is similar.

($\approx_a \subseteq \to_a$). Suppose that $(\gamma, w) \approx_a (\gamma', w')$. Then $\gamma' = \gamma + \gamma^p$ for some $\overset{\approx}{a}$-path p from w to w'. By a previous observation, every $\overset{\approx}{a}$-path is also an $\vec{a}$-path, and hence it follows that $(\gamma, w) \to_a (\gamma', w')$.

($\approx_a; \to_G \subseteq \to_G; \approx_a$, where $\to_G = \bigcap_{a \in G} \to_a$ in $\bar{M}$). Suppose that $(\gamma, w) \approx_a (\gamma', v) \to_G (\gamma'', s)$. Then, from $(\gamma, w) \approx_a (\gamma', v)$ it follows that there exists some $\overset{\approx}{a}$-path p from w to v with $\gamma^p = \gamma' - \gamma$. Thus, $(\gamma', v) \to_b (\gamma'', s)$ for all $b \in G$. It follows from this that there exists some family $\{q_b : b \in G\}$ of paths such that for all $b \in G$: q_b is a $\vec{b}$-path from v to s and $\gamma^{q_b} = \gamma'' - \gamma'$. By lemma 3, it follows that there exists some $\vec{G}$-path q from v to s, with $\gamma^q = \gamma'' - \gamma'$ and $v \to_G s$.

By applying proposition 21 to the worlds w, v, s and paths p, q, there must exist some world $v' \in S$ and there must exist paths q', p' such that:

- q' is a path from w to v'
- p' is a path from v' to s
- $\gamma^q = \gamma^{q'}$ and so q' is a $\vec{G}$-path (since q is a $\vec{G}$-path).
- $\gamma^p = \gamma^{p'}$ and so p' is an $\overset{\approx}{a}$-path (since p is an $\overset{\approx}{a}$-path).

Thus, it follows that: $(\gamma, w) \to_G (\gamma + \gamma^{q'}, v') \approx_a (\gamma + \gamma^{q'} + \gamma^{p'}, s) = (\gamma + \gamma^q + \gamma^p, s) = ((\gamma + \gamma^q) + \gamma^p, s) = (\gamma' + \gamma^q, s) = (\gamma'', s)$. Hence: $(\gamma, w) \to_G (\gamma + \gamma^{q'}, v') \approx_a (\gamma^q, s)$.

($\approx_a; \approx_b \subseteq \approx_b; \approx_a$). Suppose that $(\gamma, w) \approx_a (\gamma', v) \approx_b (\gamma'', s)$. This means that there exists some $\overset{\approx}{a}$-path p from w to v such that $\gamma^p = \gamma' - \gamma$, and that there exists some $\overset{\approx}{b}$-path q from v to s such that $\gamma^q = \gamma'' - \gamma'$. By proposition 18, it follows that $w \approx_a v$ and $v \approx_b s$. By condition $P3$ for M, it follows that there exists some v' such that $w \approx_b v' \approx_a s$. Since q is a $\overset{\approx}{b}$-path and $w \approx_b v'$, it follows by lemma 4 that there exists some path q' from w to v' such that $\gamma^{q'} = \gamma^q$. Similarly, since p is an $\overset{\approx}{a}$-path and $v' \approx_a s$, it follows by lemma 4 that there exists some path p' from v' to s such that $\gamma^{p'} = \gamma^p$. Hence $(\gamma, w) \approx_b (\gamma + \gamma^{q'}, v') \approx_a (\gamma + \gamma^{q'} + \gamma^{p'}, s) = (\gamma + \gamma^q + \gamma^p, s) = (\gamma' + \gamma^{q'}, s) = (\gamma'', s)$. Done! $\square$

Proposition 24. Let $M = (S, \to_G, \approx_a, \| \bullet \|)$ be a pseudo-model, with $w_0 \in S$. Let $\bar{M} = (\bar{S}, \to_a, \approx_a, \| \bullet \|)$ be the epistemic group model obtained by applying to M and w_0 the above construction, and let $\bar{M}' = (\bar{S}, \to_G, \approx_a, \| \bullet \|)$ be the pseudo-model obtained from $\bar{M}$ by taking $\to_G := \bigcap_{a \in G} \to_a$. Then the function $F : \bar{S} \to S$, given by $F(\gamma, w) := w$, is a bounded morphism from $\bar{M}'$ to M such that $F(\lambda_0, w_0) = w_0$ (where λ_0 is the "empty" type $\lambda(x) = 0$ for all x).

To prove this proposition, it must be shown that the valuation of basic atomic sentences of $\bar{M}'$ is preserved in M and that M satisfies the back and forth conditions for each of the relations (Blackburn *et al.* 2001: 60-63).

Proof. (Preservation of valuation). For $(\gamma, w) \in \bar{S}$: $(\gamma, w) \in \|p\|_{\bar{M}'} = \|p\|_{\bar{M}} \Leftrightarrow w \in \|p\|_M$.

("Forth" condition for $\approx_a$). Suppose that $(\gamma, w) \approx_a (\gamma', w')$. Then there is some $\overset{\approx}{a}$-path from w to w' in $\bar{M}'$. By proposition 18, it follows that $w \approx_a w'$.

("Back" condition for $\approx_a$). Suppose that $w \approx_a w'$ and let $(\gamma, w) \in \bar{S}$. Then there exists some path p_0 from w_0 to w such that $\gamma = \gamma^{p_0}$. Now take the path $p_1 = (w, \approx_a w')$ from w to w'. Since $\text{last}(p_0) = w = \text{first}(p_1)$, the composition $p_0 \cdot p_1$ is a path from w_0 to w'. So $(\gamma^{p_0 \cdot p_1}, w') \in \bar{S}$ and p_1 is an $\overset{\approx}{a}$-path from w to w' such that $\gamma^{p_0 \cdot p_1} = \gamma^{p_0} + \gamma^{p_1} = \gamma + \gamma^{p_1}$. Hence it follows that $(\gamma, w) \approx_a (\gamma^{p_0 \cdot p_1}, w')$.

("Forth" condition for $\to_G$). Suppose that $(\gamma, w) \to_G (\gamma', w')$. By definition $\to_G = \bigcup_{a \in G} \to_a$, and this it follows immediately that $(\gamma, w) \to_a (\gamma', w')$ for all $a \in G$. It follows that for every $a \in G$ there is some $\vec{a}$-path p_a from w to w' such that $\gamma' = \gamma + \gamma^{p_a}$. This means that there is a family of paths $\{p_a : a \in G\}$ such that for all $a \in G$ the following holds: p_a is an $\vec{a}$-path from w to w' and $\gamma^{p_a} = \gamma' - \gamma$. By lemma 3, there thus exists some $\vec{G}$-path p from w to w'. By proposition 20, it follows that $w \to_G w'$ (in M).

("Back" condition for $\to_G$). Suppose that $w \to_G w'$ (in M), and let $(\gamma, w) \in \bar{S}$. Then there exists some path p_0 from w_0 to w such that $\gamma = \gamma^{p_0}$. Now take $q := (w, \to_G, w')$ and $p' := p_0 \cdot q$. Then p' is a path from w_0 to w', hence $(\gamma^{p'}, w') \in \bar{S}$. Moreover, q is a $\vec{G}$-path from w to w' such that $\gamma + \gamma^q = \gamma^{p_0} + \gamma^q = \gamma^{p_0 \cdot q} = \gamma^{p'}$. Hence: $(\gamma, w) \to_G (\gamma^{p'}, w;)$. Done! $\square$

This completes the proof for proposition 24.

Theorem 1 (Completeness for epistemic group models)**.** Let Φ be a consistent set of EGL sentences. Then there exists some epistemic group model M and some world w in it such that w satisfies Φ in M.

Proof. By proposition 16 (completeness for pseudo-models) there exists some pseudo-model M and world w_0 such that $w_0 \models_M \Phi$.

Let $\bar{M}$ be the epistemic group model generated by M according to the construction above (using w_0 as origin) and let $\bar{M}'$ be the associated pseudo-model (see proposition 12). Since $F : \bar{M}' \to M$ is a bounded morphism with $F(\lambda, w_0) = w_0$, it follows that $(\lambda, w_0) \models_{\bar{M}'} \Phi$ (since bounded morphisms preserve the truth of modal formulas). Since $\bar{M}'$ and $\bar{M}$ have the same truth conditions (see proposition 12), it follows that $(\lambda, w_0) \models_{\bar{M}} \Phi$. $\square$

Chapter 8

Conclusion

In this thesis, I have presented a formal approach to representing group knowledge. This was motivated by the observation that while philosophical concepts of group knowledge often entail a particular type of group as its possible subject, the standard formal definitions from epistemic logic typically do not include this type of property as they were designed with other applications in mind.

The two main assumptions that underlie the notion of *collective knowledge* are that knowledge is question-based and that group knowledge requires epistemic groups. To capture these assumptions, I introduced new Kripke models for knowledge, called *epistemic group models*, as well as modal operators for the key concepts. I then presented an axiomatic system for a logic with *collective knowledge* that is sound and complete with respect to the class of *epistemic group models*.

With *collective knowledge*, I do not claim to have found the ultimate philosophical definition of group knowledge. This was also not the aim of this thesis. What I hope to have provided, however, is a good first step towards a formal definition of group knowledge that represents philosophical concepts as discussed in the current literature, along with an accompanying axiomatic system that is sound and complete.

Formal models for group knowledge can help philosophers gain more insight into the ramifications of the philosophical concepts that they propose by clarifying the abstract properties of these concepts and their relationship to alternative proposals. For example, they can help clarify the effect of group structure on the possibility (and extent) of group knowledge, and point to properties of individual agents that are conducive to group knowledge. Such an approach has already proven fruitful with respect to theories of individual knowledge. It is time to extend this to the group level.

There are three issues related to the notion of *collective knowledge* that are not explicitly addressed in this thesis, which to my mind offer interesting topics for further investigation. The first issue concerns the idea that questions define agents' conceptual frameworks. This idea can be implemented by "subjective models", which are constructed from "testimonial issue models" (see appendix B). However, these latter models are not epistemic group models, but a modification thereof that is made by replacing condition $P2'$ with $P2$, and by requiring the issue-relation to be an equivalence relation. I have not yet looked into whether subjective models can be constructed from epistemic group models.

The second issue concerns group belief. In this thesis, I have followed standard epistemic logic practice and have taken knowledge as a primitive notion. In so doing, I have evaded addressing the problems surrounding the notion of group belief – from both a philosophical and computational perspective. On the philosophical side, it is controversial whether groups can hold beliefs. To many philosophers, the possibility of group knowledge turns on the possibility of group belief, and consequently philosophical discussions of group knowledge often start by addressing this latter possibility (Corlett, 1996; Fagan, 2012; Gilbert, 1989). On the computational side, group belief is a problematic concept, most notably because agents' beliefs may be mutually inconsistent, while it is commonly assumed that rational beliefs must be consistent. As such, an important problem is how to derive group beliefs from individual beliefs when these individual beliefs may be mutually inconsistent (Briggs *et al.* 2012; Pettit and List, 2011).

The third issue concerns knowledge and issue dynamics. In the spirit of Dynamic Epistemic Logic, dynamic modalities can be added to the syntax of EGL, such as $[!\varphi]\psi$ (saying that ψ holds after φ is publicly announced) and $[!G]\psi$ (saying that ψ holds after all agents G have publicly announced their knowledge). Further, in this thesis, each agent's issues were fixed, though clearly agents can refine their issues or revise them altogether. It remains for future work to investigate the behavior of group knowledge under issue-changing actions.

Acknowledgements

This thesis has benefited from the insight, support and friendship of several people, which have all contributed to making the process of writing this thesis enjoyable, instructive, and motivating–in general, to making this an invaluable experience altogether.

First and foremost, I would like to thank my supervisor Alexandru Baltag for his exceptional kindness, insight and support, for sharing his ideas with me on the topic of this thesis and many others (some related, some not at all), for including me in his research group, and for his impressive patience and willingness to take the time to teach me modal logic (this list can go on). Alexandru, I feel very fortunate and grateful for having had you as my supervisor. Your enthusiasm for logic and epistemology is contagious, and your kindness, insight and patience are inspiring.

Two other people that have been indispensable in the process of writing this thesis are my parents. Myles and Anja, though I was well aware of your support, astuteness and love, both of you have surpassed any expectations I thought I could reasonably have (and actually had). Thank you for everything.

I would also like to thank the members of the thesis committee for taking the time to read my thesis, and for their helpful feedback. Last but not least, I would like to thank Kees van Berkel, Jane Boddy, Kali Martin and Chris Tyson for proof reading my thesis, and for sharing their thoughts with me on the topic. Apart from this, your friendship, enthusiasm and support has been invaluable to me. I feel greatly indebted to all of you.

Appendices

Appendix A

Other treatments of questions

A number of philosophers have explicitly linked questions to knowledge. To my mind, the most important and informative example here is (Olsson and Westlund, 2006). Olsson and Westlund argue that rational beliefs are held in response to some question from the agent's research agenda. Another important example is Schaffer's so-called *contrastive* theory of knowledge (Schaffer, 2007). His theory is a version of the *Relevant Alternatives* approach to knowledge such that questions determine the set of relevant alternatives: s knows that p as the answer to Q. His proposal is further discussed in section A.3 below. Lastly, Hintikka is another prominent reference (Hintikka, 2007). He draws attention to the fact that questions exercise regulative control over inquiry, and are used to assess our information sources. As such, the act of questioning is fundamental to epistemology.

Still, though this thesis is built on ideas found in the above-mentioned lines of research, the latter are based on interpretations of "questions" that differ from the one presented her. The term "question" is perhaps slightly misleading, as it directs attention to the role of questions in information exchange, while its intended interpretation is epistemological rather than "Socratic" or dialogical. Yet, questions can have these latter roles only because they represent the (conceptual) distinctions that we can (or are willing to) make and subsequently presuppose in deliberation and inquiry. Without distinctions there is nothing to ask. In this section, my interpretation of the term "question" is further clarified by comparing it to alternative proposals from the contemporary literature.

A.1 Olsson and Westlund's proposal: epistemic agents have a research agenda

In their paper, Olsson and Westlund argue that "any adequate representation of epistemic states must also include the agent's research agenda, i.e., the list of questions that are open or closed at any given point in time" (Olsson and Westlund, 2006: 165). This list of (theoretical) questions corresponds to the agent's epistemic interests.[1] They propose a revision to the AGM theory of belief revision, namely, to add a so-called *research agenda* – that is, a list of questions – to the AGM representation of epistemic states.[2] This extension of the AGM model is motivated on epistemological grounds: namely, an agent's research agenda (which is not represented by the classic AGM model) plays a pertinent role in rational belief change.

[1] As such, they presuppose an interpretation of "questions" that differs from the one adopted in this thesis.

[2] For a statement of the AGM theory of belief revision, see Alchourrón *et al.* (1985).

In brief, in the AGM model, epistemic states are represented as sets of sentences (in some formal language) that correspond to beliefs. Belief revision is modeled as a process of adding (expanding) and/or deleting (contracting) sentences from the epistemic state, *viz.* of deleting a belief from your epistemic state and replacing it with a new belief. A problem of belief revision is to define a contraction-operation that specifies which beliefs should be deleted from the epistemic state in order to accommodate revision. The standard solution to this problem is to define an ordering over the set of beliefs, a so-called 'entrenchment ordering', and to then assume that less entrenched beliefs should be contracted first (Olsson and Westlund, 2006: 166). According to Olsson and Westlund, however, rational belief revision should not depend only on the agent's beliefs and the entrenchment ordering thereof, but should also be determined by her research agenda. As a motivating example, they consider a lottery ticket holder who decides to retract her belief that the lottery is fair, even though she is not forced to do so in order to avoid inconsistency of her beliefs (italics added by me):

> "It seems that, given that you care about the fairness of the lottery in the first place, your retracting your belief that it is fair *should be accompanied by the reopening of the question* as to its fairness. That is, you are, in some sense, obliged to take continued interest in the matter. (...) The point is that this sort of contraction is (normally) rational just in case it is made for the purpose of opening a neutral investigation into what the best answer is." (Ibid.:168)

Their point is that belief revision often (re)opens questions, questions that an epistemic agent is (rationally) committed to addressing. These questions are goal-directed theoretical questions that are (or were) on the agent's research agenda. Moreover, Olsson and Westlund assume that beliefs are held either "habitually or with no particular reason at all" or "in response to some question" (Ibid.: 172). The question-answering beliefs have a special status, as they are epistemically more valuable to the agent: they were acquired in response to an open question from her research agenda. Accordingly, they propose to represent epistemic states as triples, consisting of a set of beliefs, an entrenchment ordering of these beliefs and a research agenda.

As Olsson and Westlund themselves mention (in a footnote), their model suggests that *rational* epistemic change is motivated only by the agent's theoretical questions. Knowledge acquisition is undoubtedly an example of rational epistemic change. Olsson and Westlund focus on the role of an agent's questions in rational belief change, and do not explicitly connect their discussion to knowledge acquisition. Yet, extending their argument to knowledge is but a little step, for knowledge is the pinnacle of rational belief. This would amount to claiming that knowledge always answers a (previously open) question from the agent's research agenda. Note that the claim of this thesis is less strong, as it presupposes an interpretation of "questions" that is more general than that of research agenda. Propositional attitudes in general presuppose questions. They only makes sense in a world in which propositions can be true or false–that is, a world of uncertaintly.[3]

A.2 Inquisitive Semantics

Perhaps the most discussed epistemic role of questions from the contemporary literature is dialogical: questions are requests for information (Groenendijk and Stokhof, 1984; Hookway, 2008). In this context, a prime example is the *Inquisitive Semantics* framework as developed in e.g. (Ciardelli, 2009; Ciardelli *et al.* 2012; Groenendijk, 1999; Groenendijk and Stofhof, 1984). The main goal of this framework is to develop a notion of semantic-meaning that captures both the descriptive content and the inquisitive content of sentences. This is motivated by the observation that language is often used to exchange information, rather than merely provide it. In brief, classical logical semantics equates sentence-meaning with truth conditions. As such, it can only model the descriptive use of language and argumentation, namely, in terms of entailment between assertions (propositions). Gricean pragmatics is based on classical semantics, and so, though it aims to capture pragmatic inferences that arise in contexts of cooperative, rational conversation, it is based on a semantics that cannot represent essential aspects of such contexts (e.g. requests for information) – at least, this is the starting point of inquisitive semantics. Inquisitive semantics proposes to model meaning as requests for information or proposals of some sort to enhance the common ground, in order to capture information exchange as a process of raising and resolving questions. The inquisitive notion of meaning makes it possible to define several logical notions, e.g. their notion of an *issue*, and to formulate pragmatic principles based on them. In its initial formulation, the semantics of questions was given in terms of the *partition-semantics* developed in (Groenendijk and Stofhof, 1984). Thus, a question was associated with a partition of the state space, such that each partition cell corresponds to a possible answer. Since then, however, several systems of inquisitive semantics have been developed that propose to model questions differently, such as (Ciardelli *et al.* 2012). It should be noted that inquisitive semantics studies questions in as far as they represent requests for particular information. Questions represent the possible (cooperative) answers that are requested of the conversation partner – that is, the information that meets the request of the questioner. So, the interpretation of "question" presupposed by inquisitive semantics is that of a request for information, and as such, it is particularly relevant to the pragmatics of cooperative information exchange – which is not a topic of this thesis.

A.3 Schaffer's proposal: S knows that p as a true answer to Q

The final alternative proposal that I address is that of Schaffer. Schaffer (2007) proposes to analyze knowledge as a ternary relation: s knows that p as a true answer to Q. In more general terms: 's knows that p rather than q', where q is the disjunction of possible, but false, answers to Q. This is a version of his so-called *contrastive* theory of knowledge

[3] For instance, to wish that p presupposes the possibility of $\neg p$.

(Schaffer, 2005). Schaffer explains that,

> [t]he view that emerges links knowledge to *inquiry* and to *discrimination.* There is no such thing as inquiring into p, unless one specifies: *as opposed to what*? There is no such thing as discriminating that p, unless one adds: *from what*? And likewise I will argue that there is no such thing as knowing that p, unless one clarifies: *rather than what*? (Schaffer, 2005: 235)

As this quote exemplifies, Schaffer also draws attention to the fact that all knowledge is an answer. In fact, to the best of my knowledge, he is the only philosopher who explicitly states this. His intended interpretation of the notion of a question, however, is actually somewhat different from the notion assumed in this thesis, and it is driven by a different purpose. Schaffer defends a version of *Relevant Alternatives* theory of knowledge in which questions specify the set of relevant alternatives. So, if an agent knows that p relative to Q, then this means that she has eliminated all alternatives to p that are in Q. By arguing that knowledge is question-relative, his proposal provides a reponse to skepticism as follows: to know that p relative to Q, the agent must *only* be infallible regarding *p relative to the other possible answers to Q*. Questions thus expand our knowledge, as opposed to limiting it. As such, though I agree with his ideas, I cannot claim that they directly support my proposal. Yet, I do believe that they provide indirect support and therefore I have included a brief summary of his argument below.

Schaffer's argument is based on the observation that interrogative knowledge ascriptions cannot be reduced to propositional knowledge acriptions–that is, without loss of content. Let me briefly explain what this means. Apart from propositional knowledge, there are several other types of knowledge, including knowledge-*how, who, what, where* and *why* (etc.), collectively referred to as 'interrogative knowledge' and abbreviated to knowledge-*wh*. We routinely ascribe knowledge to agents not only when they know that a proposition is true, but also when e.g. they know *how* to grill beef, *what* they are wearing, *why* they are wearing it. Interrogative ascriptions explicitly embed questions, and so they explicitly involves a third component (i.e., a question). For instance, "Chris knows whether puffins forage at sea" embeds the question "Do puffins forage at sea?".

Although interrogative ascriptions are in fact more common than propositional knowledge ascriptions, philosophers have nonetheless focused almost exclusively on propositional knowledge (Schaffer, 2007: 383). Schaffer points out that the few philosophers who have discussed knowledge-*wh* have concurred that it is reducible to propositional knowledge, and thus to a binary relation "in which the question Q goes missing" (Ibid.: 384).

The fact that interrogative knowledge ascriptions embed questions does not imply that the knowledge involved is itself question-relative. Such an ascription can be interpreted as stating that (1) the subject has knowledge that happens to truthfully answer the question at issue, but it can also be interpreted as stating that (2) the subject has knowledge *only* relative to the question. According to the first interpretation, knowledge-*wh* reduces to (binary) propositional knowledge. The basic idea is that knowledge-*wh* ascriptions and propositional knowledge ascriptions do not differ in terms of the knowledge that they ascribe to the subject, and thus knowledge-*wh* is reducible to (binary) propositional knowledge. As such, it enough for knowledge-*wh* that p happens to be an answer to Q. Schaffer cites Hintikka (1975), Lewis (1982), Boer and Lycan (1986), Higginbotham (1996), and Stanley and Williamson (2001) as examples of such reductive analyses.[4] Common amongst these analyses is the following schema: s knows-*wh* p whenever s knows *that* p and p happens to be the answer to the indirect question Q of the *wh*-clause (Ibid.: 385).[6] So, while p must be a true answer to Q, it is not required that s knows that p as the result of trying to resolve Q, nor is it required that s knows that p is the answer to Q. There is thus no direct link between p and Q. According to the second interpretation, there is such a link: s knows that p *as a true answer to* Q – which is the view that Schaffer defends.

In order to defend this interpretation, Schaffer first argues that the reductive analyses are mistaken, i.e., knowledge-*wh* is not reducible to a binary relation. His argument is based on the observation that different questions can have the same (true) answer. If so, then knowledge ascriptions that embed these questions must be equivalent, as they express the same knowledge. So, for example, suppose that puffins eat herring. Then, the following knowledge claims will be equivalent: (C_1) "Chris knows whether puffins forage for herring or shrimp" and; (C_2) "Chris knows whether puffins forage for herring or bread."[7] However, argues Schaffer, (C_1) and (C_2) are intuitively not equivalent.[8] [9] Since reductive analyses imply that (C_1) and (C_2) are equivalent, they must therefore be false. And so interpretation (1) should be discarded in favor of (2). Schaffer then continues that, given that the term "knowledge" is unambiguous, it follows that all knowledge is an answer (and so a ternary relation): a knows that p iff s knows that p as the true answer to Q.[10] Schaffer adopts Stalnaker's (1999) notion of a context in order to explain that propositional knowledge includes an implicit question that can be retrieved from context.[11] It should be noted that according to Schaffer's proposal, knowledge is

[4]A compelling ground to assume the reductive view is the traditional conception of knowledge as some type of belief, such that 's believes/knows that p' expresses a binary relation between s and p. Additionally, the surface form of propositional knowledge makes it very natural to assume that knowledge is a binary relation (*Ibid.*: 385).[5] Third, the term "belief" (as opposed to "knowledge") cannot take an indirect question complement. An agent cannot belief *whether* puffins forage at sea, for instance. Nor can she believe *why* they forage at sea. Hookway [41] discusses this discrepancy between knowledge and belief ascriptions in order to question the assumption that the logical grammar of sentences about knowledge runs parallel to that of sentences about belief.

[6]For instance, Lewis fills in the schema as follows: "Holmes knows whether ... if and only if he knows the true one of the alternatives presented by the 'whether'-clause, whichever one that is" (1982, p. 45).

[7]They are both reducible to 'Chris knows that puffins forage for herring'

[8]As a second example, compare the following two knowledge claims:

- (C_3) "Jane knows whether it is raining or hailing in Vienna."
- (C_4) "Jane knows whether it is raining or sunny and dry in Vienna."

[9]Schaffer call this 'the problem of convergent knowledge'.

[10]This presupposes that all knowledge is either propositional knowledge or knowledge-*wh*.

[11]At an abstract level, such contexts correspond to the issues of this thesis, as they are simply partitions of the set of possible worlds.

always relative to a particular question: s knows that p *only* as the true answer to Q. Accordingly, if s knows that p relative to Q, then she may not know that p relative to some other question Q'.[12][13]

[12] Schaffer emphasizes that his proposal is a type of contextualism in which questions determine the set of relevant alternatives. Questions fix the context by determining the denotation of $\neg p$.

[13] In terms of *constrastivism* the knowledge relation is expressed as 's knows that p rather than q'. Here q is the disjunction of false answers to Q. So knowledge involves an explicit contrast between the known proposition p and its negation q. Instead of q one could write $\neg p$, where $\neg p$ denotes the disjunction of false answers to Q.

Appendix B

Subjective models

The idea that questions define agents' conceptual framework can be formalized by *subjective models*. These models provide additional support for the idea that agents represent their knowledge, as well as the knowledge of others, in terms of their own questions. I only introduce these models here in the appendix, however, as they are not (at least not currently) constructed from epistemic group models. As such, they are not applicable to group knowledge.

Epistemic models represent the agents' epistemic states from a third-person perspective. As such, they do not formalize an agent's subjective perspective. Intuitively, agents represent states that are conceptually indistinguishable to them as the same state. Thus, if $s \approx_a t$ holds, then states s and t correspond to the same state in a's own "subjective model". This idea can be implemented by *subjective models*, which are constructed from a modified version of the testimonial models of chapter 5:

Definition 33 (Testimonial issue model)**.** Given a set A of agents and a set Φ of atomic sentences, a *testimonial issue model* over (A, Φ) is a tuple $\mathbf{S} = (S, \rightarrow_{a(a\in A)}, \approx_{a(a\in A)}, \|\bullet\|)$, such that $(S, \rightarrow_{a(a\in A)}, \|\bullet\|)$ is an epistemic model over (A, Φ) and $\approx_a$ is a map associating each agent $a \in A$ with some equivalence relation $\approx_a \subseteq S \times S$, satisfying the following three conditions: $(*)$ $\approx_a \subseteq \rightarrow_a$; $(**)$ $\approx_a \rightarrow_b \subseteq \rightarrow_b \approx_a$; $(***)$ $\approx_a \approx_b \subseteq \approx_b \approx_a$.

The difference between testimonial issue models and testimonial models is that $\approx_a$ is assumed to be an equivalence relation, and is therefore best understood as a question-relation, rather than an epistemic-relation. Subjective models are constructed from testimonial issue models as follows: given a testimonial model $\mathbf{S}$ and an agent $a \in A$, her subjective group model can be represented as a *quotient* of the model $\mathbf{S}$ given the issue-relation $\approx_a$:

Definition 34 (Subjective model)**.** Given a testimonial issue model $\mathbf{S} = (S, \rightarrow_{a\in A}, \approx_{a\in A}, \| \bullet \|)$ over (A, Φ) and an agent $a \in A$, put $\Phi_a := \{p \in \Phi : \forall s, t \in S(s \in \|p\| \wedge s \approx_a t \Rightarrow t \in \|p\|\}$ for the set of all agent-relevant atomic sentences. Then *a's subjective model* is a testimonial model $\mathbf{S_a} = (S_a, \rightarrow_a, \approx_a, \| \bullet \|)$ over (A, Φ_a), given by:

1. $S_a := \{s_a : s \in S\}$ consists of equivalence classes $s_a := \{s' \in S : s \approx_a s'\}$.

2. $s_a \rightarrow_b t_a$ iff $s \rightarrow_b \approx_a t$ (for all $b \in A$).

3. $s_a \approx_b t_a$ iff $s \approx_b \approx_a t$ (for all $b \in A$).

4. $\|p\| = \{s_a : s \in \|p\|\}$, for $p \in \Phi_a$.

Observation 2. For $b = a$ clauses 2 and 3 can be simplified as follows*:*

- $s_a \rightarrow_a t_a$ iff $s \rightarrow_a t$.
- $s_a \approx_a t_a$ iff $s \approx_a t$ (i.e. iff $s_a = t_a$).

  Proof. Suppose that $(b = a)$. The following two equalities must be shown: (1) $\rightarrow_a; \approx_a = \rightarrow_a$ and (2) $\approx_a = \approx_a; \approx_a$.
  (1) ($\Rightarrow$): Suppose $\rightarrow_a; \approx_a$. By $P1$ (ie, $\approx_a \subseteq \rightarrow_a$) it follows that $\rightarrow_a; \approx_a \subseteq \rightarrow_a \rightarrow_a$, and by transitivity of $\rightarrow_a$, this means that: $\rightarrow_a; \approx_a \subseteq \rightarrow_a \rightarrow_a = \rightarrow_a$.
  ($\Leftarrow$): Suppose $\rightarrow_a$. By veracity of $\approx_a$ (i.e., $id \subseteq \approx_a$), it follows by monotonicity of relational composition that $\rightarrow_a; id \subseteq \rightarrow_a; \approx_a$ and hence $\rightarrow_a \subseteq \rightarrow_a; \approx_a$.
  (2) Trivial. □

The set S_a of states in $\mathbf{S_a}$ is a representation of the way in which a conceptualizes the possible states. It consists of states s_a that are agent-relevant. Upon closer inspection, these states are really equivalence classes of states generated by $\approx_a$. Yet, since the states in these equivalence classes are indistinguishable from the perspective of a, she represents each equivalence class as a single state. According to the second clause, a represents all epistemically possible states that are agent-relevant (and only those). The third clause requires that if a cannot conceptually distinguish between states s and t in the (third-person) testimonial model, then she cannot distinguish between them in her subjective model either: any distinctions that go beyond her issue-relation are lost. Conditions ($P2$), and ($P3$) ensure the coherence of the subjective model construction.

Proposition 25. If $\mathbf{S}$ is an epistemic group model, then $\mathbf{S_a}$ is well-defined (in particular, the definitions of $\rightarrow$ and $\approx$ between equivalence classes do not depend on the choice of representative for these classes), and moreover $\mathbf{S_a}$ is itself an epistemic group model.

Proof. Three things need to be shown: (1) $\rightarrow_a$ is well-defined on equivalence classes s_a, (2) $\approx_a$ is well-defined on equivalence classes s_a, and (3) for $p \in \Phi_a$, $\|p\|$ is well-defined on equivalence classes s_a.
(1). Suppose that $s_a = s'_a$ i.e,. $s \approx_a s'$ and $t_a = t'_a$ i.e,. $t \approx_a t'$. To show: $s \rightarrow_b \approx_a t$ iff $s' \rightarrow_b; \approx_a t'$.

($\Rightarrow$). Suppose that $s \rightarrow_b; \approx_a t$. By definition, $s \rightarrow_b; \approx_a t$ iff $\exists w' : s \rightarrow_b w \approx_a t$. Putting this together with the other assumption, this means that $\exists w' : s' \approx_a s \rightarrow_b w \approx_a t \approx_a t'$. By $P2$ it follows that $s' \rightarrow_b; \approx_a w \approx_a t \approx_a t'$. By transitivity of $\approx_a$ it follows that $s' \rightarrow_b; \approx_a w \approx_a t'$, and also that $s' \rightarrow_b; \approx_a t'$. ($\Leftarrow$). This follows by symmetry of $s \rightarrow_b \approx_a t$ iff $s' \rightarrow_b; \approx_a t'$.
(2). Suppose that $s_a = s'_a$ i.e,. $s \approx_a s'$ and $t_a = t'_a$ i.e,. $t \approx_a t'$. To show: $s \approx_b \approx_a t$ iff $s' \approx_b; \approx_a t'$.

($\Rightarrow$). Suppose that $s \approx_b; \approx_a t$. By definition, $s \approx_b; \approx_a t$ iff $\exists w' : s \approx_b w \approx_a t$. Putting this together with the ohter assumption, this means that $\exists w' : s' \approx_a s \approx_b w \approx_a t \approx_a t'$. By $P2$ it follows that $s' \approx_b; \approx_a w \approx_a t \approx_a t'$. Then, by transitivity of $\approx_a$ it follows that $s' \approx_b; \approx_a t'$. ($\Leftarrow$). This follows by symmetry of $s \approx_b \approx_a t$ iff $s' \approx_b; \approx_a t'$.

(3). For $p \in \Phi_a$, $\|p\|$ is well-defined on equivalence classes s_a, since if $s_a = s'_a$ (and thus $s \approx_a s'$) then: $s \models p \leftrightarrow s' \models p$ (for $p \in \Phi$). $\square$

The introspective properties of the knowledge relation $\to_a$ in the original model are preserved in the subjective model. This is captured by the following two propositions:

Proposition 26. If $\mathbf{S}$ is positively introspective then $\mathbf{S_a}$ is positively introspective.

Proof. Suppose $\to_c$ is positively introspective and thus $s \to_c ; \to_c t \subseteq s \to_c t$ for any arbitrary $c \in A$. Suppose $s_a \to_b ; \to_b t_a$. This means that $\exists s'_a$: $s_a \to_b s'_a \to_b t_a$. By definition, $s_a \to_b s'_a$ iff $s \to_b ; \approx_a s'$, and similarly, $s'_a \to_a t_a$ iff $s' \to_b ; \approx_a t$. Hence: $s \to_b ; \approx_a ; \to_b ; \approx_a t$. By condition $P2'$, it follows that $s \to_b ; \to_b ; \approx_a \approx_a t$. By transitivity of $\to_b$ and $\approx_a$ (in the old model), it follows that $s \to_b ; \approx_a t$, and thus: $s_a \to_b t_a$. $\square$

Proposition 27. If $\mathbf{S}$ is fully introspective then $\mathbf{S_a}$ is fully introspective.

Proof. $\mathbf{S}$ is fully introspective means that $\to_a$ on $\mathbf{S}$ is an equivalence relation – hence reflexive, transitive ans symmetric. Suppose $s_a \to_b t_a$. By definition, $s_a \to_b t_a$ iff $s \to_b ; \approx_a t$. By symmetry of $(\to_b ; \approx_a)$ in the old model, it follows that $s \, (\to_b ; \approx_a)^{-1} \, t$, where $\to_a^{-1} = \{(s,t) \in S \times S : t \to_a s\}$ is the inverse of $\to_a$. This means that $t \to_b ; \approx_a s$. And thus, $t_a \to_a s$, which means that $s_a \to_a^{-1} t_a$. $\square$

Bibliography

[1] Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50(2), 510-530.

[2] Aloni, M., Egré, P., and De Jager, T. (2013). Knowing whether A or B. *Synthese, 190*(14), 2595-2621.

[3] Baltag, A., Boddy, R., and Smets, S. (2014). An Interrogative approach to Group Knowledge. Unpublished manuscript.

[4] Baltag, A., van Ditmarsch, H., and Moss, L. (2008). Epistemic logic and information update. *Handbook of the Philosophy of Information, Elsevier Science Publishers, Amsterdam*, 361-456.

[5] Barwise, J. (1988). Three views of common knowledge. In *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge* (pp. 365-379). Morgan Kaufmann Publishers Inc..

[6] Blackburn, P., de Rijke, M., and Venema, Y. (2001). *Modal Logic.* Cambridge: Cambridge University Press.

[7] Bovens, L., and Hartmann, S. (2003). *Bayesian epistemology.* Oxford: Oxford Univeristy Press.

[8] Bratman, M. E. (1993). *Shared intention.* Ethics, 97-113.

[9] Briggs, R., Cariani, F., Easwaran, K., and Fitelson, B. (2012). Individual coherence and group coherence.

[10] Ciardelli, I. (2009). Inquisitive semantics and intermediate logics.

[11] Ciardelli, I, Groenendijk, J., and Roelofsen, F. (2012) Inquisitive semantics. *NASSLLI lecture notes.*

[12] Coady, C. A. (1992). *Testimony: A philosophical study.* Oxford University Press.

[13] Corlett, A. J. (1996). *Analyzing social knowledge.* Rowman and Littlefield Publishers.

[14] Corlett, J. A. (2007). Analyzing Social Knowledge 1. *Social Epistemology, 21*(3), 231-247.

[15] De Ridder, J. (2014). Epistemic dependence and collective scientific knowledge. *Synthese, 191*(1), 37-53.

[16] DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review, 104*(0), 1-52.

[17] Dretske, F. I. (1970). Epistemic operators. *The journal of philosophy*, 1007-1023.

[18] Fagan, M. (2012). Collective Scientific Knowledge. *Philosophy Compass, 7*(12), 821-831.

[19] Fagin, R., Halpern, J. Y., and Vardi, M. Y. (1992). What can machines know?: On the properties of knowledge in distributed systems. *Journal of the ACM (JACM)*, 39(2), 328-376.

[20] Halpern, J. Y., Moses, Y., and Vardi, M. Y. (1995). *Reasoning about knowledge* (Vol. 4). Cambridge: MIT press.

[21] Fricker, E. (2006). SecondHand Knowledge. Philosophy and Phenomenological Research, 73(3), 592-618.

[22] Gettier, E. L. (1963). Is justified true belief knowledge?. *Analysis*: 121-123.

[23] Gilbert, M. (1987). Modelling collective belief. *Synthese*, 73(1), 185-204.

[24] Gilbert, M. (1992). *On social facts.* Princeton University Press.

[25] Goldman, A. I. (2004). Group knowledge versus group rationality: Two approaches to social epistemology. *Episteme*, 1(01), 11-22.

[26] Goldman, A., and Whitcomb, D. (Eds.). (2011). Social epistemology: essential readings. Oxford University Press.

[27] Groenendijk, J. (1999). The logic of interrogation: Classical version. In *Proceedings of SALT*, vol. 9, pp. 109-126.

[28] Groenendijk, J. (2009). *Inquisitive semantics: Two possibilities for disjunction* (pp. 80-94). Springer Berlin Heidelberg.

[29] Groenendijk, J., and Stokhof, M. (1984). Studies on the Semantics of Questions and the Pragmatics of Answers.

[30] Groenendijk, J. A., and Roelofsen, F. (2009). *Inquisitive semantics and pragmatics.*

[31] Hakli, R. (2007). On the possibility of group knowledge without belief. *Social Epistemology*, 21(3), 249-266.

[32] Halpern, J. Y., Moses, Y., and Vardi, M. Y. (1995). *Reasoning about knowledge (Vol. 4).* Cambridge: MIT press.

[33] Hardwig, J. (1985). Epistemic dependence. *The Journal of philosophy*, 335-349.

[34] Hardwig, J. (1991). The role of trust in knowledge.

[35] Hendricks, V. F. (2006). *Mainstream and formal epistemology*: Cambridge University Press.

[36] Higginbotham, J. (1996). The semantics of questions, in Shalom Lappin (ed.) *The Handbook of Contemporary Semantic Theory* (Oxford): 361 - 381.

[37] Hintikka, J. (1992). *Knowledge and Belief: An introduction to the Logic of the Two Notions*. Cornell University Press, Ithaca, 1962

[38] Hintikka, J. (1981). On the logic of an interrogative model of scientific inquiry. *Synthese, 47*(1), 69-83.

[39] Hintikka, J. (2007). *Socratic epistemology: Explorations of knowledge-seeking by questioning.*: Cambridge University Press.

[40] Holliday, W. H. (2013). Epistemic closure and epistemic logic I: Relevant alternatives and subjunctivism. *Journal of Philosophical Logic*. Forthcoming.

[41] Hookway, C. (2008). Questions, epistemology, and inquiries. *Grazer Philosophische Studien* 77, no. 1: 1-21.

[42] Lackey, J. (2008). *Learning from words: Testimony as a source of knowledge*. Oxford University Press.

[43] Lackey, J., and Sosa, E. (Eds.). (2006). *The epistemology of testimony*. Oxford University Press.

[44] Levesque, H. J. (1984, August). A logic of implicit and explicit belief. In *AAAI*: 198-202.

[45] Lewis, D. (1982). Whether Report.

[46] Lewis, D. (1996). Elusive knowledge. *Australasian journal of Philosophy*, 74(4), 549-567.

[47] Lewis, D. (2008). *Convention: A philosophical study*. Wiley. com.

[48] List, C. (2011). Group knowledge and group rationality: a judgment aggregation perspective. *Social Epistemology: Essential Readings*, 221.

[49] List, C., Pettit, P. (2011) *Group agency: The possibility, design, and status of corporate agents*. Oxford: Oxford University Press.

[50] Moore, G. E. (1959). Certainty. In Goodman, M. F. and Snyder, R. A. (eds). (1993). *Contemporary readings in Epistemology*. Englewoods Cliffs, NJ: Prentice Hall.

[51] Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.

[52] Olsson, E. J., and Westlund, D. (2006). On the role of the research agenda in epistemic change. *Erkenntnis, 65*(2), 165-183.

[53] Quinton, A. (1976). Social objects. In *Proceedings of the Aristotelian society* 76: 1-27.

[54] Pettit, P. (2010). Groups with minds of their own. *Social Epistemology: Essential Readings.*

[55] Rolin, K. (2008). Science as collective knowledge. *Cognitive Systems Research*, 9(1), 115-124.

[56] Schaffer, J. (2005). 9. Contrastive Knowledge. *Oxford studies in epistemology*, 235.

[57] Schaffer, J. (2007). Knowing the answer. *Journal*, etc.

[58] Sosa, E. (2004). Relevant alternatives, contextualism included. *Philosophical studies, 119*(1), 35-65.

[59] Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language and Information*, 7(1), 3-19.

bibitemStanley01 Stanley, J., and Williamson, T. (2001). Knowing how. *The Journal of Philosophy*, 411-444.

[60] Stalnaker, R. (2006). On logics of knowledge and belief. *Philosophical studies*, 128(1), 169-199.

[61] Stalnaker, R. (1999). Context and content: Essays on intentionality in speech and thought.

[62] Van Benthem, J., and Minică, Ş. (2009). Toward a dynamic logic of questions. In *Logic, Rationality, and Interaction* (pp. 27-41). Springer Berlin Heidelberg.

[63] Van Der Hoek, W., Van Linder, B., and Meyer, J. J. (1999). Group knowledge is not always distributed (neither is it always implicit). *Mathematical social sciences, 38*(2), 215-240.

[64] Wray, K. B. (2001). Collective belief and acceptance. *Synthese, 129*(3), 319-333.

[65] Wray, K. B. (2007). Who has scientific knowledge?. *Social Epistemology, 21*(3), 337-347.