From Pritchard & Bernecker 2011

52

LOGIC AND FORMAL SEMANTICS FOR EPISTEMOLOGY

John Symons

Epistemic expressions such as 'knows that' or 'believes that' have systematic properties that are amenable to formal study. Most obviously, statements containing epistemic expressions sometimes involve logical constants which behave in the usual way. So, for example, if you know p and q then you know q. The conceptual features of statements concerning knowledge and belief become more interesting when one begins to examine the characteristics of general principles governing the use of epistemic concepts. The behavior and interaction of these general principles has been the focus of epistemic logic. For example, as G.E. Moore pointed out, there seems to be something wrong with claiming

(1) "p and I do not believe p."

Assertions of this kind are self-defeating because of the conceptual features of knowledge or belief and not because of the syntactical features of the sentence or the character of the logical constants that are involved. As Moore noted "'I went to the pictures last Tuesday, but I don't believe that I did' is a perfectly absurd thing to say" (1952: 543). The *perfect absurdity* here is due to a violation of a principle governing epistemic concepts.

Notice that (1) is often a correct description of the state of affairs in question. For instance, since I recognize that I am fallible, I am committed to the possibility that there are cases where it is true that p and I do not believe p. Furthermore, I can assert, without paradox or contradiction, that there is some proposition p such that p and I do not believe p.

The paradox arises from the peculiarity of the agent in question attesting to particular instances of (1), where the variable p is replaced by an assertion concerning some state of affairs. Specifically, it is paradoxical insofar as it is, what John Austin called, an illocutionary act (1975: 133). While I recognize that there might be cases where replacing p with some description of some state of affairs is true, I cannot sincerely attest to both parts of the conjunction contained in (1) at a particular moment for any specific instance of (1). In this sense, Moore's paradox sheds light on the properties of epistemic agents and the concept of belief. Reflecting on the *perfect absurdity* of Moore's examples, shows us that an agent's belief and its agency are related. However, we are not restricted to relying on epistemic intuitions in our consideration of Moore's paradox. Exploration of the principles or norms governing epistemic notions can take place in an axiomatic fashion. Jaakko Hintikka, for example, provided a proof of the contradictory nature of the paradoxical form of the Moore statements, the case where the statement asserts that an agent believes *p* and *not-p* (1962: 67). I can take as a rule for instance that it is prohibited or confused to say:

(2) "I know *p* but it is not the case that *p*."

If (2) is prohibited, it is due in part to the illocutionary considerations which applied in (1) but unlike assertions of belief, (2) is prohibited by virtue of another general principle, namely the veracity of knowledge.

(3) If one knows *p* then it is true that *p*.

If we accept that knowledge implies veracity then we can consider what the implications of taking it as an axiom might be and whether it is consistent with other general epistemic principles we might hold. As discussed below, considerations of this kind have been given an elegant formal framework by epistemic logicians. Wolfgang Lenzen (1978) provided an excellent overview of arguments from the 1960s and 1970s concerning the appropriate axioms for knowledge.

Observations of the kind emphasized by Moore, concerning the behavior of the term "knows that" served as the starting points for the development of modern epistemic logic. G.H. von Wright was the first to sketch an axiomatic treatment of the behavior of epistemic concepts (1951: 29–35). However, modern epistemic logic began in earnest once Hintikka provided a semantic interpretation of epistemic and doxastic notions in the early 1960s.

Hintikka began by supplementing the language of propositional logic with two unary epistemic operators K_a and B_a such that $K_a p$ reads 'Agent *a* knows *p*' and $B_a p$ reads 'Agent *a* believes *p*' for some proposition *p*. In this way, candidate epistemic or doxastic axioms can be presented in formal terms. So, for instance, we have already seen that one intuitive axiom which we are likely to accept into our epistemic logic is:

(4) $K_c A \rightarrow A$

This is known as axiom T which we saw above as (3). With our modest addition to first-order logic in hand, we can begin to catalog other plausible epistemic axioms.

A standard list of the axioms (following Lemmon (1977), Bull and Segerberg (1984)) that are relevant for epistemic logic run as shown in Table 52.1:

Table 52.1 Axioms of Epistemic Logic

K	$K_c(A \rightarrow A') \rightarrow (K_c A \rightarrow K_c A')$
D	$K_c A \rightarrow \neg K_c \neg A$
T	$K_c A \rightarrow A$
4	$K_c A \rightarrow K_c K_c A$
5	$\neg K_c A \rightarrow K_c \neg K_c A$

.2 $\neg K_c \neg K_c$
.3 $K_c(K_c A$ -
.4 $A \rightarrow (\neg$

We can consider the
introduction of addit
epistemic notions of
to begin thinking ab
ing the familiar inter
along the lines discu

$K_c A$: *In all possib*
$B_c A$: *In all possib*

The basic assumpt
edge and belief, invo
patible with the attit
The central idea in
is a relation that is de
that some world *w* is
w'. Specifically, the r
the set of possible wo
the most basic step i
example, whether or
reflexive, or some co
one thinks about th
epistemic context, th
tives) depends on its
via the specification
express the idea that
or accessible from th
between *w* and *w'*. T
w''. The world *w'* is
depending on wheth
semantic interpretat
agent *c* considers pos
A possible world
then consists of a *fra*
set of possible worlds
W. A *model* $\mathcal{M}$ for a
assigning sets of wor
sets of possible world
be the set of atomic
powerset operation.
The model $\mathcal{M}$ =
Kripke-semantics (K
in a world *w* in $\mathcal{M}$ (
i.e., $\mathcal{M},w \models$ **a** iff *w* ∈

.2 $\neg K_c \neg K_c A \rightarrow K_c \neg K_c \neg A$

.3 $K_c(K_c A \rightarrow K_c A') \vee K_c(K_c A' \rightarrow K_c A)$

.4 $A \rightarrow (\neg K_c \neg K_c A \rightarrow K_c A)$

We can consider the philosophical merits of each axiom to a certain extent without the introduction of additional formalism. However, Hintikka's approach to the semantics of epistemic notions offers an important supplement to our intuitive reflections. In order to begin thinking about the relative merits of these axioms one can begin by considering the familiar interpretation of the K and B operators using possible world semantics along the lines discussed above:

$K_c A$: *In all possible worlds compatible with what c knows, it is the case that A*
$B_c A$: *In all possible worlds compatible with what c believes, it is the case that A*

The basic assumption is that any ascription of propositional attitudes such as knowledge and belief, involves dividing the set of possible worlds in two: Those worlds compatible with the attitude in question and those that are incompatible with it.

The central idea in possible worlds semantics is the notion of accessibility. Accessibility is a relation that is defined on the set of possible worlds. In standard modal logic we say that some world w is accessible from some world w' just in case w is *possible* relative to w'. Specifically, the relation can be characterized as a subset of the Cartesian product of the set of possible worlds. As described below, determining the accessibility relation is the most basic step in determining the properties of our semantical framework. So, for example, whether one assumes that the accessibility relation is symmetric, transitive, reflexive, or some combination of the three, will make a significant difference in how one thinks about the modal or epistemic properties of the system in question. In the epistemic context, the set of worlds accessible to an agent (its set of epistemic alternatives) depends on its informational resources at an instant. This dependency is captured via the specification of the accessibility relation, R, on the set of possible worlds. To express the idea that for agent c, the world w' is compatible with his information state, or accessible from the possible world w which c is currently in, it is required that R holds between w and w'. This relation is written Rww' and reads "world w' is accessible from w". The world w' is said to be an *epistemic* or *doxastic alternative* to world w for agent c, depending on whether knowledge or belief is under consideration. We can give this a semantic interpretation, by saying that if a proposition A is true in all worlds which agent c considers possible then c knows A.

A possible world semantics for a propositional epistemic logic with a single agent c then consists of a *frame* $\mathcal{F}$ which in turn is a pair $<W,R_c>$ such that W is a non-empty set of possible worlds and R_c is a binary accessibility relation (relative to agent c) over W. A *model* $\mathcal{M}$ for an epistemic system consists of a frame and a denotation function φ assigning sets of worlds to atomic propositional formulas. Propositions are taken to be sets of possible worlds; namely the set of possible worlds in which they are true. Let *atom* be the set of atomic propositional formulae, then $\varphi : atom \mapsto P(W)$, where P denotes the powerset operation.

The model $\mathcal{M} = <W,R_c,\varphi>$ is called a Kripke-model and the resulting semantics Kripke-semantics (Kripke 1963): An atomic propositional formula, **a**, is said to be true in a world w in $\mathcal{M}$ (written $\mathcal{M},w \models \mathbf{a}$) iff w is in the set of possible worlds assigned to **a**, i.e., $\mathcal{M},w \models \mathbf{a}$ iff $w \in \varphi(\mathbf{a})$ for all $\mathbf{a} \in atom$. The formula $K_c A$ is true in a world w (i.e.,

$\mathcal{M},w \models K_cA)$ iff $\forall w' \in W$, if $R_c ww'$, then $\mathcal{M},w' \models A$. The semantics for the Boolean connectives follow the usual recursive recipe. Similar semantics can be formulated for the belief operator. Since a belief is not necessarily true but, rather, probably true, possibly true, or likely to be true, we must modify our approach to the semantics of belief appropriately. For instance, belief can be modeled by assigning a sufficiently high degree of probability to the proposition in question and determining the doxastic alternatives accordingly. The truth-conditions for the doxastic operator are defined in a way similar to that of the knowledge operator and the model can also be expanded to accommodate the two operators simultaneously.

A modal formula is said to be *valid* in a frame if, and only if, the formula is true for all possible assignments in all worlds in the frame.

An important feature of possible world semantics is that the epistemic axioms listed above, correspond to algebraic properties of the frame in the following sense: A modal axiom is valid in a frame if, and only if, the accessibility relation satisfies some algebraic condition (see Hendricks and Symons 2006). For example, the axiom expressing the veridicality property that if a proposition is known by c, then A is true,

(5) $K_cA \rightarrow A$,

is valid in all frames in which the accessibility relation is reflexive in the sense that $\forall w \in W$: Rww. Given reflexive accessibility, every possible world is accessible from itself. Similarly if the accessibility relation satisfies the condition that

(6) $\forall w, w', w' \in W: Rww' \wedge Rw'w' \rightarrow Rww'$

which is also known as transitivity, then the axiom (7) is valid.

(7) $K_cA \rightarrow K_cK_cA$

(7) is called axiom 4 and is also known as the axiom of self-awareness, positive introspection, or the KK-thesis. In this case, the axiom captures the idea that if the agent knows p then it has knowledge of its knowledge that p. Other axioms require yet other relational properties to be met in order to be valid in all frames: If the accessibility relation is reflexive, symmetric and transitive, then

(8) $\neg K_cA \rightarrow K_c\neg K_cA$

is valid. (8) is called axiom 5, also better known as the axiom of wisdom. This is the much stronger thesis that an agent has knowledge of its own ignorance: If a does not know p, it knows that it doesn't know p. The axiom is also known as the axiom of negative introspection.

One contentious axiom which is valid in all possible frames,

(9) $K_c(A \rightarrow B) \rightarrow (K_cA \rightarrow K_cB)$,

is the closure condition for knowledge, also known as axiom K, or the axiom of deductive cogency: If the agent a knows $p \rightarrow q$, then if a knows p, a also knows q. As discussed below, this axiom leads to the most difficult philosophical problem for epistemic logicians, namely the apparent commitment to logical omniscience. It seems that if one accepts this axiom, then an epistemic agent must know everything that follows logically from its knowledge.

Other axioms of epistemic import require yet other relational properties to be met in order to be valid in all frames. When combined in various ways, these axioms make up epistemic modal systems of varying strength. Their strengths vary according to the modal formulas valid in the respective systems and given the algebraic properties assumed for the accessibility relation.

Returning to the axioms listed above, we can begin to see how we might compare their relative strengths. The weakest system of epistemic interest is usually considered to be system **T**. The reader should take care to distinguish the epistemic operator *K*, the modal axiom K and the system of axioms **K** in what follows. Similarly, we distinguish the axiom T from the system **T**. T is a system of modal logic which is characterized by reflexive frame with the axioms T and K as valid axioms.

Additional modal strength can be obtained by extending **T** with other axioms drawn from the above pool, altering the frame semantics to validate the additional axioms. By way of example, while

(10) $K_c A \to A$

is valid in system **T**,

(11) $K_c A \to A$, $K_c A \to K_c K_c A$ and $\neg K_c A \to K_c \neg K_c A$

are all valid in **S5** but not in **T**.

System **T** has a reflexive accessibility relation, **S5** has reflexive, transitive and symmetrical accessibility relations. The arrows in Table 52.2 below indicate that the system to which the arrow is pointing is included in the system from which the arrow originates and hence reflect relative strength. Then **S5** is the strongest and **S4** the weakest of the ones listed.

Table 52.2 Relative Strength of Epistemic Systems Between S4 and S5

		Epistemic Systems	
KT4	=	S4	
KT4 + .2	=	S4.2	↑
KT4 + .3	=	S4.3	↑
KT4 + .4	=	S4.4	↑
KT5	=	S5	↑

One of the important tasks of epistemic logic is to catalog all sound and complete systems of such logics in order to allow us to pick the most 'appropriate' ones. The logics range from S4 over the intermediate systems S4.2–S4.4 to S5. By way of example, Hintikka settled for S4 (1962), Kutschera argued for S4.4 (1976), Lenzen suggested S4.2 (1978), van der Hoek has proposed to strengthen knowledge according to system S4.3 (Meyer and van der Hoek 1995). Van Ditmarsch, van der Hoek and Kooi (2007) together with Fagin, Halpern, Moses and Vardi (Fagin et al. 1995) and others assume knowledge to be S5 valid.

In the doxastic context, we can also catalog the completeness properties of the alternative systems in a similar fashion. Of course in doxastic logic we drop axiom T, which is usually replaced by D. This avoids committing doxastic logic to the truth of beliefs while retaining the condition that beliefs be consistent. Replacing T with D generates systems like KD4–KD45. This approach permits the combination of epistemic and doxastic systems and for studying the interplay between knowledge and belief (see Voorbraak 1993). There are some important philosophical concerns with such combined doxastic and epistemic systems. Lenzen (1978) and Stalnaker (1996) point out that such combined systems risk conflating knowledge and belief.

How does semantic formalization relate to epistemology? By way of example, it is worth returning briefly to our discussion of Moore's problem to see what kind of light Hintikka's formalization shed on that case. In *Knowledge and Belief*, he was able to prove that statements of the sort "*p* and I do not believe *p*" are *perfect absurdities* not because they run afoul of some kind of epistemic intuition, but because, when properly analyzed, they generate a contradiction. More importantly, the analysis allows us to recognize which epistemic commitments are involved in generating the contradiction. These commitments are formulated as rules for epistemic alternatives in model systems. So, for example, the proof of the absurdity of "*p* and I do not believe *p*" (Hintikka 1962: 68) relies on the conditions governing the semantics of sentences concerning belief. The difference between the kind of reasoning we find in Moore and Hintikka with respect to "*p* and I do not believe *p*" boils down to difference with respect to the degree of explicitness and control that the philosophers aspire to in their arguments. For Hintikka, unlike Moore, the point is to achieve the same level of explicitness in epistemology as is found in logic:

> The word "logic" which occurs in the subtitle of this work is to be taken seriously. My first aim is to formulate and to defend explicit criteria of consistency for certain sets of statements—criteria which, it is hoped, will be comparable with the criteria of consistency studied in the established branches of logic.
> (1962: 3)

Logical Omniscience and Idealized Epistemic Agents

Epistemic logic inevitably traffics in idealizations. As discussed below, the problem of logical omniscience (a product of accepting the axiom of deductive cogency or axiom K and standard possible world semantics) encouraged theorists to craft formal systems which more adequately reflected the actual properties of epistemic agents. Since real epistemic agents modify their beliefs and engage in inquiry, there was some philosophical interest in attempting to formally capture the dynamical features of inquiry. Developments since *Knowledge and Belief*, principally those since Kutchera (1976) and Lenzen (1978), attempted to integrate broader insights from modal logic with epistemic logic and have made it possible to formally model some prominent features of the dynamical nature of epistemic agency. Gärdenfors' (1988) account of belief revision was particularly important in setting the stage for a slew of dynamical models of knowledge.

Logical omniscience is related to closure properties. Axiom K can, under certain circumstances, be generalized to a closure property for an agent's knowledge which is implausibly strong: Whenever an agent *c* knows all of the formulas in a set Γ and *A* follows logically from Γ, then *c* also knows A. In particular, *c* knows all theorems (letting Γ = ø), and he knows all logical consequences of any formula which he knows (letting Γ consist of a single formula).

In response to the threat of logical omniscience, some epistemologists raised the question of whether the very idea of a logic of knowledge makes any epistemological sense. For instance, Hocutt challenged the applicability of logic to any realistic account of knowledge (1972). Because there is no guarantee that a knower will recognize that it is committed to some proposition that is logically equivalent to some proposition to which it readily assents, the very idea of an epistemic logic is on slippery ground.

Some of the first proposals for solving the problem of logical omniscience introduce semantical entities which explain why the agent appears to be, but in fact is not really logically omniscient. These entities were called 'impossible possible worlds' by Hintikka (1975). The basic idea is that an agent might mistakenly count among the worlds consistent with his or her knowledge, some worlds containing logical contradictions. The mistake is simply a product of limited resources; the agent might not be in a position to detect the contradiction and might erroneously count them as genuine possibilities.

'Seemingly possible' worlds are introduced by Veikko Rantala (1975) in his urn-model analysis of logical omniscience. Rantala devised a way of alleviating the mismatch between our model theoretic reasoning about knowledge and our proof theoretic commitments: He asks us to conceive of our epistemic relationship with the world by analogy with an urn from which we can draw balls (individual units of information) one by one over time. With each new piece of information drawn from the urn, we can modify our models. The idea is, simply, that inquiry is a dynamical process in which our model of the world changes with new information. Rantala has provided a formalism which incorporates an intuitively reasonable notion of change in a model. Such change can be understood as a change in the properties of individuals of the model or a change in its universe of discourse.

Representing how the agent's model might dynamically update is one way of thinking about epistemic agency in a more realistic manner. However, on any realistic account of epistemic agency, the agent is likely to consider (albeit inadvertently) worlds in which the classical laws of logic do not hold. In this context, the general problem of establishing a set of epistemic principles for a realistic agent is unavoidable. Rantala's approach provides a way of making the appearance of logical omniscience less threatening, but at the cost of introducing a degree of arbitrariness along with impossible or seemingly possible worlds (see Rantala 1982). In Rantala's discussion of the semantics for impossible worlds (1982) the truth condition is completely free, insofar as any contradiction among an agent's beliefs can be represented by a model containing an impossible world. While logical omniscience is avoided, the price we pay is high, since no real epistemic principles hold broadly enough to encompass impossible and seemingly possible worlds (see Meyer and van der Hoek 1995: 87–88).

Some conditions must be applied to epistemic models such that they cohere with epistemic principles. Computer scientists have proposed that what is being modeled in epistemic logic is not knowledge simpliciter but a related concept which is immune to logical omniscience. The epistemic operator K_cA should be read as 'agent c knows implicitly A,' 'A follows from c's knowledge,' 'A is agent c's possible knowledge,' etc. Propositional attitudes like these should replace the usual 'agent c knows A'. While there exists some variation, the locutions all suggest modeling implicit knowledge or what is implicitly represented in an agent's information state rather than explicit knowledge

(Fagin et al. 1995, and others). The agents neither have to compute knowledge nor can they be held responsible for answering queries based on their knowledge under the implicit understanding of knowledge. Logical omniscience is an epistemological condition for implicit knowledge, but the agent might actually fail to realize this condition.

There are a variety of ways of responding to these kinds of challenges. One rather unpromising approach is to deny that epistemic logic is under any obligation to connect with more general epistemological concerns (see, for example, Lenzen 1978: 34). Rather than treating epistemic logic as a purely formal exercise, a preferable response involves maintaining that epistemic logic does carry epistemological significance but in an inevitably idealized sort of way. One restricts attention to a class of rational agents where rationality is defined by certain postulates. Thus, agents have to satisfy at least some minimal conditions to simply qualify as rational. This is, for example, what Lemmon originally suggests (Lemmon 1959). One such condition would involve assuming that rational agents should acknowledge the laws of logic. For instance, if the agent knows p and $p \rightarrow q$, it should be able to recognize that q follows validly.

These 'rationality postulates' for knowledge exhibit a striking similarity to the laws of modal and epistemic logic. One can, in turn, legitimately attempt to interpret the necessity operator in alethic axioms as a knowledge operator and then justify the modal axioms as axioms of knowledge. While Lemmon constructs the rational epistemic agent directly from the axiomatization of the logic, another way of justifying the epistemic axioms involves reference to their semantical features. This is the line of thought that Hintikka pursued in *Knowledge and Belief*. Hintikka stipulated that the axioms or principles of epistemic logic are conditions descriptive of a special kind of general (strong) rationality. The statements that can be proved false by application of the epistemic axioms are not inconsistent, meaning that their truth is logically impossible. They are, rather, rationally 'indefensible.' Indefensibility is fleshed out as the agent's epistemic laziness, sloppiness or perhaps cognitive incapacity whenever to realize the implications of what he in fact knows. Defensibility, then, means not falling victim of 'epistemic negligence' as Chisholm calls it (Chisholm 1963, 1977). The notion of indefensibility gives away the status of the epistemic axioms and logics. Some epistemic statement for which its negation is indefensible is called 'self-sustaining.' The notion of self-sustenance actually corresponds to the meta-logical concept of validity. Corresponding to a self-sustaining statement is a logically valid statement. But this will again be a statement which is rationally indefensible to deny. So, in conclusion, epistemic axioms can be understood to be descriptions of rationality. This argument is spelled out in detail by Hilpinen (2002).

Common Knowledge and Distributed Knowledge

So far, this essay has discussed the epistemic properties of individual agents. However, many recent developments in epistemic logic concern the study of the formal properties of systems of interacting agents. This section introduces two of the most prominent notions in the study of multi-agent systems: common and distributed knowledge.

When we consider agents who are connected via some network we can study the effect of new information, presented to part of the (or made public to the whole) group. Formal grasp of the role of announcements in a complex network of agents has important practical consequences for our understanding of cooperation and competition. The manner in which new information moves through a multi-agent system and how it

causes individual agent:
of the networks conne
competitive behavior o
behaviors on a prior sha
Common knowledge
the knowledge that on
because not only do all
agents know p and, furt
know that my fellow dri
that I know that they
knowledge is a very pov
knowledge in a broad r
phers from David Hum
Common knowledge
difficulties faced by pro
epistemic condition of
puter scientists to focus
Common knowledge is
social entanglements th
any fool knows (to ech
knowledge see it as can
of collaborative activity
resembling human epis
be in place. Clearly, fo
most basic social intera
munity. As Lewis note
agents that observe it.
mon knowledge with re
and games.
A detailed treatmen
edge is beyond the scop
lent overview). Howev
systems above, it is pos
tems with just a little c
ties of single- and mult
introduced. A modal sy
where it can be assume
be described by the san
for n agents consists of
logic it is possible to ex
agent knows that anotl
It is possible to dev
another agent knows
here it is possible to ex
knows that everyone k
mon knowledge.
One way of defining
the entire group of ager

causes individual agents to modify their beliefs is, in part, dependent on the character of the networks connecting those agents. However, analysis of the cooperative and competitive behavior of agents immediately brings into focus the dependency of these behaviors on a prior shared epistemic medium known as common knowledge.

Common knowledge, as distinguished from shared or mutual knowledge begins with the knowledge that one's fellow agents know that *p*. However, it is more than that, because not only do all the agents in a group know that *p*, they also know that *all other* agents know *p* and, furthermore, that they all know that they all know *p*, and so on. I know that my fellow drivers know that they ought to stop at the red light and they know that I know that they know that they should stop, etc. Thus, in one sense common knowledge is a very powerful kind of phenomenon. The social role played by common knowledge in a broad range of human activities has long been recognized by philosophers from David Hume (1740) to David Lewis (1969).

Common knowledge became a concern for theoretical computer scientists given the difficulties faced by projects in artificial intelligence which focus too narrowly on the epistemic condition of single agents. These difficulties encouraged theoretical computer scientists to focus on the social and conventional features of knowledge and belief. Common knowledge is the basic background knowledge which supports the kinds of social entanglements that are crucial for sophisticated forms of intelligence; it is what any fool knows (to echo John McCarthy (1979)). Philosophical accounts of common knowledge see it as carrying a great burden; supporting the very possibility of the kind of collaborative activity that defines human intelligence. In order to approach anything resembling human epistemic agency, a significant level of social scaffolding needs to be in place. Clearly, for example, we depend on shared epistemic starting points for most basic social interactions, including, prominently, membership in a linguistic community. As Lewis noted (1969) a convention requires common knowledge among the agents that observe it. Robert Aumann (1995) also emphasized the centrality of common knowledge with respect to norms, social and linguistic practices, agent interactions and games.

A detailed treatment of the various formal techniques for tackling common knowledge is beyond the scope of this essay (see Vanderschraaf and Sillari 2007 for an excellent overview). However, given our account of the logical landscape for single-agent systems above, it is possible to introduce some of the main features of multi-agent systems with just a little oversimplification. The primary difference between the semantics of single- and multi-agent semantics is that more than one accessibility relation is introduced. A modal system for *n* agents results from combining *n* modal logics in cases where it can be assumed that the agents are homogeneous in the sense that they can all be described by the same logical system. Thus, in the simplest case, an epistemic logic for *n* agents consists of *n* copies of a certain modal logic. In such an extended epistemic logic it is possible to express that some agent in the group knows a certain fact, that an agent knows that another agent knows a certain fact, etc.

It is possible to develop the logic even further: Not only can an agent know that another agent knows a fact, but they can all know this fact simultaneously. From here it is possible to express that everyone knows that everyone knows that everyone knows that everyone knows that . . . some fact holds. This is what is meant by common knowledge.

One way of defining common knowledge involves defining common knowledge for the entire group of agents rather than partitioning the group of agents into subsets with different common 'knowledges.' Once multiple agents have been added to the syntax, the language is augmented with an additional operator C. CA is then interpreted as 'It is common knowledge among the agents that A.' Well-formed formulas follow the standard recursive recipe with modifications that account for the multiple agents. So, for instance, the operator E is introduced such that EA means 'Everyone knows that A.' EA is defined as the conjunction $K_{1A} \wedge K_{2A} \wedge \ldots \wedge K_n A$.

To semantically interpret n knowledge operators, binary accessibility relations R_n are defined over the set of possible worlds W. A special accessibility relation, $R°$, is introduced to interpret the operator of common knowledge. The relation must be flexible enough to express the relationship between individual and common knowledge. The idea is to let the accessibility relation for C be the transitive closure of the union of the accessibility relations corresponding to the knowledge operators for the individual agents.

The model $\mathcal{M}$ for an epistemic system with n agents where the agents have common knowledge is a structure $\mathcal{M} = \langle W, R_1, R_2, \ldots, R_n, R°, \varphi \rangle$, where W is a non-empty space of possible worlds, $R_1, R_2, \ldots, R_n, R°$ are accessibility relations over W for which $R° = (R_1 \cup R_2 \cup \ldots \cup R_n)$ and φ again is the function assigning worlds to atomic propositional formula $\varphi : atom \mapsto P(W)$. The semantics for the Boolean connectives remain intact. The formula $K_i A$ is true in a world w, i.e., $\mathcal{M}, w \models K_i A$ iff $\forall w' \in W$: if $R_i ww'$, then $\mathcal{M}, w' \models A$. So, A is common knowledge in a world w, when $\mathcal{M}, w \models CA$ iff $R° ww'$ implies $\mathcal{M}, w' \models A$.

Varying the properties of the accessibility relations $R_1, R_2, \ldots, R_n$, results in different epistemic logics. For instance system **K** with common knowledge is determined by all frames, while system **S4** with common knowledge is determined by all reflexive and transitive frames. Similar results can be obtained for the remaining epistemic logics (Fagin et al. 1995).

Informally speaking, the claim that some proposition is common knowledge is extremely strong. Therefore, those propositions for which we can claim common knowledge tend to be very weak. To claim that some proposition is common knowledge, that everyone knows that everyone knows A, implies that everybody knows A, which implies individual knowledge of A.

If we think of common knowledge as involving very strong claims about the epistemic state of a group of agents, at the opposite end of the spectrum is the notion of distributed knowledge. Distributed knowledge is an epistemic property which captures the idea that there is an aggregated store of knowledge in a group, some of which might not necessarily be possessed by any individual member of the group. If even one agent in a group knows A then A is part of the distributed knowledge of the group. Where common knowledge is very strong (and its argument is rather weak), distributed knowledge is weaker, but can be obtained for much stronger facts, as we shall see.

One way to think about distributed knowledge is to recognize its relationship to a communication network. Something is distributed knowledge in a group if it could be known by the individuals were they able to talk to each other. For instance, in a crowd of 100 people, when two people have the same birthday, this might not be individually known, but could be known if the members were able to talk to each other. So, to take a simple case, if A knows that B is older than C or D and E knows that B is not older than D, then while no individual agent knows that A is older than C, that knowledge is distributed throughout the group and could be elicited given the right kinds of communication. This is a kind of knowledge that can be ascribed to some collection of agents

(given certain conditic
isolation. Of course the
gated store of distribut

When we consider
entific communities, d
of the properties of in
company as a whole k
to the distributed knov
the nature of cognitior
brain and body, those p
mal study of distribute
in a community might
until we begin to exam

The most famous ex
scenarios like the mud
father. k children hav
they cannot see whet
their father, do not ch
scenario involves the
and on the epistemic s
is at least one child w
have mud on their for
forward. When their f
step forward.

In order to explain t
the child with mud o
other child with mud
where $k = 2$ we can in

$k = 2$.
$a^* \rightarrow b^*$
↓[[points from a*
c

Muddy child a can see
After the father's first
not know whether b h
forehead on either a
the muddy forehead t
of muddy headed a tc
second announcemer
and two muddy child
children would proce
in Meyer and van der

There is a range of
multi-agent systems a
explained. The kinds
muddy children, whic

(given certain conditions) but which need not necessarily be ascribed to any agent in isolation. Of course the knowledge that an individual agent has is also part of the aggregated store of distributed knowledge.

When we consider the epistemic properties of groups such as corporations or scientific communities, distributed knowledge that the group exhibits is likely to be one of the properties of interest. For instance, if we are entitled to say that the electric company as a whole knows how to maintain the power supply, we do so by reference to the distributed knowledge that exists in the group. Similarly, in debates concerning the nature of cognition which relies on resources beyond the confines of the individual brain and body, those problems related to the possibility of extended cognition, the formal study of distributed knowledge might prove useful. The distribution of knowledge in a community might seem like a rather nebulous or metaphysically extravagant notion until we begin to examine it in a formal setting.

The most famous example of the formal study of features of group knowledge involves scenarios like the muddy children problem. The scenario involves n children and their father. k children have mud on their foreheads. The children can see each other but they cannot see whether they have mud on their own foreheads. The children trust their father, do not cheat, are rational and do not communicate with one another. The scenario involves the effect of the father's announcement on the behavior of the group and on the epistemic states of the members of the group. The father announces: 'There is at least one child with mud on its forehead. Will all the children who know they have mud on their foreheads please step forward?' If k is greater than one, no child steps forward. When their father makes his announcement the k-th time all muddy children step forward.

In order to explain this scenario, we first consider the simplest case. Where $k = 1$ then the child with mud on its head knows that it must be the muddy one since it sees no other child with mud on its head. So, when $k = 1$ the explanation is clear. In the case where $k = 2$ we can imagine the following scenario with two muddy children a and b:

$k = 2.$
$a^* \rightarrow b^*$
↓[[points from a* to c]]
c

Muddy child a can see that b is muddy, but it does not know whether it is muddy itself. After the father's first announcement, when b does not step forward, a knows that b does not know whether b has a muddy forehead. This means that b sees at least one muddy forehead on either a or c. Since a can see that c has a clean forehead, a reasons that the muddy forehead that b saw, was a's. b reasons in the same way based on the failure of muddy headed a to step forward after the father's first announcement. So, after the second announcement, both children step forward. By induction from cases with one and two muddy children, we can easily see how cases with greater numbers of muddy children would proceed. A full treatment of the muddy children problem can be found in Meyer and van der Hoek (1995: 56) or in Fagin et al. (1995: 3).

There is a range of cases like these in which we must account for the interaction of multi-agent systems and in which certain collective features of group behavior must be explained. The kinds of epistemic systems under consideration include cases, like the muddy children, which are sensitive to the introduction of new information via public announcement and in which the interactions of the agents contributes another dynamical component which has an effect on the unfolding states of the system. The muddy children problem is a simple example of the kinds of dynamical features that epistemic logic tackled in the 1980s and 1990s.

Game Theory, Belief Revision and the Properties of Agents

In multi-agent settings, it is natural to consider the role of competition and cooperation. Thus, as epistemic logic began to attend to the dynamics of groups, game theory began to play a more prominent role in reflections on epistemic agency. Aumann, van Benthem, Brandenburger, Fagin, Halpern, Keisler, Moses, Stalnaker, Vardi and others have contributed to uncovering important features of agent rationality showing how game theory adds to the general understanding of notions such as knowledge, belief and belief revision. By the end of the 1990s, Baltag, Moss and Solecki had combined epistemic logic with belief revision theory to study actions and belief updates in games (Baltag et al. 1998).

In the 1980s, Alchourrón, Gärdenfors and Makinson developed a theory of belief revision theory (AGM) which provides an account of rational change of belief in light of novel evidence (Alchourrón 1985; Gärdenfors 1988). Expansions, contractions and revisions in an agents' set of belief are characterized formally. 'Revision' here means additions of beliefs to the agent's belief-set which maintain consistency. Revision is distinguished from simple expansion, which takes place without regard for consistency, and contraction, where beliefs are removed from the set. In order to be considered rational, an agent who revises his beliefs must obey the AGM postulates. Taking K to be the agent's initial set of beliefs and * to be the revision operation and letting A be the additional information that the agent encounters, the basic postulates are presented by Robert Koons (2009) as follows:

(1) $K*A$ is closed under logical consequence.
(2) A belongs to $K*A$.
(3) $K*A$ is a subset of the logical closure of $K \cup \{A\}$.
(4) If $\neg A$ does not belong to K, then the closure of $K \cup \{A\}$ is a subset of $K*A$.
(5) If $K*A$ is logically inconsistent, then either K is inconsistent, or A is.
(6) If A and B are logically equivalent, then $K*A = K*B$.
(7) $K*(A \& B)$ is a subset of the logical closure of $K*A \cup \{B\}$.
(8) If $\neg B$ does not belong to $K*A$, then the logical closure of $K*A \cup B$ is a subset of $K*(A \& B)$.

Following Andre Fuhrmann's development of the idea of translating AGM into dynamical modal logic (1988, 1991), de Rijke also showed that the AGM postulates governing expansion and revision can be translated into the object language of dynamic modal logic (de Rijke 1994). At about the same time, Segerberg demonstrated how the theory of belief revision could be formulated in modal logic. Segerberg merged the static first generation doxastic logic with the dynamics of belief change into 'dynamic doxastic logic' (Segerberg 1995). Doxastic operators in the logic of belief like $B_c A$ can be captured by AGM in the sense that 'A is in c's belief-set T', or $\neg B_c \neg A$ becomes '$\neg A$ is not in c's belief-set T.' An immediate difference between the two perspectives is that while

AGM can express dynamic operations on belief-sets like expansions ('A is in c's belief-set T expanded by D,' i.e., $A \in T + D$), revisions ('A is in c's belief-set T revised by D,' i.e., $A \in T{*}D$), and contractions ('A is in c's belief-set T contracted by D,' i.e. $A \in T - D$), no such dynamics are immediately expressible in the standard language of doxastic logic. On the other hand, action languages include operators like $[\mu]$ and $\langle\mu\rangle$ which are prefixed to a well-formed formula A. On Segerberg's interpretation, $[\mu]A$ ($\langle\mu\rangle A$) mean that 'after every (some) way of performing action μ it is the case that A.' By introducing three new operators [+], [*], and [−] into the doxastic language, the three dynamic operations on belief-sets may be rendered as $[+D]B_cA$, $[{*}D]B_cA$ and $[-D]B_cA$.

After revising the original belief revision theory such that changes of beliefs happen in 'hypertheories' or concentric spheres enumerated according to entrencement, Segerberg (1999a, 1999b) provided several axiomatizations of the dynamic doxastic logic together with soundness and completeness results. The dynamic doxastic logic paradigm can also be extended to iterated belief revision as studied by Lindström and Rabinowicz (1997) and accommodate various forms of agent introspection. A related approach drawn up by van Ditmarsch, van der Hoek and Kooi's new 'dynamic epistemic logic' studies how information changes and how actions with epistemic impact on agents may be modeled (van der Hoek et al. 2003; van Ditmarsch et al. 2007). For a more detailed discussion of belief revision theory, see André Fuhrmann "Theories of Belief Change," Chapter 56 in this volume.

One might also choose to endow the agents with *epistemic capacities* facilitating special epistemic behaviors. Fagin, Halpern, Moses and Vardi have, for instance, considered 'perfect recall' (Fagin et al. 1995): interacting agents' knowledge in the dynamic system might increase as time goes by but the agents might still store old information. The agent's current local state is an encoding of all events that have happened so far in the run. Perfect recall is, in turn, an epistemic recommendation telling the agent to remember his earlier epistemic states.

There are other structural properties of agents being studied in the literature of dynamic epistemic logics. In an epistemic logic suited for modeling various games of imperfect information, van Benthem (2000) refers to such properties as 'styles of playing.' Properties such as 'bounded memory,' various 'mechanisms for information updates' and 'uniform strategies,' infallibility, consistency etc. have been investigated. Yoav Shoham and Kevin Leyton-Brown's *Multi-Agent Systems* (2009) provides an updated overview of the relevant literature on game theory and belief revision in a multi-agent setting. Agents as explicitly learning mechanisms are also integral parts of Kelly's (1996) computational epistemology and a related approach called modal operator epistemology (Hendricks 2001, 2003). Researchers in artificial intelligence have additionally been trying to describe and specify the behavior of intelligent/rational agents by extensions of epistemic of logic by augmenting logics of time, action and belief with modalities for desires and intentions (see Meyer 2003, in particular, his discussion of the BDI-framework of Rao and Georgeff in Section 5.2).

Acknowledgments

I owe a special debt of gratitude to Wiebe van der Hoek, whose detailed criticism of this chapter has improved it considerably. I am very grateful to Vincent Hendricks for his comments and for graciously permitting me to use some of our previously co-authored material in this chapter. Specifically, pages are largely taken from our co-authored 2006.

Thanks also to Duncan Pritchard, Clifford Hill and Emmanuel Genot for helpful criticism of an earlier draft of this essay.

References

Alchourrón, C.E., Gärdenfors, P. and Makinson, D. (1985). "On the Logic of Theory Change," *Journal of Symbolic Logic* 50: 510–530.

Aumann, R. (1995). "Backward Induction and Common Knowledge of Rationality," *Games and Economic Behavior* 8: 6–19.

Austin, J. (1975). *How To Do Things with Words*. Oxford: Oxford University Press.

Baltag, A., Moss, L.S. and Solecki, S. (1998). "The Logic of Public Announcements, Common Knowledge, and Private Suspicion." *Proceedings of TARK 1998*. Los Altos: Morgan Kaufmann Publishers, 43–56.

Benthem, J.F.A.K. van (2000). "Logic and Game Theory—Close Encounters of the Third Kind," in van Loon, I., Mints, G. and Muskens, R. (eds.), *Proceedings of LLC99*. Stanford: CSLI Publications.

Bull, R. and Segerberg, K. (1984). "Basic Modal Logic," in Gabbay, D. and Guenthner F. (eds), *Handbook of Philosophical Logic*, vol. II: *Extensions of Classical Logic*. Dordrecht: D. Reidel Publishing Company, 1–88.

Chisholm, R. (1963). "The Logic of Knowing," *Journal of Philosophy* 60, 773–795.

Chisholm, R. (1977). *Theory of Knowledge* (2nd ed.). Englewood Cliffs, NJ: Prentice Hall.

Ditmarsch, H. van, van der Hoek, W. and Kooi, B. (2007). *Dynamic Epistemic Logic*. Dordrecht: Springer.

Fagin, R., Halpern, J.Y., Moses, Y. and Vardi, M.Y. (1995). *Reasoning about Knowledge*. Cambridge: MIT Press.

Fuhrmann, A. (1988). *Relevant Logic, Modal Logic and Theory Change*, Ph.D. Thesis, Department of Philosophy and Automated Reasoning Project, Institute of Advanced Studies, Australian National University, Canberra.

Fuhrmann, A. (1991). "On the Modal Logic of Theory Change," in Fuhrmann, A. and Morreau, M. (eds.), *The Modal Logic of Theory Change*. Lecture Notes in Artificial Intelligence 465. Heidelberg-Berlin-New York: Springer.

Gärdenfors, P. (1988). *Knowledge in Flux—Modelling the Dynamics of Epistemic States*. Cambridge: MIT Press.

Gettier, E. (1963). "Is Justified True Belief Knowledge?" *Analysis* 23: 121–123.

Hendricks, V.F. (2001). *The Convergence of Scientific Knowledge—A View from the Limit*. Trends in Logic: Studia Logica Library Series, Dordrecht: Kluwer Academic Publishers.

Hendricks, V.F. (2003). "Active Agents," *Journal of Logic, Language and Information*, van Benthem, J. and van Rooy, R. (eds.), 12, 4: 469–495.

Hendricks, V.F. (2005). *Mainstream and Formal Epistemology*. New York: Cambridge University Press.

Hendricks, V.F. and Symons, J. (2006). "Where's the Bridge? Epistemic Logic and Epistemology," *Philosophical Studies* 128, 1: 2–26.

Hendricks, V.F. and Symons, J. (2007). "Epistemic Logic," *The Stanford Encyclopedia of Philosophy (Spring 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2009/entries/logic-epistemic/>.

Hilpinen, R. (2002). "Deontic, Epistemic and Temporal Modal Logics," in Jacquette, D. (ed.), *A Companion to Philosophical Logic*. Malden, MA: Blackwell, 491–509.

Hilpinen, R. (2006). "Jaakko Hintikka on Epistemic Logic and Epistemology," in Auxier, R.E. (ed.), *The Philosophy of Jaakko Hintikka*. Chicago: Open Court Publishing Company, 783–818.

Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell: Cornell University Press.

Hintikka, J. (1975). "Impossible Possible Worlds Vindicated," *Journal of Philosophical Logic* 4: 475–484.

Hintikka, J. (2007). "Epistemology without Knowledge and Without Belief," in Hintikka, J., *Socratic Epistemology: Explorations of Knowledge Seeking by Questioning*. Cambridge: Cambridge University Press, 11–37.

Hocutt, M.O. (1972). "Is Epistemic Logic Possible?" *Notre Dame Journal of Formal Logic* 13, 433–453.

Hoek, W. van der, Ditmarsch, H. van, Kooi, B. (2003). "Concurrent Dynamic Epistemic Logic," in Hendricks, V.F., Jørgensen, K.F. and Pedersen, S.A. (eds.), *Knowledge Contributors*. Synthese Library, vol. 322. Dordrecht: Kluwer Academic Publishers, 105–173.

Hume, D. (1976 [1740, 1888]). *A Treatise of Human Nature*, ed. L.A. Selby-Bigge. Rev. 2nd edn, ed. P.H. Nidditch. Oxford: Clarendon Press.

Kelly, K. (1996). *The Logic of*

Koons, R. (2009). "Defeasibl Edward N. Zalta (ed.), U ble/>.

Kripke, S. (1963). "Semantic *der Matematik* 9: 67–96.

Kutschera, F. von (1976). *Ein*

Lemmon, E.J. (1959). "Is Ther Volume XXXIII, 23–40.

Lemmon, E.J. (1977). *An Intr* Publishers.

Lenzen, W. (1978). *Recent W*

Lewis, D. (1969). *Convention*

Lindström, S. and Rabinowi and Ramsey Conditionals *Measure. Uppsala Philosoph*

McCarthy, J. (1979). "Ascr Stanford University.

Meyer, J.-J. Ch. (2003). "Mc *Handbook of Philosophical L*

Meyer, J.-J.Ch. and Hoek, W in Theoretical Computer S

Moore, G.E. (1952). *The Phil* Publishing Company.

Rantala, V. (1975). "Urn M *Symbolic Logic* 4: 455–474.

Rantala, V. (1982). "Quant *Logica* 41, 1: 41–66.

Rijke, M. de (1994). "Meetin Knowledge Representatio MIT Press, 170–196.

Segerberg, K. (1995). "Belie 535–553.

Segerberg, K. (1999a). "Th *Philosophy*, 1.

Segerberg, K. (1999b). "A C *Essays Dedicated to Ann-* *Philosophical Studies*, 49: 1

Shoham, Y. and Leyton-Br *Foundations*, Cambridge: C

Stalnaker, R. (1996). "Know 12: 133–163.

Stalnaker, R. (2006). On Lo

Vanderschraaf, P. and Sillar *(Spring 2009 Edition)*, Edw common-knowledge/>.

Voorbraak, F. (1993). *As Far* Utrecht. Published as vol

Wright, G.H. von (1951). *A*

Kelly, K. (1996). *The Logic of Reliable Inquiry*. New York: Oxford University Press.

Koons, R. (2009). "Defeasible Reasoning," *The Stanford Encyclopedia of Philosophy (Spring 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2009/entries/reasoning-defeasible/>.

Kripke, S. (1963). "Semantical Analysis of Modal Logic," *Zeitschrift für Matematische Logik und Grundlagen der Matematik* 9: 67–96.

Kutschera, F. von (1976). *Einführung in die intensional Semantik*. Berlin: W. de Gruyter.

Lemmon, E.J. (1959). "Is There Only One Correct System of Modal Logic?," *Aristotelian Society Supplementary*, Volume XXXIII, 23–40.

Lemmon, E.J. (1977). *An Introduction to Modal Logic*, in collaboration with D. Scott. Oxford: Basil Blackwell Publishers.

Lenzen, W. (1978). *Recent Work in Epistemic Logic*, in *Acta Philosophica Fennica* 30: 1–219.

Lewis, D. (1969). *Convention: A Philosophical Study*. Cambridge: Harvard University Press.

Lindström, S. and Rabinowicz, W. (1997). "Extending Dynamic Logic: Accommodating Iterated Beliefs and Ramsey Conditionals within DLL," in Lindahl, L., Needham, P. and Sliwinski, R. (eds.), *For Good Measure. Uppsala Philosophical Studies*, 46: 123–153.

McCarthy, J. (1979). "Ascribing Mental Qualities to Machines." Technical Report STAN-CS-79–725 Stanford University.

Meyer, J.-J. Ch. (2003). "Modal, Epistemic and Doxastic Logic," in Gabbay, D. and Guenthner, F. (eds.), *Handbook of Philosophical Logic* (2nd edition) Vol. 10, Dordrecht: Kluwer, 1–38.

Meyer, J.-J.Ch. and Hoek, W. van der (1995). *Epistemic Logic for AI and Computer Science*. Cambridge Tracts in Theoretical Computer Science 41. Cambridge: Cambridge University Press.

Moore, G.E. (1952). *The Philosophy of G. E. Moore*. (2nd edition) ed. Paul Arthur Schilpp, New York: Tudor Publishing Company.

Rantala, V. (1975). "Urn Models: A New Kind of Non-Standard Model for First-Order Logic," *Journal of Symbolic Logic* 4: 455–474.

Rantala, V. (1982). "Quantified Modal Logic: Non-Normal Worlds and Propositional Attitudes," *Studia Logica* 41, 1: 41–66.

Rijke, M. de (1994). "Meeting Some Neighbours: A Dynamic Modal Logic Meets Theories of Change and Knowledge Representation," in Eijck, J. van and Visser, A. (eds.), *Logic and Information Flow*, Cambridge: MIT Press, 170–196.

Segerberg, K. (1995). "Belief Revision from the Point of View of Doxastic Logic," *Bulletin of the IGPL*, 3: 535–553.

Segerberg, K. (1999a). "The Basic Dynamic Doxastic Logic of AGM," *Uppsala Prints and Preprints in Philosophy*, 1.

Segerberg, K. (1999b). "A Completeness Proof in Full DDL," in Sliwinski, R. (ed.), *Philosophical Crumbs: Essays Dedicated to Ann-Mari Henschen-Dahlqvist on the Occasion of her Seventy fifth Birthday, Uppsala Philosophical Studies*, 49: 195–207.

Shoham, Y. and Leyton-Brown, K. (2009). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge: Cambridge University Press.

Stalnaker, R. (1996). "Knowledge, Belief and Counterfactual Reasoning in Games," *Economics and Philosophy*, 12: 133–163.

Stalnaker, R. (2006). On Logics of Knowledge and Belief, *Philosophical Studies* 128, 169–199.

Vanderschraaf, P. and Sillari, G. (2007). "Common Knowledge," *The Stanford Encyclopedia of Philosophy (Spring 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2009/entries/common-knowledge/>.

Voorbraak, F. (1993). *As Far as I Know*. Epistemic Logic and Uncertainty, Dissertation, Utrecht University, Utrecht. Published as volume 7 in *Queastiones Infinitae*. Department of Philosophy, Utrecht University.

Wright, G.H. von (1951). *An Essay on Modal Logic*. Amsterdam: North-Holland Publishing Company.