Logical Omniscience vs. Logical Ignorance On a Dilemma of Epistemic Logic

Ho Ngoc Duc

Institute of Logic and Philosophy of Science, University of Leipzig
P.O. Box 920, D-04009 Leipzig, Germany
phone: +49 (0341) 97 35 775, fax: +49 (0341) 97 35 798
e-mail: duc@informatik.uni-leipzig.de

Abstract. We propose in the paper a new solution to the so-called Logical Omniscience Problem of epistemic logic. Almost all attempts in the literature to solve this problem consist in weakening the standard epistemic systems: weaker sytems are considered where the agents do not possess the full reasoning capacities of an ideal reasoner. We shall argue that this solution is not satisfactory: in this way omniscience can be avoided, but many intuitions about the concepts of knowledge and belief get lost. We shall show that axioms for epistemic logics must have the following form: if the agent knows all premises of a valid inference rule, and if she thinks hard enough, then she will know the conclusion. To formalize such an idea, we propose to "temporalize" epistemic logic, that is, to introduce a temporal component into the language. We develop a logic based on this idea and show that it is suitable for formalizing the notion of explicit belief.

Keywords: Knowledge representation, Logics of knowledge and belief, Logical omniscience, Dynamic epistemic logic, Knowledge and time

1 Introduction and Preliminaries

Epistemic logic, or the logic of the concepts of knowledge and belief, has established as an autonomous branch of logic since the work of Hintikka ([6].) The subject has been studied extensively by philosophers, linguists and, more recently, computer scientists. Most systems of epistemic logic have been developed in analogy to modal logic. Many results and methods of modal logic can be tranferred to epistemic logic, notably the techniques of the possible worlds semantics. There are, however, severe objections against the modal approach in epistemis logic. The most serious problem of this approach is perhaps the so-called "logical omniscience problem (LOP)." It can be described informally

* I thank my teachers Peter Steinacker and Heinrich Herre for their invaluable help during the last several years. I have learned much from Werner Stelzner's works on non-modal epistemic logic. I had the chance to discuss preliminary versions of this paper with several people at different occasions. I thank all of them for their comments which have helped me very much to improve the presentation of the paper.

as follows. According to the standard approach, the agent should be an ideal logician and reasoner in the following sense: she knows all logical truths, can (actually) draw all consequences of a certain sentence and can identify all logical equivalences of a given sentence. Such requirements are clearly too strong for a real agent, human or non-human. Thus, the standard systems cannot capture the notion of knowledge and belief adequately. The problem of logical omniscience is a severe obstacle for the applicability of epistemic logic. For that reason, many attempts have been undertaken to solve this problem. The goal of my paper is to assess how successful these attempts can be and then to propose another solution to the LOP.

Before going on, let us state the problem more precisely. For simplicity's sake I consider only the propositional case with only one agent. I use the terms "knowledge" and "belief" as synonyms in the paper. (It is therefore not distinguished between epistemic and doxastic logic.) Knowing p only means "having good reasons to believe that p", but does not imply that p is true. By "modal epistemic logic" I mean, as proposed by Wuttich in [13], those systems of epistemic logic which are developed in the modal tradition. The language of modal epistemic logic is built up from a set At_M of propositional letters using the usual Boolean connectives of negation and implication and the unary operator K. The set of formulae of this language will be denoted by Fml_M. (The index M indicates the modal approach.) Formally, Fml_M is the least set which includes At_M and is closed under the following rules: if p and q are members of Fml_M then so are $\neg p$, $p \supset q$ and Kp. The formula Kp is to be read: "the agent knows p". The other Boolean connectives are defined in the usual way. The minimal normal modal epistemic logic K is the least system which contains all theorems of the propositional calculus (PC) together with all instances of the axiom schema K, i.e., the schema $K(p \supset q) \supset (Kp \supset Kq)$, and is closed under modus ponens and the so-called Gödel rule: if p is a theorem then so is Kp. It is well-known that K is characterized by the class of all Kripke models (cf. [1].)

The logical omniscience problem for this minimal normal logic can be stated as follows. The following inference rules are valid for the system K:

(R1) If p is a theorem then so is Kp
(R2) If $p \supset q$ is a theorem then so is $Kp \supset Kq$
(R3) If $p \leftrightarrow q$ is a theorem then so is $Kp \leftrightarrow Kq$

The three rules (R1), (R2), (R3) are called Gödel rule, monotony rule, and congruence rule, respectively. These rules not only hold for K, but also for all of its normal extensions. An agent who is described by such a logic is said to be logically omniscient, because she knows all logical truths (according to R1), she knows all logical consequences of a sentence that she knows (according to R2), and she can identify all logical equivalent sentences of a given sentence, according to (R3). Such an agent cannot be a real one. The logic K and its normal extensions can only describe ideal agents. For modelling realistic agents we need other logics which do not suffer from the LOP. In the next section we shall discuss some common ways to solve this problem. After showing that the

strategy of weakening epistemic logic has many disadvantages I shall propose an alternative approach to the problem. The intuitions of my strategy will be explained in section 3. A formal system which can cope with the problems of the traditional approaches will be developed in section 4 following the new strategy.

2 Strategies to Avoid Logical Omniscience

An obvious strategy to solve the logical omniscience problem is to weaken epistemic logic. One denies the universal validity of the mentioned inference rules (R1)–(R3). In fact, almost all attempts to solve the LOP have in common that they consider systems that are weaker than the standard modal epistemic logics (cf. [3], [4], [7], [8], [9], [10], [11], [12], [13].) But here some care is needed: some systems solve the original version of the LOP, but not other versions. For example, the agent in Levesque's logic ([9]) does not know all truths of classical logic, but she knows all theorems of a relevance logic. In the same way, an agent in Ho Ngoc Duc's system ([8]) does not know all classical theorems, but knows all theorems of a three-valued logic. Thus, such solutions cannot be regarded satisfactory. Modal systems which are not normal can also be used to describe an agent who does not know all logical truths (cf. [11], for example.) However, if the monotony rule and/or the congruence rule are valid for a system, then the agent in this system must still be viewed as ideal: real agents simply never achieve such reasoning capacities which can be described by these rules. Thus, a solution to the LOP must provide us with a logic for which neither of the rules (R1)–(R3) is valid. We shall note that the notion of "theorem" in these rules not only applies to classical logic, but also to other logical systems as well.

A number of systems have been proposed which are not closed under any of the rules (R1)–(R3). To construct such a system we can postulate, for example, that the agent only knows some "obvious" logical truths, but not necessarily the "more complicated" ones. We can assume that the agent can draw all "obvious" consequences, but not any arbitrary consequence, of a certain sentence. We can consider logics which allow only the rules of classical logic (e.g., modus ponens) as inference rule. The properties of the knowledge concept must then be described by a set of axioms. The more axioms are assumed, the more rational is the agent. With the aid of such weak epistemic logics we can classify the agents according to their logical capacities. (cf. [12], [13].) In general those weak epistemic logics lack logical omniscience. Besides this axiomatic approach we can also pursue a more semantical approach. One can show that logical omniscience can be avoided if one allow "impossible possible worlds" in which the valuation of the sentences of the language is arbitrary. In other words, the logical laws do not hold in the "impossible possible worlds" ([11], [10]). Another solution is to introduce a new operator of awareness into the language and to require that belief include awareness ([4].) Because it is possible that the agent is aware of some sentence but she is not aware of its logical consequences or its equivalent sentences, the monotony rule (R2) and the congruence rule (R3) do not hold in general.

Although the approach with impossible possible worlds and the approach

with the awareness operator solve the LOP technically, they cannot be regarded satisfactory. New problems arise in these approaches besides the old problems of the possible worlds approach. Here I shall not discuss these problems in details, nor shall I try to improve any of these approaches. I shall rather present a more fundamental criticism of the common strategy of all these approaches, namely the strategy of weakening epistemic logic.

The discussion of the logical omniscience problem in the literature has concentrated mainly on the issue: in which way can logical omniscience be avoided. But the LOP has another aspect which is often overlooked in the discussion: what is left if one restrict the reasoning capacities of the agents, for example by denying the validity of the rules (R1), (R2) and (R3)? Is there still a reasonable way to describe the agent's knowledge if the regularities of the agent's knowledge is too weak to be described by these rules?

An attempt to cope with this challenge is to postulate axioms which describe the regularities of the agent's knowledge. Such axioms shall express the intuitive idea that the agent is somehow rational, or logical. The axioms are generally of the form: the belief set of the agent is closed under a certain valid rule of inference of logic, i.e., if all premises of the rule are known, then the conclusion is known. (This is also the general form of a theorem of a standard epistemic logic.) In this way we can get subsystems of the logic K which do not suffer from the LOP. Let us see how far such a strategy may take us. We consider a simple example of a ultra-weak epistemic logic.

Definition 1. Let L be the logic whose axioms consist of all theorems of the propositional calculus (PC) and all instances of the schema K. The only inference rule of L is modus ponens.

L is thus the system obtained from the smallest normal system K by dropping the Gödel rule (without any substitute!) L is a very weak epistemic logic lying between PC and K. It describes an agent, or more correctly, a class of agents who can only use modus ponens reliably: whenever they know the premisses of modus ponens then they can infer the conclusion (i.e., if Kp and $K(p \supset q)$ then Kq.) It is not required that such an agent know any other inference rule or axiom. It is clear that the agent lacks logical omniscience.

It is possible to find an adequate semantics for L using impossible possible worlds. However, we can also provide L with a simple truth-functional semantics.

Definition 2. An L-evaluation is a function V from the set Fml_M of formulae to the set $\{0, 1\}$ of truth-values, which satisfies the following conditions for arbitrary formulae:

- $V(\neg p) = 1$ if and only if $V(p) = 0$
- $V(p \supset q) = 0$ if and only if $V(p) = 1$ and $V(q) = 0$
- If $V(Kp) = 1$ and $V(K(p \supset q)) = 1$ then $V(Kq) = 1$

The first two clauses guarantee that all classical theorems are valid, and the last one ensures the validity of the axiom schema K. It can be seen easily that

if a formula is provable in L then it has the value 1 under all L-valuations. It turns out that the converse is also true, that is, the system L is complete with respect to the given semantics.

Theorem 3. *If $V(p) = 1$ for all L-valuations V then p is L-provable.*

Proof. (Sketch) If p is not a theorem of L, then $\neg p$ is L-consistent. Therefore $\neg p$ is contained in a maximal L-consistent set X, by the standard Henkin argument. The valuation V defined by V(p)=1 iff $p \in X$ for all $p \in Fml$ satisfies the three conditions above, as we can check easily. Thus, V is an L-valuation and V is a model of $\neg p$. By contraposition we have the desired result. □

As the example of system L shows, the strategy of weakening epistemic logic allows us to consider agents who are very restricted in their reasoning capacities. We can describe and classify them according to their rationality. We would have a hierarchy of agents: some agents are ideal, who are modelled by the logic K or its extensions, others are less ideal and can be described by, e.g., the logic L. The agents described by the intermediate systems between L and K are not as ignorant as those of L, neither are they ideal reasoners as those of K.

The strategy of weakening epistemic logic solves the logical omniscience problem. However, the disadvantages cannot be overlooked. First, this approach is only suited to analyze static knowledge, that is, we can at most describe the knowledge sets at one single time point. This is of course not a deficiency of these logics alone, but of most epistemic logics developed up to now. Second, the categories of agents we describe and classify by our logics are merely imaginary: they do not exist in reality. It is very implausible to assume that there are agents who always think in some fixed patterns which can be captured by one of our logics. Each agent represents rather some mixture of several logics, at some time point they can be described by one, at other time points by another, and still at other by none of our logics at all. Third, we have the feeling that our logics are too weak. Surely, we want to avoid logical omniscience. On the other hand, we are interested in having epistemic logics which are strong enough to allow sufficiently many conclusions from a given set of facts we know about the agent's propositional attitudes. For example, if we know that the agent know the conjunction $p \wedge q$ then we may expect that she knows p and that she knows q. This inference is, however, not sound for L and many related systems. We want to have agents who do know at least a (sufficiently) large class of logical truths, and can draw sufficiently many conclusions from their knowledge. This is the dilemma on logical omniscience on the one side and logical ignorance on the other side. That is why I ask the question before: what is left from belief logic if we deny the validity of the congruence rule and stronger principles? What we need is something between two extremes. Can we have some reasonable thing like that?

My goal is to show that we can solve this dilemma. I shall now propose another strategy to solve the logical omniscience problem which also solves the problem of logical ignorance. My strategy starts with the observation that the

laws of (classical) logic are not sentences about the world, they do not tell us anything about what is the case in the world. If we say that the epistemic agent knows the laws of logic, we do not mean that she knows some facts about the world, but rather that she is able to use these laws to draw conclusions from what she already knows. The laws of logic are what the agent knows implicitly; she does not need to possess them permanently. It suffices if she can recall them when she needs them in order to infer new information from her explicit data base. At a given time the set of logical laws that the agent has in her memory is restricted, and so is the set of logical consequences of all what she knows explicitly. In this way we can achieve a good tradeoff between logical omniscience and logical ignorance: the agent is surely not omniscient with respect to her actual or explicit knowledge, but neither is she logically ignorant. Our task is to find a suitable way to express this idea formally.

3 Temporalizing Epistemic Logic

Let us consider an inference rule, say R. It can be a valid inference rule of classical logic, or some other (non-classical) logic, for example, intuitionist logic, conditional logic or relevant logic. Assume that the agent accepts R as valid and she can use R. What does it mean? In the modal approach we formalize this idea by an axiom saying that the knowledge set of the agent is closed under this rule, that is, if all premises of the rule are known then the conclusion of R is also known. However, as we noted above, it is only true of implicit knowledge. In the context of explicit knowledge it must mean something different. It means rather that, if the agent knows all premises of the rule, and if she perform the inference according to the rule R, then she will know the conclusion. The agent does not know the conclusion automatically, but rather as the result of some action, viz. the (mental) action of performing the corresponding inference. If she does not perform this action, then we cannot require her to know (or believe) the conclusion, although this conclusion may seem to be an obvious consequences of the sentences under consideration. Especially, a logical axiom can be viewed as an inference rule without any premises. We cannot require the agent to know all axioms automatically and permanently, but she must also carry out some action before she can acquire knowledge of a certain axiom. It is possible that the agent knows all logical truths, but merely in principle. This knowledge is only implicit. Factually she never knows them all at once explicitly. As the idea of "time shift" ("before / after the reasoning") is implicit in the idea of performing an inference we call our approach the strategy of "temporalizing epistemic logic."

For formalizing the reasoning actions it is natural to use (a form of) dynamic logic (cf. [5].) We can associate each valid inference rule with an atomic action. Thus, we can add a set of basic actions to the language of epistemic logic. The set of formulae now includes formulae like $[R]Kp$ or $\langle R\rangle Kp$ with the intended meaning: "always after using rule R (or sometimes after using R) the agent knows p". The formalization of the idea that the agent accepts and is able to use an inference rule is straightforward. For example, the idea that the agent accepts

modus ponens can be formalized by the axiom: $Kp \wedge K(p \supset q) \supset \langle MP \rangle Kq$. This axiom says no more than if the agent knows p and she also knows that p implies q, then after a suitable inference step she will know q.

As the axioms can be viewed as special inference rules we can introduce an action corresponding to each axiom of the basis logic, which describes the ability of the agent to use this axiom in her reasoning. By means of the familiar program connectives for dynamic logic (such as composition or iteration) we can formalize the idea that the agent may know the consequences of some sentence which she already knows explicitly, provided that she performs the right reasoning steps. For example, assume that the agent knows the conjunction of p and $p \supset q$, that is, $K(p \wedge (p \supset q))$. There is no guarantee that she will know $p \wedge q$ automatically, as the logic K suggests. We can only say that if the agent reasons correctly then she will know $p \wedge q$. In our concrete case, let CE, CI, MP be the conjunction elimination rule, the conjunction introduction rule, and modus ponens, respectively, and let the symbol ";" denote the composition of actions. Then our theorem must be: $K(p \wedge (p \supset q)) \supset \langle CE; MP; CI \rangle K(p \wedge q)$, and not $K(p \wedge (p \supset q)) \supset K(p \wedge q)$ as in the standard modal approach.

In gerenal, assume that q follows from p in some basis logic (which is accepted by the agent) and that the agent knows p. For explicit knowledge we cannot assume that the agent automatically knows q. Let a proof of q from p be given, where the axioms and inference rules used in the proof are R_1,...,R_n (in this order, where the same axiom or inference rule may occur at different places in the sequence.) Then, instead of the monotonicity rule in the standard modal approach we have the axiom: $Kp \supset \langle R_1; \ldots; R_n \rangle Kq$. This axiom says that if the agent perform the sequence of actions from R_1 to R_n then she may know q under the given circumstances. Whether or not the agent can come to this conclusion depends crucially on her logical ability. In this way we see that the logical omniscience problem can be solved easily in a natural way. On the other hand we can still say that the agent thinks rationally, that she is not logically ignorant.

Often we do not need to care about what course of actions the agent just carried out; we are only interested in the result of the actions, so to speak. We only need to know that the agent has carried out some reasoning steps, and after that she gain certain new information. In this case it is convenient to introduce an auxiliary action F with the following intended reading: do any one of the atomic actions (we don't know what action;) repeat the non-deterministic choice finitely many times (we don't know how many!) The action F could be interpreted as a course of thought. (For dynamic logicians: if the set of all atomic actions associated with the basis logic is a finite set $\{r_1, \ldots, r_n\}$, then F can be viewed as $(r_1 \cup r_2 \cup \ldots \cup r_n)^*$, where the symbols $\cup$ and * denote choice and iteration, respectively.) The choice of the symbols F is not accidental at all: in temporal logic its stands for the operator "Future". It turns out that our auxiliary action behaves in the same manner as the mentioned operator of temporal logic. The formal language in which our dynamic epistemic logics are formulated is called DE and will be defined in the following section.

4 A System of Dynamic Epistemic Logic

Definition 4 (The language DE). We assume the propositional calculus (PC) as our basis logic. Atomic formulae of DE are of the form Kp, where p is a formula of PC. The set of DE-formulae is the least set which contains all atomic formulae and is closed under the usual Boolean connectives and the following condition: if A is a DE-formula then so is $\langle F\rangle A$. Let At denote the set of atomic formulae and Fml be the set of all DE-formulae.

The formula Kp is read: p is known. $\langle F\rangle A$ means "A is true after some course of thought"; $[F]A$ means "A is true after any course of thought". (We could think of $\langle F\rangle$ and $[F]$ as the modalities "at some future times" and "at all future times" of temporal logic.) Note that we consider only sentences about the agent's current and future knowledge such as $K(\neg p)$ or $[F]K(p \vee q)$, but no sentences about "the world" (i.e., sentences of PC) or "mixed" sentences such as $Kp \supset p$. Thus, we must be careful to distinguish between the basis language (i.e., the language of PC) and the language DE of our epistemic logic. Moreover, nested knowledge is not allowed. We do it just for simplicity's sake, but the restrictions are not essential. We use the symbols A, B, ... as meta-variables to denote DE-formulae.

Definition 5 (A minimal logic of growing knowledge). Let $TAUT$ be the set of all tautological PC-formulae. The logic BDE (Basic Dynamic-Epistemic Logic) has the following axiom schemata:

A1. All classical tautologies of the language DE.
A2. $[F](A \supset B) \supset ([F]A \supset [F]B)$
A3. $[F]A \supset [F][F]A$
A4. $Kp \wedge K(p \supset q) \supset \langle F\rangle Kq$
A5. $Kp \supset [F]Kp$
A6. $\langle F\rangle Kp$, where $p \in TAUT$

The rules of inference are:

R1. Modus ponens: if A and $A \supset B$ are theorems then B is a theorem.
R2. Necessation: if A is a theorem then so is $[F]A$

The first three axioms together with the rules R1, R2 axiomatize the minimal temporal logic of transitive time. Axiom A4 says that the agent is capable of using modus ponens. Axiom A5 says that the agent does not forget what she knows. Finally, axiom A6 says that the agent is able to use the theorems of classical logic in her reasoning. (Note that the elements of $TAUT$ belong to the basis language and not to the language DE.) The notions of a proof, a theorem, and a consistent formula or set of formulae are defined as usual. Unless otherwise stated these notions are used relatively to the system BDE.

The following list comprises some provable formulae of BDE. They say that if all premises of a valid inference rule are known or will be known, then after

some steps of reasoning the agent will know the conclusion. (Compare them with the corresponding theorems of the standard approach.) The proof of them is not difficult and is omitted.

Theorem 6. *The following formulae are theorems of BDE:*

1. $Kp \wedge \langle F \rangle K(p \supset q) \supset \langle F \rangle Kq$
2. $Kq \supset \langle F \rangle K(p \supset q)$
3. $Kp \wedge Kq \supset \langle F \rangle K(p \wedge q)$
4. $Kp \wedge \langle F \rangle Kq \supset \langle F \rangle K(p \wedge q)$
5. $K(p \wedge q) \supset \langle F \rangle Kp \wedge \langle F \rangle Kq$
6. $K(p \wedge q) \supset \langle F \rangle (Kp \wedge Kq)$
7. $Kp \supset \langle F \rangle K(p \vee q)$
8. $Kq \supset \langle F \rangle K(p \vee q)$

Instead of the monotony rule in modal epistemic logic we have now the theorem: if $p \supset q$ is a theorem of classical logic, then $Kp \supset \langle F \rangle Kq$ is a theorem of BDE. This is a simple consequence of the axioms of BDE. (Note that this rule is formulated cross the two languages of PC and BDE.) It is obvious that BDE solves the logical omniscience problem. Now we define a semantics for our basic dynamic epistemic logic. Recall that a temporal frame is a structure $\mathcal{F} = (S, R)$ where S is a non-empty set and R is a binary relation on S. We assume that R is transitive. A valuation on the frame $\mathcal{F}$ is a function V from S to the set of subsets of At. A temporal model is a pair $(\mathcal{F}, V)$ comprising a temporal frame and a valuation. Models of the Basic Dynamic Epistemic are temporal models which satisfy certain conditions.

Definition 7 (Models for BDE). A BDE-model M consists of a temporal frame $\mathcal{F}$ and a valuation V which satisfies the following conditions:

1. For all $s \in S$, if $Kp \in V(s)$ and sRt then $Kp \in V(t)$
2. For all $s \in S$, if $Kp, K(p \supset q) \in V(s)$ then $Kq \in V(t)$ for some $t \in S$ with sRt.
3. If $p \in TAUT$ then for all $s \in S$ there is a $t \in S$ such that sRt and $Kp \in V(t)$.

We can interpret $V(s)$ as the agent's explicit knowledge (or her information state) at time s. From the definition we see that if sRt then $V(s) \subseteq V(t)$, that is, the agent's knowledge always grows.

Definition 8. For any temporal model M the relation "the formula A is true at state s in M", in symbol $M, s \models A$, is defined recursively for each $s \in S$ as follows:

1. $M, s \models Kp$ iff $Kp \in V(s)$ for any atomic formula Kp
2. $M, s \models \neg A$ iff $M, s \not\models A$
3. $M, s \models A \supset B$ iff $M, s \models B$ or $M, s \not\models A$
4. $M, s \models \langle F \rangle A$ iff there is some t such that sRt and $M, t \models A$

A formula A is said to be BDE-satisfiable if and only if it is true at some state in some BDE-model. We say that A is true in M just in case A is true at all states in M. A is valid (with respect to BDE-models) iff it is true in all BDE-models.

Theorem 9 (Soundness and Completeness). *A formula A is a theorem of BDE if and only if A is valid.*

Proof. Soundness is straightforward: we have only to show that all BDE-axioms are valid and the application of the inference rules leads from valid premises to valid consequences. Completeness follows immediately from the next lemma. □

Lemma 10. *Every consistent set of formulae is satisfiable.*

Proof. Assume that the set X is consistent. We show that it has a model, i.e., that it is satisfiable. We define a model for X as follows. By the standard Lindenbaum argument, X is contained in a maximal consistent set. Let M be the following structure: $M = ((S, R), V)$ where

- S is the set of maximal consistent sets of BDE-formulae.
- sRt if and only if $\langle F\rangle A \in s$ whenever $A \in t$ (if and only if $A \in t$ whenever $[F]A \in s$).
- $V(s) = \{Kp : Kp \in s\}$ for all atomic formulae Kp.

We show that M is a BDE-model. First, we show that (S, R) is a transitive temporal frame. Assume that sRt and tRu, and let $A \in u$. From tRu we have $\langle F\rangle A \in t$, therefore $\langle F\rangle\langle F\rangle A \in s$, because sRt. By axiom (A3) we can infer $\langle F\rangle A \in s$, hence sRu, i.e., R is transitive.

Next, we prove that for all states s in the temporal model M and all formulae A, $M, s \models A$ if and only if $A \in s$. The base case (A is atomic) is trivial. If A is $\neg B$ then $M, s \models A$ iff $M, s \not\models B$ iff $B \notin s$ (by induction hypothesis) iff $\neg B \in s$ (by maximality of s,) i.e., $A \in s$. If A is another Boolean connection then the argument is similar. Let A be the formula $\langle F\rangle B$. Assume that $M, s \models A$. Then there is some t such that sRt and $M, t \models B$. By induction hypothesis, $B \in t$. By definition of M, $\langle F\rangle B \in s$. Conversely, assume that $A \in s$, i.e., $\langle F\rangle B \in s$. We show that there is some t such that sRt and $M, s \models t$. It suffices to show that the set $Y = \{B\} \cup \{C : [F]C \in s\}$ is consistent. Assume that Y is inconsistent. In this case there must be some $C_1, \ldots, C_n$ such that $[F]C_i \in s$ and $\neg(B \wedge C_1 \wedge \ldots \wedge C_n)$ is provable. It follows that $(C_1 \wedge \ldots \wedge C_n) \supset \neg B$ is provable. By the rule of necessation and the distribution axiom, $([F]C_1 \wedge \ldots \wedge [F]C_n) \supset [F]\neg B$ is provable, so this formula belongs to s. It follows that $[F]\neg B \in s$, contradicting the assumption that $\langle F\rangle B \in s$. Thus, X is consistent. Let t be any maximal consistent extension of X. Then sRt by definition of R, and $B \in t$. By induction hypothesis $M, t \models B$ and so $M, s \models A$.

It remains to show that the temporal model M is a BDE-model. Assume that $Kp \in V(s)$ and sRt. We show that $Kp \in V(t)$. By definition of V we have

$Kp \in s$. As $Kp \supset [F]Kp$ is an axiom it must belong to s, thus $[F]Kp \in s$. It follows that $M, s \models [F]Kp$, hence $M, t \models Kp$, so $Kp \in t$, therefore $Kp \in V(t)$.

Now let $Kp \in V(s)$ and $K(p \supset q) \in V(s)$. We show that $Kq \in V(t)$ for some t such that sRt. By maximality of s all theorems of BDE belong to s, thus $(Kp \wedge K(p \supset q) \supset \langle F \rangle Kq) \in s$. But $V(s) \subseteq s$, so $Kp \in s$ and $K(p \supset q) \in s$. It follows that $\langle F \rangle Kq \in s$, therefore $M, s \models \langle F \rangle Kq$, that is, there must be some t such that sRt and $M, t \models Kq$. From the latter we conclude $Kq \in t$, hence $Kq \in V(t)$.

Finally, we show that if $p \in TAUT$ then for all $s \in S$ there is a $t \in S$ such that sRt and $Kp \in V(t)$. But this is a simple consequence of the fact that $\langle F \rangle Kp \in s$, by axiom A6 and the maximality of s.

Thus, the canonical model is a BDE-model. If X is a consistent set then all formulae of X are satisfied at a state of the canonical model, which is a maximal consistent extension of X, i.e., X has a model. □

5 Conclusions

We have shown how to solve the logical omniscience problem of epistemic logic while preserving the intuition that the agents *are* logical beings. Our strategy consists in taking the dynamic aspect of knowledge into account. We have argued that the correct form of an axiom for epistemic should be: if the agent knows all premisses of a valid inference rule, and if she performs the right reasoning, then she will know the conclusion as well. This intuitive idea has been captured formally within an axiomatic system, and the calculus has been given an adequate semantics. Our strategy can do justice to the intuition that the agent is neither logically omniscient nor logically ignorant. She is non-omniscient, because her actual (or explicit) belief at a single time point need not be closed under any law. It is even possible that she does not believe any logical truth at some of her information states. On the other hand, she is non-ignorant, because she is capable of logical thinking. She can use her reasoning capacities to infer new information from what she already knows. If she performs the correct inferences and if she has enough time, then she might arrive at an ideal information state where all logical consequences of her current beliefs have been drawn. This ideal state can never been achieved if the underlying logic is rich enough, but this is another matter.

On the basis of our basis logic BDE we can develop further epistemic logics which capture other desirable properties of knowledge. It is not hard to write down, for example, axioms saying that the agent is (in principle) capable of introspection and to develop the corresponding model conditions. (In this case we have to extend our language to allow nested belief, of course.) A more interesting problem is to develop dynamic epistemic logics based on non-monotonic logic, where the agent can revise her knowledge when she finds out that her knowledge is inconsistent. We may expect to find interesting connections with two other, very active fields of AI research, viz. to non-monotonic reasoning and to the

logic of belief revision. This seems to be a promising field of research and needs further investigations.

References

1. Chellas, B.: Modal Logic: An Introduction. Cambridge 1980
2. Cresswell, M. J.: Logics and Languages. London 1973
3. Eberle, R.: A Logic of Believing, Knowing and Inferring. Synthese **26** (1974) 356–382
4. Fagin, R., Halpern, J.: Belief, Awareness and Limited Reasoning. Artificial Intelligence **34** (1988) 39–76
5. Harel, D.: Dynamic Logic. In: Handbook of Philosophical Logic, ed. D. Gabbay and F. Guenthner, vol. II, 1984, 497–604
6. Hintikka, J.: Knowledge and Belief. New York 1962
7. Hintikka, J.: Impossible Possible Worlds Vindicated. Journal of Philosophical Logic **4** (1975) 475–484
8. Ho Ngoc Duc: Ein System der epistemischen Logik. In: Philosophie und Logik. Frege-Kolloquien Jena 1989/1991, ed. W. Stelzner, Berlin 1993, 205–214
9. Levesque, Hector: A Logic of Implicit and Explicit Belief. Proceedings AAAI-84, 198–202
10. Rantala, V: Impossible Worlds Semantics and Logical Omniscience. Acta Philosophica Fennica **35** (1982) 18–24
11. Steinacker, Peter: Superschwache Modalkalküle und einige epistemische Anwendungen. Dissertation. Leipzig, 1979
12. Stelzner, Werner: Epistemische Logik. Berlin 1984
13. Wuttich, Klaus: Glauben, Zweifel, Wissen. Eine logisch-philosophische Studie. Berlin 1991

This article was processed using the LaTeX macro package with LLNCS style