JOSEPH Y. HALPERN

SHOULD KNOWLEDGE ENTAIL BELIEF? *

ABSTRACT. The appropriateness of S5 as a logic of knowledge has been attacked at some length in the philosophical literature. Here one particular attack based on the interplay between knowledge and belief is considered: Suppose that knowledge satisfies S5, belief satisfies KD45, and both the *entailment property* (knowledge implies belief) and *positive certainty* (if the agent believes something, she believes she knows it) hold. Then it can be shown that belief reduces to knowledge: it is impossible to have false beliefs. While the entailment property has typically been viewed as perhaps the least controversial of these assumptions, an argument is presented that it can plausibly be viewed as the culprit. More precisely, it is shown that this attack fails if we weaken the entailment property so that it applies only to *objective* (nonmodal) formulas, rather than to arbitrary formulas. Since the standard arguments in favor of the entailment property are typically given only for objective formulas, this observation suggests that care must be taken in applying intuitions that seem reasonable in the case of objective formulas to arbitrary formulas.

KEY WORDS: knowledge, belief, entailment property, logic of knowledge and belief.

1. THE PROBLEM

Ever since philosophers started studying formal models of epistemic logic, there has been discussion as to what is the "right" set of properties for knowledge. In particular, the axioms of S5 have been attacked at some length.[1] My own feeling is that there is no unique "right" notion of knowledge; the appropriate notion is application dependent. Nevertheless, the S5 notion of knowledge has proved to be quite useful (and, indeed, is used routinely) in distributed systems and game theory applications (see (Fagin, Halpern, Moses, and Vardi 1995; Halpern 1987, 1993) for a discussion and overview). Recently, there has been significant interest in logics of knowledge and belief (Friedman and Halpern 1994; van der Hoek 1993; Kraus and Lehmann 1988; Lenzen 1978, 1979; Moses and Shoham 1993; Voorbraak 1991, 1992). In this context, it becomes of interest to reexamine a particular attack against the S5 interpretation of knowledge that seems to be due originally to Lenzen (1978), and has been revived recently by Levesque (private communication, 1992)

* Research sponsored in part by the Air Force Office of Scientific Research (AFSC), under Contract F49620-91-C-0080. The United States Government is authorized to reproduce and distribute reprints for governmental purposes.

Journal of Philosophical Logic **25:** 483–494, 1996.

and Lamarre and Shoham (1994). Roughly speaking, this attack says that if knowledge satisfies S5, belief satisfies KD45, and we make some seemingly reasonable assumptions about the relationship between belief and knowledge, including what Lenzen (1978) has called the *entailment property* (that knowledge implies belief), then an agent cannot hold false beliefs. This is certainly highly counterintuitive!

While the entailment property has typically been viewed as perhaps the least controversial of the assumptions made, I shall argue that it can plausibly be viewed as the culprit here. More precisely, I shall show that Lenzen's attack fails if we weaken the entailment property so that it applies only to *objective* (nonmodal) formulas, rather than to arbitrary formulas. Since the standard arguments in favor of the entailment property (see (Lenzen 1978) for an overview) are typically given only for objective formulas, this observation suggests that care must be taken in applying intuitions that seem reasonable in the case of objective formulas to arbitrary formulas. To understand this issue, it is instructive to examine this attack in more detail. We first need to review some definitions.

There are many interpretations that one could give to knowledge and belief. For the purposes of this argument, I shall give some intuition behind the way I intend these notions to be used (without making any pretense that the intuition completely characterizes these notions). We say that an agent *knows* φ if φ is true in any world consistent with the information he has received. Our notion of belief corresponds to Lenzen's (1978) *strong belief*, Voorbraak's (1992) *rational belief*, and Lamarre and Shoham's (1994) *certainty*. It can be identified with "holding with probability 1" (Halpern 1991). Roughly speaking, an agent believes φ if φ holds at all the worlds in a set to which the agent assigns probability 1; for him also to know φ, φ must be true at all worlds. An alternate reading of belief, used by Moses and Shoham (1993) is that the agent believes φ if he knows that some fixed formula α (which intuitively characterizes the assumptions being made by the agent) implies φ.

Suppose we use the modal operator K for knowledge and B for belief.[2] Given our intuitive reading of belief, it seems reasonable to assume the entailment property: knowledge implies certainty. Formally, this amounts to the axiom

$$K\varphi \Rightarrow B\varphi.$$

Another standard assumption about the relationship between knowledge and belief is captured by the intuition that "to the agent, the facts of which he is certain appear to be knowledge" (Lamarre and Shoham 1994);

for future reference, we call this the *(positive) certainty property*. Formally, it amounts to

$$B\varphi \Rightarrow BK\varphi.$$

If we further assume that B satisfies the axioms of KD45 and that K satisfies the properties of S5, then from the entailment property and positive certainty property we also get the *(negative) certainty property*:

$$\neg B\varphi \Rightarrow B\neg K\varphi.$$[3]

Here are the details of Lenzen's attack: It does not require the negative certainty property, nor does it require that belief satisfies all of KD45. It suffices to assume the entailment property, the positive certainty property, that knowledge satisfies the axioms of S5, and that an agent does not believe both a fact and its negation, i.e., that $\neg(B\varphi \wedge B\neg\varphi)$ is valid. Let p be a fact that the agent is certain is true, but is in fact false; i.e., $\neg p \wedge Bp$ holds. By the certainty property, the agent is certain he knows that p; i.e., BKp holds. But since p is false, the agent cannot know p. By the introspective properties of knowledge (according to S5), the agent knows that he does not know p; i.e., $K\neg Kp$ holds. By the entailment property, he is certain that he does not know p; i.e., $B\neg Kp$ holds. Thus, the agent is certain both that he knows p and that he does not know p. This is a contradiction.

There are several ways of dealing with this attack:

- We can drop the assumption that knowledge satisfies S5. This is the approach taken by Lenzen (1978, 1979) and Lamarre and Shoham (1994), and one of the approaches considered by Voorbraak (1991, 1992) and van der Hoek (1993). The focus has tended to be on logics that are between S4.2 and S4.4 in strength.[4]
- We can drop the positive certainty property. This is the approach taken by Kraus and Lehmann (1988), Moses and Shoham (1993), and Friedman and Halpern (1994).
- We can drop or weaken the entailment property, while keeping both positive certainty and S5 for knowledge. The logics OK&RIB and OK&ORIB of (Voorbraak 1991) do this.

Which is the best approach? Again, my feeling is that there is not necessarily one best approach; the answer is application-dependent. My goal here is not to argue in favor of any particular approach, but to understand to what extent we have to give up the entailment property in order to maintain S5 knowledge and positive certainty.

This is not the place for a general discussion of the entailment property. Lenzen (1978) presents a spirited defence of it, and I have no particular quarrel with any of his arguments. But it must be stressed that all his arguments apply only to objective formulas. At the point in the paper where he makes them, he has not yet started to consider non-objective formulas. My claim is that the intuitions on which Lenzen's arguments are based become much shakier when we start talking about *subjective* formulas. Of course, if we define knowledge as true, justified belief (or, more generally, as belief + truth + some missing ingredient), then the entailment property follows trivially. However, while this view has some intuitive appeal, its justification is again almost always given in terms of objective formulas. It is far from clear what notions like justification mean once we start talking about subjective formulas. For example, what counts as justification for the belief that I do not know p? Are we entitled to assume that the beliefs in subjective formulas are automatically justified? In general, I would argue that our intuitions for notions like justification are far less grounded in everyday natural language when we try to apply them to modal formulas.

This suggests that while it may be appropriate to define knowledge as something like true, justified belief for objective formulas, there is no reason to suspect that this definition should extend to arbitrary formulas. In particular, there is no obvious reason why a property like

$$K\neg Kp \Rightarrow B\neg Kp$$

should be valid. Certainly none of Lenzen's arguments for the entailment property apply to it. But it is precisely this property that causes problems in the attack described above.

Voorbraak's logics OK&RIB and OK&ORIB both have the entaiment property for *belief formulas*; that is, $K\varphi \Rightarrow B\varphi$ holds in these logics if φ is a Boolean combination of formulas of the form $B\psi$. In these logics, the entailment property does not hold in general for objective formulas. In the remainder of this note, I shall sketch a semantics for knowledge and belief that seems reasonably well motivated, where knowledge satisfies the S5 properties, belief satisfies the KD45 properties, and we have both positive and negative certainty, as well as the entailment property for belief formulas and, most importantly, for objective formulas. Needless to say, the attack does not apply to this logic.

2. A SOLUTION

Formally, I consider a propositional modal logic, with modal operators B and K. Starting with a set Φ of primitive propositions, let $\mathcal{L}$ be the smallest set of formulas containing all the propositions in Φ that is closed under conjunction, negation, and application of the modal operators B and K. A formula is *objective* if it is a Boolean combination of the primitive propositions in Φ.

We want to capture all the properties described above (with entailment restricted to objective formulas). To be consistent with the standard literature on belief, we shall take B to satisfy the axioms of KD45. In addition, the semantics enforces the property that an agent knows what he believes, since this has been argued (for example, by Lenzen (1978)) to be reasonable. Thus, we shall have

$$B\varphi \Rightarrow KB\varphi.$$

Since the properties of KD45 imply that $\neg B\varphi \Leftrightarrow B\neg B\varphi$, it is easy to see that we also get

$$\neg B\varphi \Rightarrow K\neg B\varphi.$$

We take as our starting point a logic that has the full entailment property, as well as having knowledge satisfy S5 and belief satisfy KD45. Techniques for capturing these axioms semantically are well known (see (van der Hoek 1993) for a thorough discussion). Since we are considering the single-agent case, we can provide a particularly simple possible-worlds semantics for this logic. A *KB structure* is a pair $M = (W, W')$, where W is a non-empty set of truth assignments to the primitive propositions, and W' is a non-empty subset of W. Intuitively, W is the set of worlds that are consistent with the agent's information, and W' is a subset of W to which the agent assigns probability 1. Alternatively, using the intuitions suggested by Moses and Shoham (1993), we can think of W' as the subset of W satisfying some assumption α. Given a structure $M = (W, W')$, we define what it means for a formula to be true (or false) of a *situation* (M, w), where $w \in W$. As usual, for a primitive proposition p, we define

$$(M, w) \models p \quad \text{if } p \text{ is true under the truth assignment } w.$$

Negation and conjunction are given their standard interpretations. Knowledge is identified with truth in all worlds in W, while belief is identified with truth in all worlds in W'. Thus,

$$(M, w) \models K\varphi \quad \text{if } (M, w') \models \varphi \quad \text{for all } w' \in W,$$

$$(M, w) \models B\varphi \quad \text{if } (M, w') \models \varphi \quad \text{for all } w' \in W'.$$

To show that this semantics has the claimed properties, let AX_1 be the following collection of axioms and inference rules:

A1. All instances of propositional tautologies
A2. $(K(\varphi \Rightarrow \psi) \wedge K\varphi) \Rightarrow K\psi$
A3. $K\varphi \Rightarrow \varphi$
A4. $K\varphi \Rightarrow KK\varphi$
A5. $\neg K\varphi \Rightarrow K\neg K\varphi$
A6. $(B(\varphi \Rightarrow \psi) \wedge B\varphi) \Rightarrow B\psi$
A7. $\neg B$ (*false*)
A8. $B\varphi \Rightarrow BB\varphi$
A9. $\neg B\varphi \Rightarrow B\neg B\varphi$
A10. $K\varphi \Rightarrow B\varphi$
A11. $B\varphi \Rightarrow KB\varphi$
R1. From φ and $\varphi \Rightarrow \psi$ infer ψ
R2. From φ infer $K\varphi$
R3. From φ infer $B\varphi$.

A1–A5 and R1, R2 form the familiar S5 axiom system; thus our notion of knowledge satisfies S5. A1, A6–A9, R1, and R3 form the KD45 axiom system; thus our notion of belief satisfies KD45. A10 is the entailment property, and A11 says an agent knows what he believes.

Soundness and completeness follow by standard techniques of modal logic (see (van der Hoek 1993) for a proof):

THEOREM 2.1. *AX_1 is a sound and complete axiomatization for this language with respect to KB structures.*

Of course, what this logic does not have is positive certainty. This does not make it bad; indeed, of all the approaches to modeling knowledge and belief, this is the one with which I have the most sympathy. My reasons are pragmatic rather than philosophical; this is the approach that I suspect will be most useful in applications, because it is relatively easy to model situations this way. As I now show, by keeping the same class of structures but slightly changing the semantics of belief, we can obtain a logic that gives us both positive and negative certainty. This comes at a price: we have to give up the full entailment property (for otherwise, by Lenzen's argument knowledge would collapse to belief). On the other hand, we do maintain the entailment property for belief formulas.

The idea is to capture the intuition that an agent believes that his beliefs coincide with his knowledge. There are a number of ways of doing this, but they all amount to enforcing the intuition that the agent believes that there are no worlds of probability 0 (or, under the Moses–Shoham reading, that all worlds satisfy the assumption α). Given $M = (W, W')$, let $M^=$ be the structure (W', W'). Intuitively, $M^=$ is a structure where the agent's beliefs coincide with his knowledge. We define a notion of satisfaction $\models'$ that agrees with $\models$ on all clauses except the one for belief formulas. In this case, we define

$$(M, w) \models' B\varphi \quad \text{if } (M^{=}, w') \models' \varphi \quad \text{for all } w' \in W'.$$

Thus, while the standard approach (using $\models$) would be to use M in computing he agent's beliefs, this approach uses $M^{=}$.

The motivation for using $\models'$ rather than $\models$ is easy to explain: $\models'$ enforces the intuition that the agent believes that his knowledge coincides with his beliefs. Otherwise, $\models$ and $\models'$ are identical. In particular, $\models$ and $\models'$ agree on all formulas where there are no K's in the scope of a B. It follows that the entailment property holds under the $\models'$ semantics for objective formulas. Moreover, it follows that the entailment property holds for all belief formulas, just as in Voorbraak's OK&ORIB. However, it does not hold for all formulas in $\mathcal{L}$; in particular, $K\neg Kp \Rightarrow B\neg Kp$ is not valid. For example, suppose that the only primitive proposition in the language is p. Suppose $w(p) = \mathbf{false}$ and $w'(p) = \mathbf{true}$. Let $W = \{w, w'\}$, let $W' = \{w'\}$, and let $M = (W, W')$. Notice that at world w in structure M, the agent falsely believes p; i.e., $(M, w) \models' \neg p \wedge Bp$. Thus, the attack does not apply to this framework. Not surprisingly perhaps, the formula $K\neg Kp \Rightarrow B\neg Kp$ does not hold at (M, w) either.

Van der Hoek (1993) attempts to provide a thorough analysis of how close we can get to having the entailment property, positive certainty, S5 for knowledge, and KD45 for belief. However, he considers only variants of the "standard" semantics. In particular, he does not consider changing the set of possible worlds inside the scope of a B operator. As we now show, the logic obtained by using the nonstandard semantics comes closer to having the properties we want than any of his logics.

We can characterize the $\models'$ semantics axiomatically quite easily. Let AX_2 consist of all the axioms and rules of AX_1, except that A10 is weakened to

A10$'$. $K\varphi \Rightarrow B\varphi$ if φ is an objective formula

and we have in addition positive and negative certainty:

A12. $B\varphi \Rightarrow BK\varphi$
A13. $\neg B\varphi \Rightarrow B\neg K\varphi$.[5]

Voorbraak's logic OK&ORIB satisfies all these axioms except for A10$'$; OK&RIB satisfies all but A10$'$ and A13. We would obtain the logic OK&ORIB if we dropped the assumption that in a KB structure $M = (W, W')$, we must have $W' \subseteq W$.

THEOREM 2.2. *AX_2 is a sound and complete axiomatization for this language under the $\models'$ semantics with respect to the class of KB structures.*

Proof. Soundness is straightforward, and left to the reader.

For completeness, we first show that every formula is provably equivalent to a *depth-one* formula, where a depth-one formula is one without nested modal operators. Thus, a depth-one formula is a Boolean combination of objective formulas and formulas of the form $B\varphi$ and $K\varphi'$, where φ and φ' are objective.

This follows easily from the following equivalences, all of which are provable in AX_2:

E1. $KK\varphi \Leftrightarrow K\varphi$
E2. $K\neg K\varphi \Leftrightarrow \neg K\varphi$
E3. $KB\varphi \Leftrightarrow B\varphi$
E4. $K\neg B\varphi \Leftrightarrow \neg B\varphi$
E5. $K(\varphi \wedge \psi) \Leftrightarrow (K\varphi \wedge K\psi)$
E6. $K(K\varphi \vee \psi) \Leftrightarrow (K\varphi \vee K\psi)$
E7. $K(\neg K\varphi \vee \psi) \Leftrightarrow (\neg K\varphi \vee K\psi)$
E8. $K(B\varphi \vee \psi) \Leftrightarrow (B\varphi \vee K\psi)$
E9. $K(\neg B\varphi \vee \psi) \Leftrightarrow (\neg B\varphi \vee K\psi)$
E10. $BB\varphi \Leftrightarrow B\varphi$
E11. $B\neg B\varphi \Leftrightarrow \neg B\varphi$
E12. $BK\varphi \Leftrightarrow B\varphi$
E13. $B\neg K\varphi \Leftrightarrow \neg B\varphi$
E14. $B(\varphi \wedge \psi) \Leftrightarrow (B\varphi \wedge B\psi)$
E15. $B(B\varphi \vee \psi) \Leftrightarrow (B\varphi \vee B\psi)$
E16. $B(\neg B\varphi \vee \psi) \Leftrightarrow (\neg B\varphi \vee B\psi)$
E17. $B(K\varphi \vee \psi) \Leftrightarrow (B\varphi \vee B\psi)$
E18. $B(\neg K\varphi \vee \psi) \Leftrightarrow (\neg B\varphi \vee B\psi)$.

Since all these proofs are straightforward (and have a similar flavor), I shall just sketch the proof of E17, leaving the remaining details to the reader.

Using $\vdash$ to denote provability in AX_2, since $\vdash B\varphi \Rightarrow BK\varphi$ by A12, we clearly have

$$\vdash B\varphi \Rightarrow B(K\varphi \vee \psi)$$

and

$$\vdash B\psi \Rightarrow B(K\varphi \vee \psi).$$

Hence, by propositional reasoning (A2 and R1)

$$\vdash (B\varphi \vee B\psi) \Rightarrow B(K\varphi \vee \psi).$$

For the other direction, since $\vdash \neg B\varphi \Rightarrow B\neg K\varphi$ by A13, we have

$$(1) \qquad \vdash (B(K\varphi \vee \psi) \wedge \neg B\varphi) \Rightarrow (B(K\varphi \vee \psi) \wedge B\neg K\varphi).$$

Now by E14,

$$(2) \qquad \vdash (B(K\varphi \vee \psi) \wedge B\neg K\varphi) \Rightarrow B((K\varphi \vee \psi) \wedge \neg K\varphi).$$

By standard KD45 reasoning,

$$(3) \qquad \vdash B((K\varphi \vee \psi) \wedge \neg K\varphi) \Rightarrow B\psi.$$

From (1), (2), and (3), we get

$$\vdash (B(K\varphi \vee \psi) \wedge \neg B\varphi) \Rightarrow B\psi.$$

Using standard propositional reasoning, we can conclude

$$\vdash B(K\varphi \vee \psi) \Rightarrow (B\varphi \vee B\psi).$$

This gives us E17.

Now using standard propositional reasoning, E5, E14, and the fact that *Ktrue* and *Btrue* are provable (by R2 and R3), it follows that any formula φ is provably equivalent to a disjunction of conjunctions of the form

$$\begin{aligned}\varphi_0 \ &\wedge\ K\varphi_1 \wedge \neg K\varphi_{21} \wedge \cdots \wedge \neg K\varphi_{2m}\\ &\wedge\ B\varphi_3 \wedge \neg B\varphi_{41} \wedge \cdots \wedge \neg B\varphi_{4n},\end{aligned}$$

where each of the formulas φ_i or φ_{ij} is objective.

As usual, to show completeness, it suffices to show that if φ is consistent with AX, then it is satisfiable in some KB structure. By the above, it suffices to assume that φ is a conjunction of the form above. I now show how to construct a KB structure satisfying such a consistent φ. Let W consist of all truth assignments satisfying φ_1, let W' consist of all truth assignments satisfying $\varphi_1 \wedge \varphi_3$, and let $M = (W, W')$. There must be some truth assignment w satisfying φ_0, for otherwise $\varphi_0 \wedge \varphi_1$ is inconsistent, contradicting A3. I claim that $(M, w) \models' \varphi$.

By construction, $(M, w) \models' \varphi_0 \wedge K\varphi_1 \wedge B\varphi_3$. It remains to show that $(M, w) \models' \neg K\varphi_{2j}$, $j = 1, \dots, m$ and $(M, w) \models' \neg B\varphi_{4k}$, $k = 1, \dots, n$. If $(M, w) \not\models' \neg K\varphi_{2j}$, then there is no truth assignment satisfying $\varphi_1 \wedge \neg\varphi_{2j}$. Thus, $\varphi_1 \Rightarrow \varphi_{2j}$ is a propositional tautology, and by R2, it follows that $\vdash K(\varphi_1 \Rightarrow \varphi_{2j})$. Using A2, we get that $\vdash K\varphi_1 \Rightarrow K\varphi_{2j}$. But this contradicts the consistency of $K\varphi_1 \wedge \neg K\varphi_{2j}$, and hence of φ.

If $(M, w) \not\models' \neg B\varphi_{4k}$, then there is no truth assignment satisfying $\varphi_1 \wedge \varphi_3 \wedge \neg\varphi_{4k}$. Thus, $(\varphi_1 \wedge \varphi_3) \Rightarrow \varphi_{4k}$ is a propositional tautology.

Using R3 and B6, arguments similar to those above show that $\vdash B(\varphi_1 \wedge \varphi_3) \Rightarrow B\varphi_{4k}$. But, using A10′, we can easily show $\vdash (K\varphi_1 \wedge B\varphi_3) \Rightarrow B(\varphi_1 \wedge \varphi_3)$. This contradicts the consistency of $K\varphi_1 \wedge B\varphi_3 \wedge \neg B\varphi_{4k}$, and hence φ.

This completes the proof of completeness. □

Using standard arguments due originally to Ladner (1977), it can also be shown that the satisfiability problem for this logic is NP-complete, no worse than that of propositional logic. I shall not pursue this point here.

3. CONCLUSION

My goal here was not to present yet another logic of knowledge and belief and argue for its superiority, but rather to examine Lenzen's attack and to understand better the reasonableness of the assumptions it requires. As I have shown, the attack depends on a strong form of the entailment property. Is this strong form justified, that is, should knowledge entail belief? That depends. Words like "knowledge" and "belief" admit more than one interpretation and, as I said earlier, I do not think there is a unique right one. Rather than asking "Should knowledge entail belief?", we should ask whether a particular model of knowledge and belief is both useful and true to the spirit of the way the words are used in a particular class of applications. The entailment property should certainly not serve as a defining test for the reasonableness of the model.

In any case, what I do hope my arguments have shown is that Lenzen's attack does not necessarily rule out S5 as an appropriate logic of knowledge. More importantly, they suggest that we must be careful when trying to extend to subjective formulas intuitions – such as those that motivate the entailment property – that may seem reasonable when applied to objective formulas.

ACKNOWLEDGEMENTS

I would like to thank Gerhard Lakemeyer, Yoram Moses, and Yoav Shoham for useful discussions on the subject of this paper, and an anonymous referee for useful comments.

NOTES

[1] The systems S5 and KD45 are formally described in the next section. For the purposes of this introduction, it suffices to know that S5 presumes that the agent's knowledge

is closed under logical implication, the agent has perfect introspection, and that the agent knows only true propositions. KD45 is similar, except it weakens the requirement that the agent knows only true propositions to the requirement that the agent not hold inconsistent beliefs.

[2] Lenzen (1978) and Lammarre and Shoham (1994) actually use C rather than B for the notion we are considering, and reserve B for a weaker notion.

[3] Proof sketch: from $\neg B\varphi$, we get $B\neg B\varphi$ by negative introspection. The positive certainty property then allows us to conclude $BK\neg B\varphi$. From the entailment property we can conclude $\neg B\varphi \Rightarrow \neg K\varphi$, which gives us $BK\neg K\varphi$. Using standard S5 reasoning, $K\neg K\varphi$ is equivalent to $\neg K\varphi$. This gives us $B\neg K\varphi$, as desired.

[4] To be fair, the authors of these papers present positive reasons for moving away from S5, and do not view the move as simply a way to avoid Lenzen's attack.

[5] Although, as shown above, the negative certainty property follows from the positive certainty property, the KD45 axioms for belief, the S5 axioms for knowledge, and the full-fledged entailment property, it does not follow given the weaker version of the entailment property that we have here, so we take it as an axiom.

REFERENCES

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

Friedman, N. and Halpern, J.Y. (1994). A knowledge-based framework for belief change. Part I: Foundations, in R. Fagin (ed.), *Theoretical Aspects of Reasoning about Knowledge: Proc. Fifth Conference*, pp. 44–64. San Francisco, Calif., Morgan Kaufmann.

Halpern, J.Y. (1987). Using reasoning about knowledge to analyze distributed systems, in J. F. Traub, B. J. Grosz, B. W. Lampson, and N. J. Nilsson (eds), *Annual Review of Computer Science*, Vol. 2, pp. 37–68. Palo Alto, Calif., Annual Reviews Inc.

Halpern, J.Y. (1991). The relationship between knowledge, belief, and certainty, *Annals of Mathematics and Artificial Intelligence* **4**, 301–322.

Halpern, J.Y. (1993). Reasoning about knowledge: a survey circa 1991, in A. Kent and J. G. Williams (eds), *Encyclopedia of Computer Science and Technology*, Vol. 27 (Supp. 12), pp. 275–296. New York, Marcel Dekker.

Hoek, W. van der (1993). Systems for knowledge and belief, *Journal of Logic and Computation* **3**(2), 173–195.

Kraus, S. and Lehmann, D.J. (1988). Knowledge, belief, and time, *Theoretical Computer Science* **58**, 155–174.

Ladner, R.E. (1977). The computational complexity of provability in systems of modal propositional logic, *SIAM Journal on Computing* **6**(3), 467–480.

Lamarre, P. and Shoham, Y. (1994). Knowledge, certainty, belief, and conditionalisation, in J. Doyle, E. Sandewall, and P. Torasso (eds), *Principles of Knowledge Representation and Reasoning: Proc. Fourth International Conference (KR '94)*, pp. 415–424. San Francisco, Calif., Morgan Kaufmann.

Lenzen, W. (1978). Recent work in epistemic logic, *Acta Philosophica Fennica* **30**, 1–219.

Lenzen, W. (1979). Epistemische betrachtungen zu [S4, S5]. *Erkenntnis* **14**, 33–56.

Moses, Y. and Shoham, Y. (1993). Belief as defeasible knowledge, *Artificial Intelligence* **64**(2), 299–322.

Voorbraak, F. (1991). The theory of objective knowledge and rational belief, in *Logics in AI, European Workshop JELIA '90*, Berlin–New York, pp. 499–515. Springer-Verlag.

Voorbraak, F. (1992). Generalized Kripke models for epistemic logic, in Y. O. Moses (ed.), *Theoretical Aspects of Reasoning about Knowledge: Proc. Fourth Conference*, pp. 214–228. San Francisco, Calif., Morgan Kaufmann.

IBM Research Division,
Almaden Research Center, Dept. K53/802,
San Jose, CA 95120, U.S.A.
e-mail: halpern@almaden.ibm.com