December 1988

Report No. STAN-CS-88-1237

Belief as Defeasible Knowledge

by

Yoav Shoham and Yoram Moses

Department of Computer Science

Stanford University
Stanford, California 94305

Belief as Defeasible Knowledge

Yoav Shoham ✝
Computer Science Department
Stanford University
Stanford, CA 94305

and

Yoram Moses
Department of Applied Mathematics
The Weizmann Institute of Science
Rehovot, Israel 76100

Abstract

We investigate the relation between the notions of knowledge and belief. Contrary to the well-known slogan about knowledge being "justified, true belief," we propose that belief be viewed as defeasible knowledge. Specifically, we offer a definition of belief as knowledge-relative-to- assumptions, and tie the definition to the notion of nonmonotonicity. Our definition has several advantages. First, it is short. Second, we do not need to add anything to the logic of knowledge: the right properties of belief fall out of the definition and the properties of knowledge. Third, the connection between knowledge and belief is derived from one fundamental principle, which is more enlightening than a collection of arbitrary-seeming axioms relating the two

1 Introduction

In recent years there has been much interest in formal reasoning about knowledge and belief, in both the AI community and the distributed computation community. Motivated by the need to represent the knowledge of intelligent agents, and drawing on Hintikka's work in philosophy [5], Moore was the first to introduce the logic of knowledge to AI [9]. Further work in AI on formal modeling of agents' epistemic states includes [6, 8]. An introduction to modal epistemic logic in AI can be found in [2]. Motivated by the need to represent the local knowledge of individual processors, the distributed computation community too has taken great interest in the topic. There has been a tremendous amount of work there, best represented by the proceedings of the two conferences devoted to the subject [4, 12]. This is by no means an exhaustive list of references, either in AI or in distributed computation.

Our aim in this paper is to clarify the relation between knowledge and belief, two epistemic notions that play an important role in AI. The old slogan is that "knowledge is justified, true belief," suggesting that belief should be taken as basic, and knowledge defined in terms of it. In fact, we know of few attempts to capture the two notions in the' same formal system, and none which take this particular tack.

One could imagine a second way of combining the notions of knowledge and belief, which starts by defining them each separately. It has become standard to define epistemic notions through Kripke models, (although see [1] for an alternative to Kripke semantics). In the case of knowledge, one usually uses the simple S5 system (but see, e.g., [11] for arguments against the S5 system and for the S4 system). This yields intuitive properties for knowledge: if I know φ then φ is true, if I know φ then I know that I know it, and (in the case of S5) if I don't know φ then I know that I don't know it. As is well known, these three properties are achieved by the axioms $K\varphi \supset \varphi$, $K\varphi \supset KK\varphi$, and $\neg K\varphi \supset K\neg K\varphi$, respectively (or, equivalently, by the requirement that the accessibility relation on possible worlds be reflexive, transitive and Euclidean). In addition one has the 'normality' condition, $K(p \supset q) \wedge Kp \supset Kq$: An agent knows the tautological consequences of his knowledge.

Of these properties, the first one is clearly inappropriate for belief: I may believe something false. Indeed, since this seems to be the property distinguishing knowledge from belief, the common logic for belief is so-called *weak S5*, or KD45, in which the $K\varphi \supset \varphi$ axiom (equivalently, the reflexivity requirement) is omitted.

One could thus imagine combining knowledge and belief as follows: create two modalities, one (say) regular S5, the other (say) weak S5, and write enough axioms relating them to one another. Indeed, we know of at least one such attempt [7]. Although reasonable, this approach suffers a serious disadvantage: There is no theoretical basis for those added axioms, and therefore we have no guarantee that we have indeed captured the full connection between the two notions.

In this paper we offer a simple alternative: start with only a definition of knowledge, any definition that you find acceptable, and define belief as a defeasible version of it.

Roughly speaking, we will translate each occurrence of "the agent believes that φ" into "the agent knows that either φ, or else something specific unusual is the case.', For example, if the robot's vision system reports an obstacle then the robot believes that an obstacle exists, since it knows that either the obstacle indeed exists, or else its vision system is malfunctioning (the latter considered unusual). The full definition we will adopt is only slightly more complex.

This definition of belief has several advantages. First, it is short. Second, we do not need to add anything to the logic of knowledge: the right properties of belief fall out of the definition and the properties of knowledge. Third, not only do we get a connection between knowledge and belief, but we get it from one fundamental principle, in a way that is more enlightening than a collection of arbitrary-seeming axioms relating the two notions. Finally, and most surprisingly, we note an added benefit of our definition: it suggests a close connection between the notion of belief and, of all things, nonmonotonic reasoning.

The paper is organized as follows. In Section 2 we give our definition of belief. In fact, we will give two such definitions, of which we will adopt one. Both are a general reduction of the notion of belief to that of knowledge, and do not assume any particular definition of knowledge. In Section 3 we explore a particular instance of our definition, the one in which knowledge is taken to be defined by the S5 system. In Section 4 we tie the discussion to the notion of nonmonotonicity. We end in Section 5 with some concluding remarks and discussion of related work.

2 Defining belief

As was said in the introduction, we will reduce the notion of belief to that of knowledge. We will not assume anything about the definition of knowledge, only that we have a language for describing the world, and that if φ is a wff in that language then so is $K\varphi$, meaning "the agent knows that φ" (we restrict the discussion in this article to a single agent, but the extension to multiple agents is straightforward).

Each formula that is believed will be believed only with respect to some other "assumption" formula. We start with a definition that is close to the one we will actually adopt. This initial definition is a direct translation of the sentence given in the introduction: we say that φ is believed just in case it is known that either φ holds or else that the assumption is violated.

Definition 1 $B'(\varphi, \varphi_{ass}) =_{def} K(\varphi_{ass} \supset \varphi)$

We will see later that this simple definition actually has some attractive properties. However, it also has some properties that, under certain circumstances, we might not find acceptable. In particular, most logics of knowledge allow us to derive $K\neg\varphi_{ass} \supset B'(\varphi, \varphi_{ass})$: whenever we know that our assumption is violated, we must believe φ.

This motivates our final definition of belief. In the following definition, we add the condition that beliefs cannot be grounded in assumptions that are actually known to be violated. Specifically, we require that a formula be believed relative to an assumption that is known not to hold, only if that formulamisexplicitly known.

Definition 2 $B(\varphi, \varphi_{ass}) =_{def} K(\varphi_{ass} \supset \varphi) \wedge (K\neg\varphi_{ass} \supset K\varphi)$

We first note the following easily-seen connection between the two definitions of belief:

Proposition 1 $\neg K\neg\varphi_{ass} \supset (B'(\varphi, \varphi_{ass}) \equiv B(\varphi, \varphi_{ass}))$

Indeed, as we explore the ramifications of our definition of B in the next section, we will also see that the difference between ***B*** and ***B'*** hinges on the possibility of knowing the negation of one's assumptions.

3 Properties of belief

Our definition of belief in the previous section was justified on intuitive grounds, if at all. Indeed, it was this intuition alone that originally led us to the definition. We now put it to a test by verifying that it has formal consequences that make sense. We have not discovered any undesirable consequences of our definition (but see discussion in the summary section). On the other hand, many desirable properties do follow from them. These include all the properties of belief that we have seen in previous formalisms, and some new ones.

In order to present crisp results, we will explore the ramification of our definitions in the context of a particular logic of knowledge, the S5 system described in the introduction. As was said, S5 is the most popular system for defining knowledge. Although it is very well known, for completeness we repeat the definition of (single-agent, propositional) ***S5*** here. (Throughout this article we will restrict the discussion to the single-agent, propositional case, but it will be apparent that the discussion extends easily to the multiple-agent, first-order case.)

Definition 3
Syntax. ***Given a set of primitive propositions P, the formulas of the logic consist of the members of P, their boolean closure (closure under*** $\wedge$ ***and*** $\neg$***), and their closure under the modal operator K.***
Axiom system. ***Beside the axiom schemas of propsotional calculus, the S5 axiom schemas consist of*** $(K(\varphi_1 \supset \varphi_2) \wedge K\varphi_1) \supset K\varphi_2$, $K\varphi \supset \varphi$, $K\varphi \supset KK\varphi$, ***and*** $\neg K\varphi \supset K\neg K\varphi$. ***The rules of inference are*** modus ponens ***and*** generalization: ***from*** φ ***infer*** $K\varphi$.
Semantics. ***S5 Kripke Structures are pairs (M, w), where M is a set of (total) valuations of P, and w is a member of M. The notion of a formula being satisfied in a Kripke***

structure is defined as follows. For primitive propositions p, $\langle M, w\rangle \models p$ iff $w \models p$, where the latter $\models$ denotes standard propositional satisfaction. $\langle M, w\rangle \models \varphi_1 \wedge \varphi_2$ iff $\langle M, w\rangle \models \varphi_1$ and $\langle M, w\rangle \models \varphi_2$. $\langle M, w\rangle \models \neg\varphi$ iff it is not the case that $\langle M, w\rangle \models \varphi$. Finally, $\langle M, w\rangle \models K\varphi$ (or simply $M \models K\varphi$) iff $\langle M, w'\rangle \models \varphi$ for all $w' \in M$.

As has been discussed extensively in the literature, K can be viewed as the knowledge operator, and thus $K\varphi$ is read "φ is known" (or, in the general case, $K_i\varphi$ is read "agent i knows φ").

<u>Notational convention.</u> When φ_{ass} can be inferred by context, or when it is not important what the particular assumption is, we will replace $B(\varphi, \varphi_{ass})$ by $B(\varphi)$, or simply $B\varphi$. We will do the same for B'. When several belief operators occur in the same sentence with omitted assumption arguments, we assume that all the assumption formulas are the same. For example, $Bp \wedge Bq$ stands for $B(p, r) \wedge B(q, r)$, where r is either understood from context, or else is any arbitrary formula.

We will concentrate on the properties of B, although at some points, when some properties of B' are illuminating, we will make reference to B' too. We start by noting an equivalent definition of B. The definition reads "φ is believed iff it is known that either φ is indeed true, or else the assumption is violated without this violation being known":

Proposition 2 $B(\varphi, \varphi_{ass}) \equiv K(\varphi \vee (\neg\varphi_{ass} \wedge \neg K \neg\varphi_{ass}))$

Proof.

$B(\varphi, \varphi_{ass}) \equiv$
$K(\varphi_{ass} \supset \varphi) \wedge (K\neg\varphi_{ass} \supset K\varphi) \equiv$
$K(\varphi \vee \neg\varphi_{ass}) \wedge (K\varphi \vee \neg K\neg\varphi_{ass}) \equiv$
$K(\varphi \vee \neg\varphi_{ass}) \wedge K(\varphi \vee \neg K\neg\varphi_{ass}) \equiv$
$K((\varphi \vee \neg\varphi_{ass}) \wedge (\varphi \vee \neg K\neg\varphi_{ass})) \equiv$
$K(\varphi \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass}))$ ■

It is also **easy** to see that one cannot believe in contradictory statements (and recall the notational convention of suppressing the assumption argument):

Proposition 3 $\neg(B\varphi \wedge B\neg\varphi)$

Proof.

$B\varphi \wedge B\neg\varphi \equiv$
$K(\varphi \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \wedge K(\neg\varphi \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$K((\varphi \wedge \neg\varphi) \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$K(\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass}) \equiv$
$K\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass}) \equiv$
$false$ ■

Interestingly, for B' we get the following:

Proposition 4 $B'(\varphi, \varphi_{ass}) \wedge B'(\neg\varphi, \varphi_{ass}) \equiv K \neg\varphi_{ass}$

Proof.
$B'(\varphi, \varphi_{ass}) \wedge B'(\neg\varphi, \varphi_{ass}) \equiv$
$K(\varphi \vee \neg\varphi_{ass}) \wedge K(\neg\varphi \vee \neg\varphi_{ass}) \equiv$
$K((\varphi \vee \neg\varphi) \wedge \neg\varphi_{ass}) \equiv$
$K\neg\varphi_{ass}$ ■

We now begin to explore the connection between knowledge and belief. The first connection is obvious:

Proposition 5 $K\varphi \supset B\varphi$

The converse implication of course does not hold. Note that this proposition is true of B' too. Note also that the validity of the generalization inference rule, from φ infer $B\varphi$, follows from the last proposition.

Proposition 6 $BK\varphi \equiv K\varphi$

. ***Proof***
$BK\varphi \equiv$
$K(K\varphi \vee(\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$K\varphi \vee K(\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$K\varphi \vee \mathit{false} \equiv$
$K\varphi$ *a*

The transitivity property of belief, or the property of "positive introspection," follows immediately:

Corollary 7 $BB\varphi \equiv B\varphi$

Proposition 8 $B\neg K\varphi \equiv \neg K\varphi$

Proof
$B\neg K\varphi \equiv$
$K\neg(K\varphi \vee(\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$\neg K\varphi \vee K(\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$\neg K\varphi \vee \mathit{false} \equiv$
$\neg K\varphi$ ■

As a corollary we get the "negative introspection" property:

Corollary 9 $B\neg B\varphi \equiv \neg B\varphi$

Recall that knowledge was defined essentially by the four axioms of S5. The last two corollaries show that two of the axioms hold for belief too. Another axiom, that of distributivity, also holds:

Proposition 10 $B(\varphi_1 \supset \varphi_2) \wedge B\varphi_1 \supset B\varphi_2$

Proof.
$B(\varphi_1 \supset \varphi_2) \wedge B\varphi_1 \equiv$
$K((\varphi_1 \supset \varphi_2) \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \wedge K(\varphi_1 \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$K(((\varphi_1 \supset \varphi_2) \wedge \varphi_1) \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \supset$
$K(\varphi_2 \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$B\varphi_2$ ■

Thus, through our one definition of ***B***, ***we*** have obtained a *weak-S5* logic of belief!

The one remaining axiom $S5$, $K\varphi \supset \varphi$, we would not want belief to have, and indeed it does not. We do, however, have that if our assumptions hold, then indeed our beliefs are correct:

Proposition 11 $B(\varphi, \varphi_{ass}) \wedge \varphi_{ass} \supset \varphi$

This is true of ***B'*** too, and in both cases the proof is immediate.

We now continue to explore the connection between ***B*** and ***K***. We have seen that ***BK*** and $B\neg K$ collapse to ***K*** and $\neg K$, respectively. We now show that ***KB*** and $K\neg B$ collapse to ***B*** and $\neg B$, respectively.

Proposition 12 $KB\varphi \equiv B\varphi$

Proof.
$KB\varphi \equiv$
$KK(\varphi \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$K(\varphi \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$B\varphi$ ■

Proposition 13 $K\neg B\varphi \equiv \neg B\varphi$

Proof.
$K\neg B\varphi \equiv$
$K\neg K(\varphi \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$\neg K(\varphi \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$
$\neg B\varphi$ ■

4 Knowledge, belief and nonmonotonic reasoning

So far we have paid little attention to the second argument to the belief operator. It turns out to be instructive to look at the assumptions more closely.

We first note the "belief weakening" property (some of the easy proofs below are omitted):

Proposition 14 $B(p, q) \wedge K(p \supset r) \supset B(r, q)$

This property holds also for B'. One might have expected an analogous "assumption strengthening" property. For B' this expectation is met:

Proposition 15 $K(q \supset r) \wedge B'(p, r) \supset B'(p, q)$

The same fact, however, does not hold for B. Indeed, the way to understand the role of the assumption in B is as an assumption in nonmonotonic logics. On the basis of certain assumptions we are willing to adopt certain beliefs, but given more evidence we may discard some of them.

This point is further made when we examine the conditions under which one can believe in, or even know, the very assumption in which the belief is grounded. In the following we assume a specific implicit assumption formula φ_{ass}. Thus, $B\varphi$ stands for $B(\varphi, \varphi_{ass})$. In particular, $B\varphi_{ass}$ stands for $B(\varphi_{ass}, \varphi_{ass})$.

The B' operator holds few surprises in this regard. We have the following two easy facts:

Proposition 16 $B'\varphi_{ass}$

Proposition 17 $B'\neg\varphi_{ass} \equiv K\neg\varphi_{ass}$

Thus, according to B', we always believe in the truth of our assumption, and the only time we believe in the negation of our assumption is also the only time beliefs can become inconsistent – when we actually know the assumption to be violated.

As far as believing the negation of the assumption, the B operator behaves identically:

Prbposition 18 $B\neg\varphi_{ass} \equiv K\neg\varphi_{ass}$

Proof

$B\neg\varphi_{ass} \equiv$

$K(\neg\varphi_{ass} \vee (\neg\varphi_{ass} \wedge \neg K\neg\varphi_{ass})) \equiv$

$K\neg\varphi_{ass}$ ∎

When it comes to believing that the assumption does hold, however, B behaves quite differently from B'. Obviously, if we know that our assumption holds then we also believe that it does (knowledge entails belief), but in fact we believe in the assumption even under weaker conditions:

Proposition 19 $B\varphi_{ass} \equiv \neg K \neg \varphi_{ass}$

Proof.
$B\varphi_{ass} \equiv$
$K(\varphi_{ass} \vee (\neg\varphi_{ass} \wedge \neg K \neg \varphi_{ass})) \equiv$
$K(\varphi_{ass} \vee \neg K \neg \varphi_{ass}) \equiv$
$K(\neg K \neg \varphi_{ass}) \equiv$
$\neg K \neg \varphi_{ass}$ ∎

Now, this is an eminently nonmonotonic inference: as long as you don't know your assumption to be false, believe it to be true. Notice, however, that our logic is entirely monotonic.[1]

As a corollary of the last two propositions, we get that we always take a stance towards our assumption; we either know that it does not hold, or we believe that it does:

Corollary 20 $B\varphi_{ass} \vee K \neg \varphi_{ass}$

Of course, this last property does not hold for arbitrary φ's. Nor can we replace the ***B*** by a ***K***: we do not necessarily have complete knowledge about the truth of our assumptions. Intuitively speaking, if we know whether our assumption is true or false, it is no longer an assumption. Indeed, we have the following property:

Proposition 21 $(B\varphi \wedge (K\varphi_{ass} \vee K\neg\varphi_{ass})) \supset K\varphi$

Proof.
$B\varphi \wedge (K\varphi_{ass} \vee K \neg \varphi_{ass}) \equiv$
$B\varphi \wedge K(\varphi_{ass} \vee K \neg \varphi_{ass}) \equiv$
$B\varphi \wedge K\neg(\neg\varphi_{ass} \wedge \neg K \neg \varphi_{ass}) \equiv$
$K(\varphi \vee (\neg\varphi_{ass} \wedge \neg K \neg \varphi_{ass})) \wedge K\neg(\neg\varphi_{ass} \wedge \neg K \neg \varphi_{ass}) \supset$
$K\varphi$ ∎

5 Summary and Discussion

We have offered a definition of belief as a defeasible form of knowledge, or as knowledge-relative-to-assumptions. The definition is short, and, we find, intuitive, deriving the connection between knowledge and belief from one basic principle. The definition, which was based on intuitive understanding of knowledge and belief, turns out to be very robust. We have shown it to have formal consequences that one would expect from the concept of belief. We have also shown that the notion of belief strongly exhibits properties of nonmonotonic reasoning.

[1] Who was it that once said: "nonmonotonicity - si, nonmonotonic logics - *no*!"

Our notion of belief originates from a definition given by the second author in [10] for belief in the context of distributed computation. There the definition for -processor i believing in φ was
$K_i((i \in \mathcal{N}) \supset \varphi)$, where $\mathcal{N}$ is the set of nonfaulty processors. This is clearly a special case of our $\boldsymbol{B}$. Otherwise, we know of little previous work on incorporating knowledge and belief within the same framework, and none which takes our approach. The most closely related work of which we are aware is by Kraus and Lehmann [7], who indeed introduce two modalities and relate them through axioms. Their axioms turn out to be a proper subset of the propositions proved in this paper. In addition, Halpern has recently proposed a probabilistic account of both knowledge and belief, in which knowledge is equated with certainty and belief with "almost certainty" [3]. Although on the one hand Halpern's account does not explicate the assumption underlying belief, and on the other hand we do not explain the connection between statistical information and belief, there appears to be full compatibility between the two accounts.

Finally, we note that our definition of belief does not rely on assuming a particular logic of knowledge. In this article we adopted the S5 system for two reasons: it is by far the most widely adopted system, and we needed some system in order to present crisp results.. There are those who object to an S5 definition of knowledge. Some objections are mild, for example to the axiom of negative introspection $(\neg K\varphi \supset K\neg K\varphi)$. Other objections are more extreme, for example to the assumption of deductive closure (the axiom ($K(\varphi_1 \supset \varphi_2) \wedge K\varphi_1) \supset K\varphi_2)$. However, our definition of $B(\varphi, \varphi_{ass}) =_{def} K(\varphi_{ass} \supset \varphi) \wedge (K\neg\varphi_{ass} \supset \boldsymbol{K}\ \varphi)$ holds regardless of the meaning assigned to the $\boldsymbol{K}$ operator. For example, one could drop the axiom of negative introspection (and thus adopt the S4 system), and remain with our definition of belief. Or, one might adopt the notion of resource-bounded knowledge in [10]. Of course, the specific properties of belief will change as we assume different logics of knowledge.

Indeed, several fascinating issues remain to be investigated:

- It will be interesting to explore the properties of belief that result from other notions of knowledge.

- We are interested in repercussions of our definition on notions such ***as common belief.*** It can be shown to be be both intuitive and useful to base the notion of common belief on the definition of belief as defeasible knowledge.

- We have not given completeness results for the dervied notion of belief. For example, we have not shown that if we start with an S5 logic of knowledge, then the resulting notion of belief is complete for, e.g, weak S5. We believe that such a result will be relatively easy to attain.

- All formal theories of commonsense leave open pragmatic questions. For example, the theory of circumscription does not tell us which predicates to minimize and vary. Similarly, we have not presented guidelines for choosing which assumptions to make.

Although these are not questions about logic itself, answers are crucial if we hope to integrate our results with other work in AI.

Acknowledgements

We thank Joe Halpern, Vladimir Lifschitz, Fangzhen Lin, John McCarthy and Amir Pnueli for their comments, as we do quite a number of others who have discussed the paper with us.

References

[1] R. Fagin and M. Y. Vardi. Internal semantics for modal logic. In *Proc. 17th ACM Symp. on Theory of Computing,* pages 305-315, 1985.

[2] J. Y. Halpern and Y. Moses. A guide to the modal logics of knowledge and belief? Preliminary draft. In *Proc. IJCAI-85,* pages 480-490, 1985.

[3] J.Y. Halpern. Knowledge, belief and uncertainty. IBM technical report (to appear), 1988.

[4] J.Y. Halpern, editor. *Proceedings of the First Conference on Theoretical issues in Reasoning about Knowledge.* Morgan-Kaufmann, Monterey, CA, 1986.

[5] J. Hintikka. *Knowledge and Belief* Cornell University Press, 1962.

[6] K. Konolige. *A Deduction Model of Belief* Morgan-Kaufmann/Pitman, 1986.

[7] S. Kraus and D. Lehmann. Knowledge, belief and time. *Theoretical Computer Science,* 58:155-174, 1988.

[8] H. Levesque. A logic of implicit and explicit belief. In *Proceedings of AAAI,* pages 198-202, Austin, TX, August 1984.

[9] R.C. Moore. *Reasoning about Knowledge and Action.* PhD thesis, M.I.T. Artificial Intelligence Laboratory, 1975.

[10] Y. Moses. Resource-bounded knowledge. In M.Y. Vardi, editor, *Proc. of the Second Conference on Theoretical Apsects of Reasoning about Knowledge,* pages 261-275. Morgan-Kaufmann, May 1988.

[11] M.Y. Vardi. A model-theoretic analysis of monotonic knowledge. In *Proc. 9th IJCAI,* pages 509-512, Los Angeles, CA, 1985.

[12] M.Y. Vardi, editor. *Proceedings of the Second Conference on Theoretical issues in Reasoning about Knowledge.* Morgan-Kaufmann, Monterey, CA, 1988.