Possible Worlds, Belief, and Modal Logic: a Tutorial

Anthony H. Dekker*

October 5, 2004

1 Introduction

This short tutorial had its genesis in a coincidence of two events: I was sorting through a pile of notes dating from the time when I was engaged as a professional logician [4], and at the same time a friend of a friend mentioned in passing the complexities of modal logic, thus motivating me to bring order to those notes, rather than throwing them away.

In *modal logic* we provide extensions to the concept "X is true." For example, we define concepts (or *modalities*) such as:

- "X is believed to be true"
- "X is known to be true"
- "X ought to be true"
- "X is eventually true"
- "X is necessarily true"

These extensions makes sense in the context of *possible worlds* or *alternate universes*. An alternate universe is one whose characteristics or history differs from our own. For example, works of fiction generally describe some kind of alternate universe, which differs from our own to a greater or lesser extent.

However, we require that these alternate universes are *logically consistent*. There may be alternate universes where Elizabeth II is not the Queen of England, but there are no alternate universes where $2 + 2 = 5$.

The extent to which alternate universes actually *exist* is a deep metaphysical question with strong connections to theology and physics. However, for our purposes we will assume that anything that can be imagined without contradiction is a valid alternate universe.

2 Belief

The easiest road to exploring possible universes is perhaps the modality of *belief*. Let us imagine a particular individual (call him John). To say that John believes X is to say that in the various imaginary universes in which John may think he lives, X is true. We recognise the cosmological metaphor when we ask: "what planet is he on?"

 The author can be contacted at dekker@acm.org

The standard logic of belief makes rather strong assumptions about John's rationality, because we have two basic axioms about belief. First:

$$\text{In any possible world } w\text{, John believes } T\text{, for every tautology } T \qquad \text{(B1)}$$

These tautologies (which we assume John believes) include all the laws of logic, such as:

- P and Q implies P or Q (i.e. P and $Q \implies P$ or Q)
- $((P \implies R)$ and $(Q \implies R)$ and $(P$ or $Q)) \implies R$
- P or not-P
- etc.

These tautologies also include all the true statements of mathematics, such as $2 + 2 = 4$ (since mathematics has essentially absorbed logic as a discipline, most mathematicians would take Kant to be incorrect in distinguishing mathematical truths as analytic *a priori* and logical truths as synthetic *a priori*).

The second axiom (or rather, rule of inference) is that John accepts all the logical implications of his beliefs. Specifically:

$$\text{In any world } w\text{, if John believes both } X \text{ and } X \implies Y\text{, then he believes } Y \qquad \text{(B2)}$$

These two axioms may be a trifle unrealistic, since they have made John out to be an almost God-like mathematician. However, in the absence of any easy way of putting formal bounds on John's mental ability, assuming B1 and B2 is our way of saying that John is rational. We *do* permit John to have incomplete or incorrect beliefs about truths that are contingent on the state of the world. For example, we permit John to believe that Michael Jackson is the president of the USA.

3 Kripke Semantics

We explain the meaning of belief statements in terms of alternate worlds in a manner due to Kripke. We invent a relationship between alternate worlds which is represented by arrows, e.g. $w_1 \to w_2$. This relationship is usually called *accessibility*, but since the meaning of the relationship depends on context, it is less confusing to simply call it "arrow."

In the context of John's beliefs, the arrow $w_1 \to w_2$ means that, if John was living in the alternate world w_1, then w_2 is one of the imaginary worlds that he would think of as possible.

If we assume that John is living in world v, and $v \to w_1$, $v \to w_2$, etc. (i.e. the w_i are all the worlds that John can imagine as possible), then we say that "John believes X" is true in world v exactly when X is true in all the imaginary worlds w_i:

$$\text{In any world } v\text{, John believes } X \iff X \text{ is true in all } w_i \text{ with } v \to w_i \qquad \text{(Kripke)}$$

In other words, John believes X exactly when X is true in all the worlds he can imagine as possible. Notice that there is no requirement that X actually be true in the world v in which John is living.

Now we can immediately see that, with this interpretation of belief, the axioms B1 and B2 work. Tautologies are true in every possible world, so in particular they are true in the w_i, and so John believes them (axiom B1). If John believes X and John believes $X \implies Y$, then X and $X \implies Y$ are true in all the w_i, and so (by the logical consistency of possible worlds) Y is true in all the w_i, and so John believes Y (axiom B2).

In general, we are able to specify a variety of modal logics, either in terms of the arrows, or in terms of axioms like B1 and B2. As we will see, the two forms of specification are equivalent.

No doubt this is all clearer with an example. Imagine four alternate worlds v, w_1, w_2, and w_3, in which the following statements are true:

v:	A	$(2 + 2 = 4)$
	B	(Bush is President of the USA)
	G	(Gold is an electrically conductive metal)
w_1:	A	
	B	
	D	(Democratic Socialist Republic is the USA's new name)
	E	(Earth was thought to be flat in the Middle Ages)
	F	(First to discover America was Columbus)
w_2:	A	
	B	
	C	(Cod liver oil tastes nice)
	E	
	F	
w_3:	A	
	C	
	E	
	H	(Hillary Clinton is the greatest leader in the USA)

Let the arrow relationships be:

$$w_1 \longleftarrow v \longrightarrow w_2 \longrightarrow w_3$$

Then if John is living in world v, the two worlds w_1 and w_2 are the worlds he can imagine as possible, and so he believes the things that are true in both of those worlds:

- In the world v, John believes A
- In the world v, John believes B
- In the world v, John believes E
- In the world v, John believes F
- In the world v, John believes not-G
- In the world v, John believes not-H

Of course, if John was living in world w_2, then world w_3 would be the only world he could imagine as possible, and so he would believe the things that were true in w_3:

- In the world w_2, John believes A

- In the world w_2, John believes C
- In the world w_2, John believes E
- In the world w_2, John believes H
- In the world w_2, John believes not-B
- In the world w_2, John believes not-D
- In the world w_2, John believes not-F
- In the world w_2, John believes not-G

If John was living in world w_1, then there are would be *no* world he could imagine as possible, and so *any* statement X would be trivially true in "all the worlds he imagines as possible," and so:

- In the world w_1, John believes X, for every statement X

Chains of arrows which stop always lead to this situation, which is why some logics exclude them, as we will see later.

To make matters even more complex, "John believes X" is itself a statement which is true or false in the worlds w_1 and w_2 that John can imagine as possible from world v, and hence John may or may not believe it. Looking at the statements believed in both w_1 and w_2, we get:

- In the world v, John believes "John believes A"
- In the world v, John believes "John believes C"
- In the world v, John believes "John believes E"
- In the world v, John believes "John believes H"
- In the world v, John believes "John believes not-B"
- In the world v, John believes "John believes not-D"
- In the world v, John believes "John believes not-F"
- In the world v, John believes "John believes not-G"

Notice that these beliefs about beliefs are different from what John actually believes. That is, not only is John mistaken about the world in which he lives (he believes E and F, which are not true in world v), but he is also mistaken about his own beliefs (he thinks that he believes C and H, when in fact he does not do so). This may portray John to be much more confused than is realistic, and we can make John's beliefs far more sensible by placing conditions on the arrow relationship between alternate worlds. We address this in the next section.

4 Conditions on Arrows

There are three conditions on the arrow relationship between possible universes that are of interest:

1. **Reflexivity** says that there is an arrow $w \to w$ from every world to itself.
2. **Transitivity** says that if $w_1 \to w_2 \to w_3$ are arrows, then there is also an arrow $w_1 \to w_3$, i.e. chains of arrows are treated like arrows too.
3. **Continuation** says that chains of arrows don't stop: in every world v there is at least one arrow $v \to w$ (this condition fails for w_1 and w_3 in our example above). One easy way of satisfying this condition is to have reflexivity, but that is not necessary.

Reflexivity means that the actual world is always one of the worlds that is imagined as possible. Let us assume that our example worlds above describe the beliefs of another person (call her Anne), but that the arrows:

$$w_1 \longleftarrow v \longrightarrow w_2 \longrightarrow w_3$$

are supplemented by reflexivity (the reader may draw in looping arrows from each world to itself, if desired). In world v, Anne believes X exactly when X is true in all three worlds she can imagine as possible (w_1, w_2, and her current world v). So Anne believes less than John does in the same circumstances:

- In the world v, Anne believes A
- In the world v, Anne believes B
- In the world v, Anne believes not-H

Reflexivity means that Anne always imagines the actual world as possible, and so Anne's beliefs are always true in the world that she finds herself in. In some modal logics, this is taken as an axiom:

$$\text{In any world } w, \text{“Anne believes } X\text{”} \implies X \qquad \text{(B3)}$$

5 Doxastic Logic

Transitivity in our example means that there is an arrow $v \to w_3$ as well as:

$$w_1 \longleftarrow v \longrightarrow w_2 \longrightarrow w_3$$

(the reader may draw in the extra arrow, if desired). Let us assume that Peter has a transitive (but not reflexive) belief system. From world v, Peter can imagine three worlds as possible (w_1, w_2, and w_3), and so:

- In the world v, Peter believes A
- In the world v, Peter believes E
- In the world v, Peter believes not-G

From worlds w_1 or w_2, Peter's beliefs are the same as John's, and so:

- In the world v, Peter believes "Peter believes A"
- In the world v, Peter believes "Peter believes C"
- In the world v, Peter believes "Peter believes E"
- In the world v, Peter believes "Peter believes H"
- In the world v, Peter believes "Peter believes not-B"
- In the world v, Peter believes "Peter believes not-D"
- In the world v, Peter believes "Peter believes not-F"
- In the world v, Peter believes "Peter believes not-G"

Transitivity means that worlds accessible from v include those accessible from worlds one step away (w_1 or w_2), and so transitivity corresponds to the rule of *introspection*:

$$\text{In any world } w\text{, "Peter believes } X\text{"} \implies \text{Peter believes "Peter believes } X\text{"} \quad \text{(B4)}$$

i.e. if Peter believes something, then he believes that he believes it. His beliefs may still be wrong, however. Peter thinks that he believes C and H, but he doesn't really believe them. He also believes E, which isn't true in the world v in which he lives.

This kind of logic of belief with transitivity is called *doxastic logic* (from the Greek *doxa*, belief), and it is quite useful. In the past [4], I have used it for reasoning formally about electronic commerce: if I am trying to buy something from you over the Internet, what does it take to make you believe that I am who I say I am?

6 Epistemic Logic

If we combine reflexivity, as in Anne's beliefs (all beliefs are true) with transitivity, as in Peter's beliefs (i.e. introspection), we obtain *epistemic logic*, the logic of knowledge (from the Greek *episteme*, knowledge). If Cathy's belief system is both reflexive and transitive, then we get new versions of the axioms B1 to B4 (replacing "believes" with "knows," to acknowledge the truth of her beliefs, since knowledge can be defined as true belief):

$$\text{In any world } w\text{, Cathy knows } T\text{, for every tautology } T \quad \text{(B1}'\text{)}$$

$$\text{In any world } w\text{, if Cathy knows both } X \text{ and } X \implies Y\text{, then she knows } Y \quad \text{(B2}'\text{)}$$

$$\text{In any world } w\text{, "Cathy knows } X\text{"} \implies X \quad \text{(B3}'\text{)}$$

$$\text{In any world } w\text{, "Cathy knows } X\text{"} \implies \text{Cathy knows "Cathy knows } X\text{"} \quad \text{(B4}'\text{)}$$

From the last two axioms, it follows that "Cathy knows X" and "Cathy knows "Cathy knows X" " are exactly equivalent.

If we add reflexivity and transitivity to our example:

$$w_1 \longleftarrow v \longrightarrow w_2 \longrightarrow w_3$$

(the reader may draw in the extra arrows, if desired), then the only remaining belief that we have is:

- In the world v, Cathy knows A

But although everything that Cathy knows is true, her knowledge is incomplete. Reflexivity has tamed Cathy's wild imagination: having imagined the worlds w_1, w_2, and w_3, she doesn't believe the statements C, D, E, F, or H. However, B and G are true in her current world v, but she doesn't know that.

For every statement X that Cathy knows, if we can provide a derivation of "Cathy knows X" using axioms and rules, then we will have provided an explanation as to *how* Cathy knows X.

7 Deontic Logic

Let us forget reflexivity for now, and consider other ways of ensuring that chains of arrows don't stop (continuation, i.e. in each world w_1, there is at least one arrow $w_1 \rightarrow w_2$). In *deontic logic*, we interpret the arrow $w_1 \rightarrow w_2$ as saying that w_2 is a *better version* of the world w_1, and we interpret beliefs as beliefs about what *ought to be done* (rather than about what is true).

Continuation makes good sense in this context: if we can't imagine how things could be better, how can we know what ought to be done? We also require transitivity, so we get variations of the axioms B1, B2, and B4:

$$\text{In any world } w, T \text{ is obligatory, for every tautology } T \qquad (\text{B1}'')$$

$$\text{In any world } w, \text{ if both } X \text{ and } X \Longrightarrow Y \text{ are obligatory, then } Y \text{ is obligatory} \qquad (\text{B2}'')$$

$$\text{In any world } w, \text{“}X \text{ is obligatory”} \quad \Longrightarrow \quad \text{“}X \text{ is obligatory” is obligatory} \qquad (\text{B4}'')$$

Continuation gives us a new axiom:

$$\text{In any world } w, \text{“}X \text{ is obligatory”} \quad \Longrightarrow \quad \text{not “not-}X \text{ is obligatory”} \qquad (\text{D})$$

This axiom provides a degree of logical consistency for obligation. It follows because, if "X is obligatory" in world v, then X is true in all worlds w_i where $v \rightarrow w_i$ (by continuation there is at least one such w_i). So not-X must be false in all the w_i, and so not-X cannot be obligatory in world v.

The double-not combination on the right-hand side of axiom D is quite useful, and we abbreviate it is as "X is permissible":

$$\text{In any world } w, \text{“}X \text{ is permissible”} \quad \Longleftrightarrow \quad \text{not “not-}X \text{ is obligatory”} \qquad (\text{Permissibility})$$

Standard deontic logic is believed by some people to be inconsistent (see the article on "deontic paradoxes" in [2]). Consider the following set of seemingly obvious statements:

(a) Some individual (call him Fred) is a criminal.

(b) It is obligatory that criminals be punished, i.e.
Fred is a criminal $\Longrightarrow$ "Fred will be punished" is obligatory.

(c) It is obligatory that innocent people not be punished, i.e.
(Fred is not a criminal $\Longrightarrow$ Fred will not be punished) is obligatory.

(d) It is obligatory that Fred not be a criminal.

From (a) and (b) we infer that it is obligatory that Fred be punished, while from (c) and (d) we inferred that it is obligatory that he not be punished. But this contradiction is inconsistent with the D rule!

Possible-world semantics casts considerable light on this apparent paradox. Consider the current (imperfect) world v, connected to better worlds $v \to w_i$. Statement (d) says that in every better world w_i, Fred is a good citizen, and not a criminal. This is reasonable, but then by statement (c) he will be free from punishment in every better world.

Statement (a) says that in the imperfect world v, Fred is a criminal. This is also reasonable. But statement (b) is less reasonable: it says that, because Fred is a criminal in the current imperfect world v, he should be punished in the alternate universes w_i where he is innocent! Clearly, we would not wish this to be the case.

The paradox is thus revealed to be due to sloppy formulation of Fred's situation, and possible-world semantics has made this clear. A more reasonable version of (b) is:

(e) It is obligatory that criminals be punished, i.e.
(Fred is a criminal $\implies$ Fred will be punished) is obligatory.

This says that that in every better world w_i, Fred is punished if he is a criminal *in that world*. Together with (c), this means that he is punished exactly when he is a criminal. This allows either for some better worlds in which Fred is a criminal (and is punished), or for statement (d) where he is never a criminal in a better world (and is therefore never punished).

8 Temporal Logic

Temporal logic is the logic of *time*. In temporal logic, we interpret the possible worlds w_t as different versions of the current world, at various times t. We interpret the arrows $w_t \to w_{t+1}$ as "world w_{t+1} follows after world w_t," i.e. the possible worlds are exactly the future states of the present world. We therefore have a *history* of worlds:

$$w_0 \longrightarrow w_1 \longrightarrow w_2 \longrightarrow w_3 \longrightarrow \cdots$$

together with reflexivity and transitivity. The result is similar to epistemic logic, but instead of "believes" or "knows," we have the modal operator "henceforth," which means "true from now on." The Kripke rule then becomes:

$$\text{In any world } w_t\text{, "henceforth } X\text{"} \iff X \text{ is true in } w_t, w_{t+1}, w_{t+2}, \ldots \qquad \text{(Kripke}')$$

We also get new versions of rules B1 to B4:

$$\text{In any world } w_t \text{ (i.e. at any time } t\text{), henceforth } T\text{, for every tautology } T \qquad \text{(B1}''')$$

$$\text{At any time } t\text{, ("henceforth } X\text{" and henceforth "}X \implies Y\text{")} \implies \text{henceforth } Y \qquad \text{(B2}''')$$

$$\text{At any time } t\text{, "henceforth } X\text{"} \implies X \qquad \text{(B3}''')$$

$$\text{At any time } t\text{, "henceforth } X\text{"} \implies \text{henceforth "henceforth } X\text{"} \qquad \text{(B4}''')$$

The statement "henceforth X" means that X is true now and at all future times. The axioms B3 and B4 tell us that this is exactly equivalent to "henceforth "henceforth X" ". Analogously to the "permissible" operator in deontic logic, we can define "eventually X" to be:

not (henceforth not-X)

This makes sense, because if it is not the case that "henceforth not-X", then it is not the case that not-X is true now and into the future, which means that X must be true at least once, either now or at some future time.

We can easily show that "eventually "eventually X" " is exactly equivalent to "eventually X", but other combinations of operators are more interesting. The combination "eventually "henceforth X" " means that there will come a time when X is true and remains true forever. The combination "henceforth "eventually X" " means that at every future time, X is eventually true, i.e. X is true infinitely often.

The version of temporal logic presented here is due to Pnueli [8], who also includes the operator "next," defined by:

$$\text{At any time } t\text{, “next } X\text{”} \iff X \text{ is true at time } t+1 \qquad \text{(Next)}$$

i.e. "next X" means that X is true at the next "tick of the clock." Alternatively, we can capture the meaning of the "next" operator with five axioms:

$$\text{At any time } t\text{, next not-}X \iff \text{not “next } X\text{”} \qquad \text{(N1)}$$

$$\text{At any time } t\text{, (“next } X\text{” and next “}X \Longrightarrow Y\text{”)} \Longrightarrow \text{next } Y \qquad \text{(N2)}$$

$$\text{At any time } t\text{, “henceforth } X\text{”} \Longrightarrow \text{next } X \qquad \text{(N3)}$$

$$\text{At any time } t\text{, “henceforth } X\text{”} \Longrightarrow \text{next “henceforth } X\text{”} \qquad \text{(N4)}$$

$$\text{At any time } t\text{, (}X \text{ and henceforth “}X \Longrightarrow \text{next } X\text{”)} \Longrightarrow \text{henceforth } X \qquad \text{(N5)}$$

The last axiom is particularly useful to computer scientists, in proving that some property of a computer program is always true. For software inside heart pacemakers, telephone switches, railway control computers, and nuclear power plants, such proofs are of critical importance.

More traditional versions of temporal logic (e.g. tense logic) are not reflexive, so that "henceforth" and "eventually" refer solely to the future, and do not include the present. However, those other versions of temporal logic are not fundamentally different in terms of what they can prove.

Reasoning about time goes back to Aristotle and Diodorus Cronus. Mediaeval Scholastics also devoted considerable thought to time, particularly the Latin verbs *incipit* (it begins) and *desinit* (it ends). With the Renaissance, these distinctions fell out of favour, and (the computer not having been invented yet) seemed useless. Juan Luis Vives strongly ridiculed the "subtle rigourism" of mediaeval thought on time in his 16$^{\text{th}}$ century *Adversus pseudodialecticos* [6], and later books on logic (such as the great 17$^{\text{th}}$ century book by Antoine Arnauld and Pierre Nicole [1]) did not include any formalisation of time.

9 Necessity

If we connect arrows between *every* pair of possible worlds, we obtain a logic of *necessity.* The statement X is *necessarily true* if it is true in every possible world. The Kripke rule then becomes:

$$\text{“necessarily } X\text{”} \iff X \text{ is true in every world } w \qquad \text{(Kripke}''\text{)}$$

We no longer need to specify in which world the statement "necessarily X" is supposed to be true, because (from the definition), if "necessarily X" is true in one world, it will be true in every world. We again get variations of the axioms B1 to B4:

$$\text{necessarily } T\text{, for every tautology } T \qquad \text{(B1}''''\text{)}$$

$$(\text{“necessarily } X\text{” and necessarily “} X \Longrightarrow Y\text{”}) \Longrightarrow \text{necessarily } Y \qquad (\text{B2}'''')$$

$$\text{“necessarily } X\text{”} \Longrightarrow X \text{ (in any world } w) \qquad (\text{B3}'''')$$

$$\text{“necessarily } X\text{”} \Longrightarrow \text{necessarily “necessarily } X\text{”} \qquad (\text{B4}'''')$$

Analogously to the “permissible” operator in deontic logic, or the “eventually” operator in temporal logic, we can define “possibly X” to be:

not (necessarily not-X)

This makes sense, because if it is not the case that “necessarily not-X”, then it is not the case that not-X is true in every possible world, which means that X must be true in at least one possible world. It is easy to show that if X is necessary, then X is possible.

The statement “possibly X” doesn't need to be referred to a particular world either, since, if “possibly X” is true in one world, it will be true in every world. In fact, we can add a new axiom:

$$\text{“possibly } X\text{”} \Longrightarrow \text{necessarily “possibly } X\text{”} \qquad (\text{B5})$$

We can divide the true statements in our current world into necessary truths (which are true in every world), and *contingent* truths (which are possibly false). The necessary truths include the laws of mathematics and logic. But what else? The answer depends on what we consider possible, which brings us back to the question of exactly what is a possible world. Is the future determined, for example? What about the past—could it have happened differently? Could the laws of physics have been different? These are all questions for serious debate.

We have also skipped over some serious questions of epistemology. The truths of geometry are necessary, for example (insofar as they are statements about abstract Euclidean or non-Euclidean spaces), and we have treated them as analytic. But it *is* true to say that, in practice, those truths were recognised by examining the real world, and only then thinking abstractly. So was Kant right after all? Plato would have said no—he would have said that the truths of geometry are *remembered* from a past spiritual existence, and the real world at best triggers those memories. Christian philosophers influenced by Neoplatonism would have argued that the truths of geometry are understood by divine illumination. Again, these are serious issues. However, having a formal mechanism for reasoning about belief, knowledge, time, and necessity, we are better equipped for handling some of these perennial questions. Table 1 summarises the logics we have considered.

References

[1] Arnauld, A. and Nicole, P., “La logique ou l'art de penser,” Paris, 1662. Partly available electronically at `http://www.ff.iij4u.or.jp/~yyuji/library/port-royal/lartdepenser.html`

[2] Audi, R., “The Cambridge Dictionary of Philosophy,” 2nd ed, Cambridge University Press, 1999.

[3] Bowen, K. A., “Model Theory for Modal Logic,” D. Reidel, Dordrecht, 1979.

[4] Dekker, A. H., *C3PO: a Tool for Automatic Sound Cryptographic Protocol Analysis*, Proceedings of 13th IEEE Computer Security Foundations Workshop, Cambridge, England, July 3–5, 2000. Available electronically at `http://www.acm.org/~dekker/c3po.pdf`

System	Axioms	Arrows	Standard Name
John's beliefs	B1,B2	arbitrary	K modal logic
Anne's beliefs	B1–B3	reflexive	T modal logic
Peter's beliefs (doxastic logic)	B1,B2,B4	transitive	K4 modal logic
Cathy's beliefs (epistemic logic)	B1–B4	reflexive, transitive	S4 modal logic
Deontic logic	B1,B2,B4,D	continuation, transitive	D4 modal logic
Temporal logic	B1–B4,N1–N5	linear sequence	DX temporal logic
Logic of necessity	B1–B5	fully connected	S5 modal logic

Table 1: Summary of Modal Logics

[5] Fitting, M., "Proof Methods for Modal and Intuitionistic Logics," D. Reidel, Dordrecht, 1983.

[6] Hasle, P. and Øhrstrøm, P., *The Flow of Time into Logic—and Computer Science*, EATCS Bulletin **82** (Feb 2004) 191–226.

[7] Meyer, J.-J. Ch., van der Hoek, W. and Vreeswijk, G. A. W., *Epistemic Logic for Computer Science: A Tutorial (Parts 1 and 2)*, EATCS Bulletin **44** (June 1991) 242–270 and **45** (October 1991) 256–287.

[8] Pnueli, A., *The Temporal Semantics of Concurrent Programs*, Theoretical Computer Science **13** (1981) 45–60.

[9] Wallen, L. A., "Automated Deduction in Nonclassical Logics," MIT Press, 1990.