Three Interactions between Context and Epistemic Locutions

Richmond H. Thomason
Philosophy Department
University of Michigan
Ann Arbor, MI 48109-2110, USA

email: rich@thomason.org

Note: This paper was presented in Roskilde at Context 2007. The May 20, 2007 version of the paper was published in *Modeling and Using Context: 6th International and Interdisciplinary Conference*, edited by Boicho Kokinov, Daniel C. Richardson, Thomas R. Roth-Berghofer, and Laure View, Springer-Verlag, Berlin, 2007, pp. 467–481.

This is the August 31, 2007 version of the same paper, prepared for public distribution over the Internet. It is an expanded version of the published paper, and contains some corrections. In particular, some typographical errors and oversimplifications in the definition of $\models_1$ in Section 2 are corrected in this version.

Abstract

I motivate and formalize three interactions between context and epistemic modalities, discussing (1) the effects of temporal perspective on knowledge of the future, (2) the effects of perceived risk on belief, and (3) the effects of presupposed background information on epistemic 'might'.

1. Introduction

From the start, philosophy has been full of arguments that start off plausibly enough, but that lead, through a series of apparently sound steps, to implausible or even totally unacceptable conclusions. These arguments strike many people—philosophers and nonphilosophers—as puzzles. If the conclusion is unacceptable, there must be a mistake somewhere in the argument. The puzzle will be resolved when the mistake has been identified and explained.

Two sorts of resolution methods stand out as particularly useful in this enterprise: appeals to ambiguity and to context dependence. It is easy to see why a syntactically correct argument—say, an application of *modus ponens*—can be mistaken if an expression is used with one sense in the premises and with quite another sense in the conclusion, or if expressions are used in one context in the premises and another context in the conclusion.[1] And both explanations have been exploited in philosophy.

Of course, the ambiguities and contextual dependencies that are important for philosophy have to be subtle. If they were obvious, there would be no puzzle in the first place; the arguments would obviously be wrong and the solutions would be ready to hand. But although the skeptical arguments of Sextus Empiricus, for example, have conclusions that are obviously wrong, it is very difficult to identify the mistake, and to give a satisfactory explanation of the flaw in the reasoning.

Cartesian demons and evil neurological experimenters are typical protagonists in the philosophy classrooms where problems of this kind are introduced to students. We can imagine experiments where disembodied brains are provided with experiences just like those of normal people. Since we can't exclude these possibilities, we ourselves might be just such a brain. So, after all, we don't know that we have hands.

Philosophers' unfortunate fondness for ridiculously far-fetched examples of this kind can create a misleading impression that knowledge is problematic only in introductory philosophy classrooms. But the problem is much more general than this. Let's consider some more realistic examples.

Take, for instance, the following two quotations from the Brown Corpus (my underlining added).[2]

> (1.1) As for food, Mrs. Henry Louchheim, chairman of this phase, is a globetrotter who knows good food. "New Orleans"? she says,"Of course I've had the best. It is just bad luck that we are having the party in a month with no 'R's, so no oysters. But we have lots of other New Orleans specialties. I <u>know</u> they will be good. We've tried them out on the club chef—or say, he has tried them out on us and we have selected the best".

> (1.2) Just a month after the Korean War broke out, the 7th Cavalry was moving into the lines, ready for combat. From then on the Fighting Seventh was in the thick of the bitterest fight in Korea. One night on the Naktong River, Mel Chandler called on that fabled *esprit de corps*. The regiment was dug in on the east side of the river and the North Koreans were steadily

[1] The variable-fixing expression 'let x be' provides a simple and unproblematic example of a contextual fallacy. *Premises:* 'Let x be 2. Then x is even. Let x be 3. Then x is odd.' *Conclusion:* 'x is even and odd.'

[2] The Brown Corpus was collected in 1961. It contains over a million words of representative English prose from various genres. I have provided long quotations because the point I want to make, I believe, is better supported by naturally occurring examples that also supply background detail about the reasoning on which the knowledge is based.

> building up a concentration of crack troops on the other side. The troopers knew an attack was coming, but they didn't know when, and they felt lonely and depressed.

You can quibble with these claims to knowledge. You could point out that food can spoil unexpectedly—for instance if the refrigeration fails. You could say that the attack would not have happened if the North Korean government suddenly told their troops to stay in place, perhaps because they had arbitrarily and unexpectedly decided to negotiate for peace. Imagine an articulate and persuasive critic pointing out possibilities to Mrs. L in which the food will not be good, citing instance after instance where expert caterers had ordered good food from reliable chefs, but where, through some misfortune, bad food had been served. Under determined criticism of this sort, Mrs. L may well have to withdraw her knowledge claim, even though she may feel that she has been tricked or manipulated into doing so.

Examples of this kind show convincingly that the extreme and far-fetched skeptical examples of which philosophers are so fond are limiting cases of natural cases where perfectly ordinary knowledge claims can be undermined. Why is knowledge useful, if it is so fragile? A diagnosis of the flaw in a troublesome argument is more than a logical exercise; in order to resolve the difficulty, it needs to clarify the nature of knowledge and the role of knowledge claims in communication and reasoning.

In [Lewis, 1979], David Lewis proposed the idea of a "conversational score," and described ways in which conversational acts depend on the score, and can influence it. Two of these dependencies (the possibilities over which a modal auxiliary like 'can' are taken to range, and the standards of accuracy appropriate for an adjective like 'smooth') suggest ways of trying to disarm some well-known philosophical puzzles. These ideas have led to a great deal of activity in the intersection of epistemology and philosophy of language; see [Preyer and Peter, 2005a] for a collection of papers on this topic, and references to others.

As you would expect, the linguistic evidence sheds some additional light, but at the same time adds another layer of complexity to the already complex issues. So the extent to which the philosophical problems are resolved or even illuminated is controversial. Genuine philosophical problems are recalcitrant, even to multi-disciplinary solutions.

Nevertheless, I believe that our best hope for making progress on these problems lies in a multi-disciplinary approach that brings to bear all the areas of Cognitive Science. What I miss in the recent philosophical work on context and epistemology is sufficient attention to reasoning, and to the role of epistemic locutions in realistic examples where the epistemic notions are actually engaged. In fact, many examples in the current literature go beyond typical philosophical examples in their complexity.[3] They involve some minimal stage setting, more than one sentence, and judgments about whether the sentences are appropriate or true. But few of these examples actually show the characters actually engaged in reasoning, and many of them are not entirely realistic.

I also suspect that we might gain more in the long run by relaxing our concentration on large-scale philosophical topics, like the refutation of skepticism. The linguistic work on context and epistemology has revealed unexpected features of locutions like 'know' and 'might', but the most interesting things we have learned seem to be at best distantly related to the original problem. I believe that we can improve our theories by combining models

[3] I am thinking of examples like: "I direct my eyes at a ripe tomato. I see red."

of compositional semantics and domain reasoning with a systematic study of the work done by epistemic concepts in realistic, moderately complex examples of reasoning. These things are well worth doing, even if we remain unable to settle the larger question of philosophical skepticism.

In this paper I will try to illustrate this methodology by considering three cases: knowledge of the future, the effects of risk on belief, and epistemic 'might'.

2. Uncertain Knowledge of the Future

Indeterminist tense logics seek to provide a semantics for tense operators in models that allow multiple future outcomes, giving equal weight to each alternative. Supervaluations, a technique proposed in [van Fraassen, 1969], can be used for this purpose; the idea was first proposed in [Thomason, 1970] and has been taken up and developed by later authors. (See [Belnap, Jr. *et al.*, 2001] and the references therein.) The idea is to first define truth relative to an arbitrarily selected future history; a sentence is then said to be true *simpliciter* if it is true relative to all future histories.

To formalize this, we will work with a propositional modal language $\mathcal{L}_1$ that has past and future modalities [P] and [F][4] and a historical necessity operator [H]. Frames and models for this language are defined as follows.

Definition 1. Frames, histories, valuations, models, evaluation points.

A *frame* for indeterministic logic is a pair $\mathcal{F} = \langle \mathrm{M}, \preceq \rangle$, where M is a nonempty set (of moments) and $\prec$ is a transitive, antireflexive relation over M such that if $m_1, m_2 \prec m$, then $m_1 \prec m_2$ or $m_2 \prec m_1$ or $m_1 = m_2$.

A *history* h on a frame $\mathcal{F} = \langle \mathrm{M}, \preceq \rangle$ is a maximal $\prec$ chain on $\mathcal{F}$. $\mathcal{H}_m$ is the set of histories h of $\mathcal{F}$ such that $m \in h$. $\mathcal{H}$ is the set $\bigcup_{m \in \mathrm{M}} \mathcal{H}_m$ of all histories of $\mathcal{F}$.

A *valuation* of a set $\mathcal{P}$ of propositional atoms on a frame $\mathcal{F} = \langle \mathrm{M}, \preceq \rangle$ is a function V that inputs an atomic formula $p \in \mathcal{P}$ and outputs a subset $V(p)$ of M.

A *model* on a set $\mathcal{P}$ of propositional atoms and a frame $\mathcal{F}$ is a pair $\langle \mathcal{F}, V \rangle$, where $\mathcal{F}$ is a frame and V is a valuation of $\mathcal{P}$ on $\mathcal{F}$.

An *evaluation point* (or *e-point*) in a model $\mathcal{M} = \langle \mathrm{M}, \preceq \rangle$ is a pair $\langle m, h \rangle$, where h is a history over $\langle M, \preceq \rangle$ and $m \in h$.

The satisfaction relation $\mathcal{M}, m, h \models \phi$ between a model $\mathcal{M} = \langle \mathrm{M}, \preceq \rangle$, an e-point in $\mathcal{M}$, and a formula ϕ of $\mathcal{L}_1$ is defined recursively as follows.

Definition 2. $\mathcal{M}, m, h \models \phi$.

[4] [F] is "always in the future", [P] is "always in the past." The duals are ⟨F⟩, "sometimes in the future," and ⟨P⟩, "sometimes in the past."

1. **Basis:** $\mathcal{M}, m, h \models p$ iff $m \in V(p)$.
2. **Booleans:** Boolean conditions are routine.
3. **Past:** $\mathcal{M}, m, h \models \langle \text{P} \rangle \phi$ iff for some $m' \prec m$, $\mathcal{M}, m', h \models \phi$.
   $\mathcal{M}, m, h \models [\text{P}]\phi$ iff for all $m' \prec m$, $\mathcal{M}, m', h \models \phi$.
4. **Future:** $\mathcal{M}, m, h \models \langle \text{F} \rangle \phi$ iff for some m', $m \prec m'$ and $m' \in h$, $\mathcal{M}, m', h \models \phi$.
   $\mathcal{M}, m, h \models [\text{F}]\phi$ iff for all m', $m \prec m'$ and $m' \in h$, $\mathcal{M}, m', h \models \phi$.
5. **Historical Necessity:** $\mathcal{M}, m, h \models [\text{H}]\phi$ iff for all h' such that $m \in h'$, $\mathcal{M}, m, h' \models \phi$.

Applying van Fraassen's supervaluation idea to this notion of satisfaction relative to a history provides the following definition of satisfaction *simpliciter*.

Definition 3. $\mathcal{M}, m \models \phi$.
$\mathcal{M}, m \models \phi$ iff for all $h \in \mathcal{H}_m$, $\mathcal{M}, m, h \models \phi$.

In Model $\mathcal{M}_1$, for instance, neither $\langle \text{F} \rangle p$ nor $\neg\langle \text{F} \rangle p$ is true at m_0 (neither $\mathcal{M}_1, m_0 \models \langle \text{F} \rangle p$ nor $\mathcal{M}_1, m_0 \models \neg\langle \text{F} \rangle p$) because $\mathcal{M}_1, m_0, h_1 \models \langle \text{F} \rangle p$, while $\mathcal{M}_1, m_0, h_1 \models \neg\langle \text{F} \rangle p$. But, for the same reason, $[\text{H}]\langle \text{F} \rangle p$ and $[\text{H}]\neg\langle \text{F} \rangle p$ are both false at m_0. Related to this is the fact that a formula like $\langle \text{F} \rangle p$ is *future-dependent*—its truth at a moment has to be evaluated with respect to a postulated future—whereas $[\text{H}]\langle \text{F} \rangle p$ is future *in*dependent.

So far, I have simply restated ideas that go back to [Thomason, 1970] and which I hope the reader will be willing to assume as background. Now, however, I want to add a knowledge operator $[\text{K}]$ to this temporal language, and to ask whether $[\text{K}]\langle \text{F} \rangle p$ is future dependent. In effect, we are asking whether, in an indeterministic setting, someone's claim to know something may turn out to have been true in one subsequent scenario, and to have been false in another.

Model $\mathcal{M}_1$:

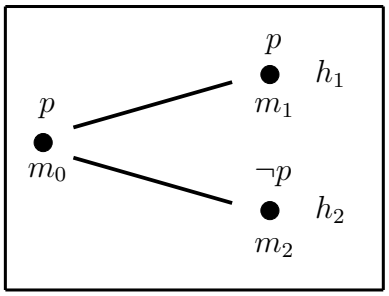

Before we can set out the satisfaction conditions for knowledge, we will need some motivation and philosophical background.

If we are interested in ordinary, human knowledge of the future (rather than the "certain knowledge" that some theologians talk about), it certainly seems as if knowledge can be future-dependent in this way, and hence can be fallible.

Returning to Examples (1.1) and (1.2), cited in Section 1, in which Mrs. Louchheim knows the food will be good and the 7th Cavalry knows the enemy will attack, let's suppose

that in both cases things turned out just as expected. Mrs. L found the food to be very good, and the North Koreans attacked the next day across the Naktong. We are told that Mrs. L knew her food, and had sampled the specialties to be served at the party; we can assume that the 7th Cavalry were observing the enemy and the observations were supplemented with other information and evaluated by military intelligence experts.

Consider again the objections that a food-ordering skeptic or a military intelligence skeptic might raise to Mrs. L or the 7th Cavalry. The food skeptic, says, for instance, that New Orleans is hot and the refrigeration can always fail. Depending on details that I haven't supplied, in some cases the objections may be useful; but in many cases we would say they are *beside the point*. I would like to claim that, in calling them beside the point, we are classifying them as inappropriate, rather than (say) as simply mistaken, untrue, or without force, and that the sort of mistake we have in mind here is related to context. I'm sympathetic to those who want to say that attributions of knowledge are relative to standards of justification. Raising an objection to a knowledge claim that only has force if these standards are inappropriately strict is one way of being beside the point.

But others have discussed this matter at some length.[5] I want to investigate a different and, I think, unnoticed additional way in which knowledge has to depend on context. At least, this sort of context dependence is needed if we, and Mrs. L, and the 7th Cavalry, are ever to be in a position to have knowledge about genuine future contingencies.

The problem is this. In each case, the attack on the knowledge claim involves a scenario that apparently is possible, even if far-fetched. In an indeterminist setting, it is plausible to suppose that these are historically possible scenarios. For instance, nothing about the world at the time Mrs. L made her prediction ruled out a refrigerator failure at just the right time to ruin the food for her party. Let us grant this.

At m_0, Mrs. L says she knows the food will be good. But now, let's follow the improbable scenario until we reach a moment m_2 at which spoiled food is served at the party. From the standpoint of m_2, we *must* say that Mrs. L didn't know at m_0 that the food was going to be good. She was wrong because at m_0, the food was going to be bad. Her knowledge claim fails, not because she lacked an adequate justification, but because what she claimed to know was false.

But if later on, at m_2, we would be correct in saying she didn't know that the food would be good, don't we have to conclude that in fact, she didn't know it when she made her claim? At m_2 we can say with certainty that she didn't know at m_0 that the food would be good. But this means that at m_0 her claim to know it would be good was false. So, it seems that at m_0 she can't have known, after all, that the food would be good, even if (luckily) it does turn out to be good.

Note that this attack on the possibility of knowledge has nothing to do with standards of justification—it uses only temporal reasoning, and a semantic characteristic of knowledge—that you can't know things that are false.

This sort of argument against the possibility of uncertain knowledge of the future can be disarmed, but doing so requires developing a rather complex and delicate account of the semantics of future-dependent claims. I do not think that this account could have been

[5]See the discussion of this issue in the papers collected in [Preyer and Peter, 2005b], and in the works cited in that volume.

worked out without the clarity that comes with formalization and the use of model-theoretic techniques.

First, notice that an argument similar to the one we used to attack knowledge at m_0 can be deployed in terms of truth rather than knowledge. We refer again to Model $\mathcal{M}_1$, and simply consider $\langle \text{F} \rangle p$ at m_0. Notice that from the standpoint of m_2, $\neg\langle \text{P} \rangle\langle \text{F} \rangle p$ is true. But if $\neg\langle \text{P} \rangle\langle \text{F} \rangle p$ is true at m_2, then (looking back from m_2), $\langle \text{F} \rangle p$ should be *false* at m_0, whereas the supervaluational theory we endorsed makes it neither true nor false.

In [MacFarlane, 2002], John MacFarlane addresses the problem of after-the-fact truth by making satisfaction relative to not one, but two moments: a "context of evaluation" and a "context of assessment." The satisfaction scheme then becomes $\mathcal{M}, m, m', h \models \phi$, where m is the moment of evaluation and m' the moment of assessment. And ϕ is true *simpliciter* in $\mathcal{M}$ at $\langle m, m' \rangle$ if $\mathcal{M}, m, m', h \models \phi$ for all $h \in \mathcal{H}_{m'}$—supervaluational truth is reckoned using the histories through the moment of assessment rather than the larger set of histories through the (possibly earlier) moment of evaluation.

In Model $\mathcal{M}_1$, for instance, $\langle \text{F} \rangle p$ will be true *simpliciter* at $\langle m_0, m_1 \rangle$ and false *simpliciter* at $\langle m_0, m_2 \rangle$. It will be neither true nor false, however, at $\langle m_0, m_0 \rangle$.

In recent unpublished work [Thomason, 2007], I have proposed a slightly different approach, which uses double-indexing techniques of the sort that go back to [Kaplan, 1978]. According to this theory, the recursive definition of satisfaction is modified as follows, to incorporate MacFarlane's two temporal indices.

Definition 4. $\mathcal{M}, m, m', h \models_1 \phi$.

1. **Basis:** If $\phi \in \mathcal{P}$ then $\mathcal{M}, m, m', h \models_1 \phi$ iff $m \in V(\phi)$.
2. **Booleans:** Boolean conditions are routine.
3. **Past:** $\mathcal{M}, m, m', h \models_1 \langle \text{P} \rangle \phi$ iff for some $m_1 \prec m$,
$\mathcal{M}, m_1, m', h \models_1 \phi$.
$\mathcal{M}, m, m', h \models_1 [\text{P}]\phi$ iff for all $m_1 \prec m$, $\mathcal{M}, m_1, m', h \models_1 \phi$.
4. **Future:** $\mathcal{M}, m, m', h \models_1 \langle \text{F} \rangle \phi$ iff for some m_1, $m \prec m_1$ and
$m_1 \in h$, $\mathcal{M}, m_1, m_2, h \models_1 \phi$, where
$m_2 = \max(m_1, m')$.
$\mathcal{M}, m, m', h \models_1 [\text{F}]\phi$ iff for all m_1, $m \prec m_1$ and $m_1 \in h$,
$\mathcal{M}, m_1, m_2, h \models_1 \phi$, where $m_2 = \max(m_1, m')$.
5. **Historical Necessity:** $\mathcal{M}, m, m', h \models_1 [\text{F}]\phi$ iff for all h' passing
through m, $\mathcal{M}, m, m', h' \models_1 \phi$.

We now define truth *simpliciter* by first insisting, as Kaplan does, that evaluation must start at a *normal index*, where the moment of evaluation is the same as the moment of assessment. This gives us a notion of truth relative to a single e-point. We then use supervaluations, considering a formula to be simply true at m when it is true relative to all histories passing through m.

Definition 5. $\mathcal{M}, m, h \models_1 \phi$ and $\mathcal{M}, m \models_1 \phi$.
$\mathcal{M}, m, h \models_1 \phi$ iff $\mathcal{M}, m, m, h \models_1 \phi$.
$\mathcal{M}, m \models_1 \phi$ iff $\mathcal{M}, m, h \models_1 \phi$ for all $h \in \mathcal{H}_m$.

For a language with only future and past tenses and historical necessity, this account of truth at a moment does not differ from MacFarlane's. But if truth is added to the language, there are substantive differences from the sort of semantics that MacFarlane recommends. This two-stage approach to satisfaction also helps when a knowledge operator [K] is added.

The standard semantic treatment of knowledge in temporal settings is presented in [Fagin *et al.*, 1995]. This theory uses possible-worlds semantics rather than branching time. A possible world in the Fagin-Halpern-Moses models is analogous to an e-point or moment-history pair in our models. Therefore, in a branching setting, we need to treat an epistemic possibility as an e-point.

In the simplest case (and we do not want to make things more complicated than they need to be), we interpret knowledge at world w using a set K_w of worlds containing w; this interpretation yields the modal logic **S5**. Therefore, in a branching-time structure in which we neglected the sensitivity of knowledge to an after-the-fact perspective, an agent's knowledge would be represented by a function from e-points to sets of e-points.

We want, however, to model knowledge in such a way that the interpretation of a sentence like 'Mrs L knows that the food will be good' may depend a future perspective from which it is assessed. At a moment in time at which the food is served and is good, 'Mrs L knew

that the food would be good' may be true, while at the moment prior to the catering event 'Mrs. L knows that the food will be good' may not be true.

To take the potential effect of the moment of assessment into account, we need first to relativize the function representing knowledge to moments: rather than one function K from e-points to sets, we need a family of functions $\mathrm{K}_{m'}$. Here, m' is a later moment from which an agent's knowledge at an earlier moment is assessed in view of the subsequent history.

Second, we need to take into account that the satisfaction relation $\mathcal{M}, m, m', h \models \phi$ for a model $\mathcal{M}$ involves three parameters, and that we must supply values for all of these parameters in order to interpret a formula $[\mathrm{K}]A$. Therefore, in general and before we make any simplifying assumptions, the function $\mathrm{K}_{m'}$ inputs an e-point $\langle m, h\rangle$ and returns a set of triples of the form $\langle m_1, m_1', h_1\rangle$—that is, a subset of $\mathrm{M} \times \mathrm{M} \times \mathcal{H}$. And the satisfaction conditions for knowledge formulas are these: $\mathcal{M}, m, m', h \models [\mathrm{K}]\phi$ iff for all $\langle m_1, m_1', h_1\rangle \in K_{m'}(\langle m, m', h\rangle)$, $\mathcal{M}, m_1, m_1', h_1 \models \phi$.

We can clarify this general formulation, and at the same time arrive at a more illuminating account of how knowledge should be interpreted in an indeterminist setting, by considering how an agent's knowledge can depend on circumstances. Even though we will not incorporate it in our models, the idea of a *local state* from [Fagin *et al.*, 1995] is helpful here. At each moment, a *global state* determines all the things that are true or false at that moment, whether or not an agent knows them. The things that can bear on an agent's knowledge belong to its local state.

In our theory, an agent's local state can in principle be determined by an e-point $\langle m, h\rangle$ and an assessment moment m', where $m \preceq m'$. If an agent's local state at m_1, m_1' and h_1 is the same as its local state at m_2, m_2' and h_2, then the agent's knowledge will not be able to discriminate between these cases, so a formula $[\mathrm{K}]\phi$ will be true in one case iff it is true in the other one. We make the indeterminist assumption that an agent's knowledge cannot discriminate between genuine future alternatives. This implies that if two histories h_1 and h_2 both pass through m_1', then an agent's local state at m_1, m_1' and h_1 will be the same as the agent's local state at m_1, m_1' and h_2.

The function $K_{m'}$ determines the possibilities that are open to an agent's knowledge at a moment m, from the perspective of an after-the-fact moment m'. It follows from what we just said that if h_1 and h_2 are two different histories passing through m', then $K_{m'}(\langle m, h_1\rangle) = K_{m'}(\langle m, h_2\rangle)$. Since $K_{m'}$ ignores its second argument in this way, we can simplify things by letting $K_{m'}$ be a function from moments to subsets of $\mathrm{M} \times \mathrm{M} \times \mathcal{H}$.

An element of counterfactuality arises in the when the modal logic knowledge is combined with indeterminism. In earlier versions of this paper, including the one published in Context 2007 (which had to be abbreviated due to space limitations) I tried to avoid this issue. But if justice is to be done to the technical and philosophical issues, it can't really be ignored.

We evaluate an agent's knowledge relative to a moment m at which the knowledge state obtains and a future perspective m' from which that knowledge is evaluated. Suppose that $\langle m_1, m_1', h_1\rangle \in K_{m'}(m)$. Then m_1 is an epistemic alternative to m for our agent—a moment that, for all the agent knows, might obtain. Circumstances in the past of m about which the agent has no knowledge, for instance, may turn out differently in m_1. The moment m_1', on the other hand, represents an alternative perspective for assessment, and should somehow correspond to m'.

But since $m_1 \preceq m_1'$ (a perspective on a moment is always a future possiblity of that

moment), m_1' must share the same past as m_1. If moments m and m_1 are incomparable, involving different present circumstances or different pasts, then m_1' cannot be a future possibility for m—it has to be the moment in the future of m_1 that *would* have come about instead of m' if m_1 had been the case. This is where counterfactuality comes in.

Counterfactuality can be combined with indeterminist models of time; see [Thomason and Gupta, 1980]. But combining the two can make for models that are rather complicated. For purposes of presenting the present model theory, I'll make some rather strong assumptions that simplify things somewhat.

First, I'll assume an equivalence relation $\sim$ on moments yielding the moments that could be epistemic alternatives for m. If $\langle m_1, h_1\rangle \in K_{m'}(m)$, then $m \sim m_1$. Second, I'll assume an isomorphism between the futures of similar worlds. That is, where $m_1 \sim m_2$, I'll postulate a one-one, order-preserving function s_{m,m_1} from $\{m' \,/\, m \preceq m'\}$ onto $\{m_1' \,/\, m_1 \preceq m_1'\}$. We can assume that if $\langle m_1, m_1'\rangle \in K_{m'}(m)$ then m_1' corresponds to m'—that is, that $s_{m,m_1}(m') = m_1'$.

These ideas lead to the following revised definition of a frame, and to a satisfaction clause for knowledge along the following lines.

Definition 6. Epistemic frames, models.

An (epistemic) *frame* for indeterministic logic is a structure $\mathcal{F} = \langle \mathrm{M}, \prec, \sim, \mathrm{s}, \mathrm{K} \rangle$, where:

(1) M is a nonempty set (of moments).
(2) $\prec$ is a transitive, antireflexive relation over M such that if $m_1, m_1 \preceq m$, then $m_1 \preceq m_2$ or $m_2 \preceq m_1$ or $m_1 = m_2$.
(3) $\sim$ is an equivalence relation over M.
(4) For each $m, m' \in \mathrm{M}$ such that $m \preceq m'$, $\mathrm{K}_{m'}$ is a set of e-points such that (i) if $\langle m_1, h_1 \rangle \in \mathrm{K}_{m'}$ then $m_1 \sim m_2$ and $h \in \mathcal{H}_{m_1}$, and (ii) for each h passing through m', $\langle m, h \rangle \in \mathrm{K}_{m'}$.
(5) For each m_1, m_2 such that $m_1 \sim m_2$, s_{m_1,m_2} is a one-one function from $\{m' \,/\, m_1 \preceq m'\}$ onto $\{m' \,/\, m_2 \preceq m'\}$, subject to the condition that for all $m_{,1}, m_2, m, m'$, $m \preceq m'$ iff $s_{m_1,m_2}(m) \preceq s_{m_1,m_2}(m)$.

A *model* on a set $\mathcal{P}$ of propositional atoms and a frame $\mathcal{F}$ is a pair $\langle \mathcal{F}, V \rangle$, where $\mathcal{F}$ is a frame and V is a valuation of $\mathcal{P}$ on $\mathcal{F}$.

Definition 7. $\mathcal{M}, m, m', h \models_1 \phi$, $\mathcal{M}, m, h \models_1 \phi$, and $\mathcal{M}, m \models_1 \phi$.

Clauses 1–5 as in Definition 4.

6. **Knowledge:** $\mathcal{M}, m, m', h \models_1 [\mathrm{K}]\phi$ iff for all $\langle m_1, h_1 \rangle \in K_{m'}(m)$, $\mathcal{M}, m_1, s_{m,m_1}(m'), h_1 \models_1 \phi$.

$\mathcal{M}, m, h \models_1 \phi$ and $\mathcal{M}, m \models_1 \phi$ are defined as in Definition 5.

The following example will help to clarify how this semantics works in simple cases. We confine histories to only two points of time; the earlier point at which Mrs. L claims to know that the food at the party will be good, and the time at which the food is served. We will use p to mean 'The food at the party is served, and it is good' and q to mean 'The chef is married'. Mrs. L has no idea at any time whether the chef is married. In the actual initial moment, the chef is married and the party has not yet taken place. In all histories, the chef's marital status remains unchanged. Whether or not the chef is married, there is an outcome in which the food is good, and a (far-fetched) outcome in which it is not good. This model actually is a forest consisting of two unconnected trees, but it is easiest to diagram the knowledge relations by unfolding the four histories, as in Model $\mathcal{M}_2$, below.

The first diagram shows the knowledge relations for Mrs L at m_0, as assessed from m_1. Dashed arrows labeled with 'K_{m_1}' relate m_0 to the e-points that are compatible with what Mrs. L knows, as assessed from a moment m_1 at which the food is served, and is good. In this model, we assume that $K_{m_1}(m_0) = \{\langle m_0, h_1 \rangle, \langle n_0, h_3 \rangle\}$. This is how we represent our intuition that Mrs K knows at m_0 that the food will be good.

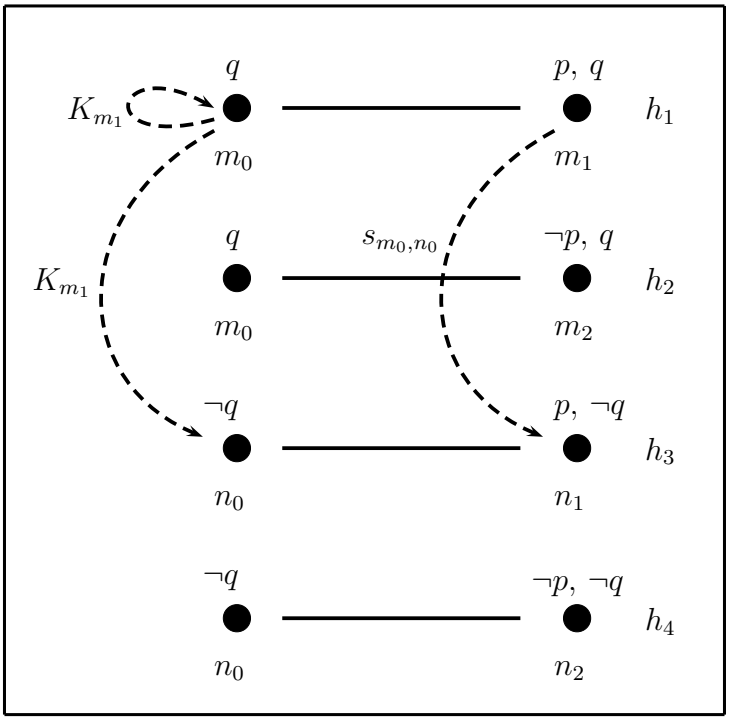

Model $\mathcal{M}_2$, knowledge at m_0 assessed from m_1

The second diagram shows the knowledge relations for Mrs L at m_0, as assessed from m_2. We assume that $K_{m_2}(m_0) = \{\langle m_0, h_1\rangle, \langle m_0, h_2\rangle, \langle n_0, h_3\rangle, \langle n_0, h_4\rangle\}$. From the anomalous perspective of m_2, where the food is not good, we assume that in fact Mrs L knew nothing about whether the food would be good. From every assessment perspective, she knows nothing about whether the chef is married. Whether the food is good and whether the chef is married are independent. Therefore, from this perspective all possibilities are compatible with what Mrs L knew at m_0.

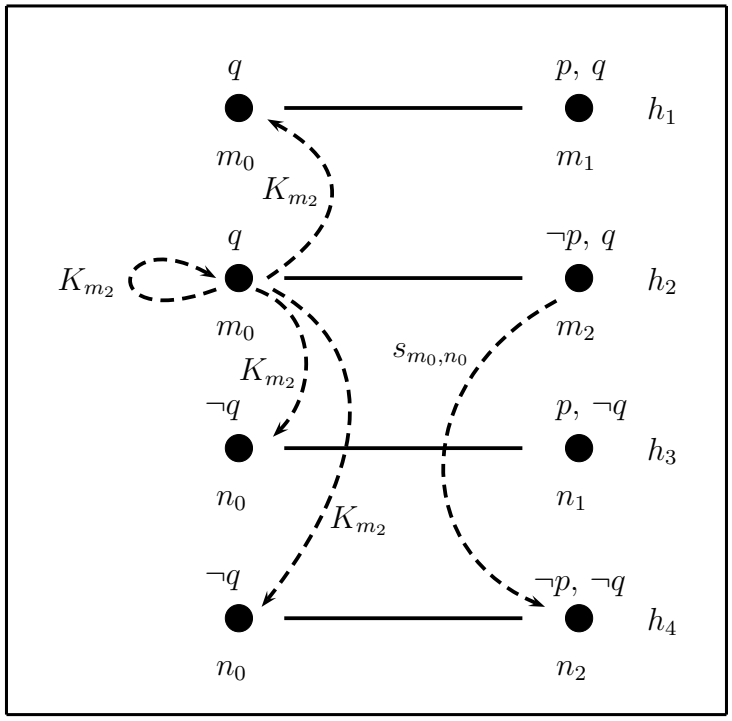

Model $\mathcal{M}_2$, knowledge at m_0 assessed from m_2

According to the satisfaction definition, $\mathcal{M}_2, m_0, m_2, h_2 \models$ [K]⟨F⟩p iff for all $\langle m, h\rangle \in K_{m_1}(m_0)$, $\mathcal{M}_2, m_1, s_{m_0,m}(m_1), h \models_1$ ⟨F⟩p. Therefore, $\mathcal{M}_2, m_0, m_2, h_1 \models$ [K]⟨F⟩p iff (i) $\mathcal{M}_2, m_0, m_2, h_1 \models_1$ ⟨F⟩p, (ii) $\mathcal{M}_2, m_0, m_2, h_2 \models_1$ ⟨F⟩p, (iii) $\mathcal{M}_2, n_0, n_2, h_3 \models_1$ ⟨F⟩p, and (iv) $\mathcal{M}_2, n_0, n_2, h_4 \models_1$ ⟨F⟩p. But, for instance, $\mathcal{M}_2, m_0, m_2, h_2 \not\models_1$ ⟨F⟩p. Thus, $\mathcal{M}_2, m_0, m_2, h_2 \not\models$ [K]⟨F⟩p. As assessed from m_2, Mrs L does not know at m_0 that the food will be good.

I believe that this theory is on the right track (at least, if it is granted that the project of modeling knowledge in an indeterminist setting is appropriate), but this formulation has to be regarded as preliminary. It doesn't incorporate all the appropriate constraints on knowledge, and it may well require fundamental revision. But I do think that something like this account can be maintained, and that it shows how to develop and defend a view of how knowledge of the future can be uncertain.

Note that it is an essential part of this theory that the semantics of an agent's knowledge is not merely a matter of the agent's psychological state. Mrs. L's psychological state at the e-point where the food will be bad is indistingishable from her state at the e-point where it will be good.

3. Belief, Risk, and Impatience

Belief is (apparently) less dependent on context than knowledge. I might be able to convince Mrs. L that she didn't know the food would be good by elaborating scenarios where it turns

out to be bad; but I certainly can't convince her by such arguments that she didn't believe it would be good.

Nevertheless, I think that belief depends on context in other ways, and that sometimes we are able to manipulate these dependencies in order to create or destroy beliefs. I discussed these matters in [Thomason, 1987], so I will be fairly brief here.

Consider a nervous driver at a stop sign at a busy intersection on a dark night. He needs to drive across the intersection. He looks left. A car zooms by from that direction. He looks right. It's clear. He looks left, it's clear. But wait—he can't see what's going on to the right, and doesn't believe it's clear anymore. So he looks right. He repeats the process until he realizes that he'll never get across this way. Time is pressing. But he can't move unless the road is clear. So he lowers his standards, saying to himself "If it was clear to the right a second ago it's clear now." And he hits the gas, hopefully without getting into an accident.

Moral: There are occasions when we can't act without a belief, and in which high standards for belief prevent us from having an appropriate belief. In these cases, an urgent need to act can cause us to lower our standards.

Consider a normal driver approaching an intersection on a dark night. Traffic on the crossroad is very infrequent. There is no sign of a car coming, and as the driver approaches the intersection, she believes there is no car on the other road. But as she get closer she thinks "What if there were—we would crash!" Instantly, her belief disappears, and she brakes her car.

Moral: There are occasions when we have a belief that is well justified, but the consequences of acting on this belief if we are wrong are very harmful. In these cases, we can destroy the belief by changing our standards.

In a theory of practical reasoning where actions are determined by beliefs and desires (rather than by probabilities and utilities) mechanisms of this sort are essential in order to deal with uncertainty and risk.

We can model these effects by supposing that beliefs are local. Rather than appealing to a global, monolithic attitude, we construct belief-like attitudes for the occasion at hand out of a large stock of potential beliefs that can be combined much as we might select and combine propositional axioms for some *ad hoc* purpose. Potential beliefs come with features indicating their provenance and, for instance, the circumstances under which we learned them, and they are sorted according to their entrenchment or plausibility. When we combine potential beliefs into a modality that will guide our actions in a given situation, we can manipulate the beliefs by filtering out less plausible proto-beliefs in the presence of risk, or allowing them in when it is urgent to have a belief of some sort.

The context-dependence that I have indicated in this section is different from the context-dependence I discussed in Section 2. It belongs to a theory of philosophical psychology, or to an agent architecture, rather than to semantics; for semantic purposes, we might well want to relativize the truth of $\textsc{Believe}(a, p)$ to a world w, without distinguishing the features of w that bear on whether $\textsc{Believe}(a, p)$ is true.

However, in understanding how belief can play an effective and flexible role in practical reasoning, it may be important that an agent's belief that p may depend on more than the agent's dispassionate assessment of whether or not p is true. It depends on other psychological factors, and these factors can to some extent be manipulated.

4. Epistemic 'Might'

So-called epistemic modals, and especially the epistemic employment of the modal auxiliary 'might', have attracted attention in the literature on context and epistemology, and the recent literature has disclosed some hitherto unnoticed and interesting features. Here, I want to discuss the apparent dependence of locutions involving 'might' on the speaker's epistemic state. This case was raised as an aside in [Hawthorne, 2004], and has been discussed in [DeRose, 1991], [Egan *et al.*, 2005], [MacFarlane, 2006], and elsewhere. Here I will mainly refer to the account in [Egan *et al.*, 2005].

The examples that are usually mobilized to illustrate the speaker-dependence of 'might' involve overhearers. In the example that [Egan *et al.*, 2005] develop, for instance, Myles (a reporter being interviewed on the radio), asked where Professor Granger is, says "We don't know—she might be in Prague." Professor G, listening to the broadcast while on a beach somewhere in the South Pacific, denies that she might be in Prague. You can elaborate this: if Prof. G calls the station and tells Myles where she is, he will have to admit that he was wrong when he said she might be in Prague.

Various theories have been advanced to account for this. Egan, Hawthorne and Weatherson favor a theory proposed earlier by DeRose (but without some of De Rose's added qualifications, so that the relationship between the speaker and the relevant group is somewhat loose):

> S's assertion of 'it is possible that P' is true if and only if (1) no member of the relevant community knows that P is false and (2) there is no relevant way by which members of the relevant community can come to know that P is false. [DeRose, 1991][pp. 593–594]

I find this account, with its multiple 'relevant's, too vague. As we will see, there are cases when, for practical purposes, there is no way the individual members of the community can come to know P, but, in the appropriate sense, the group knows P. We will need to see how these cases interact with epistemic 'might'.

I will suppose along with many other authors, but without arguing for it, that 'might' is connected as in DeRose's analysis with some sort of group knowledge. The problem, given a group G and knowledge operators $\texttt{[K,a]}$ for each agent $a \in G$, to define the relevant group epistemic operator $\texttt{[K,G]}$.

DeRose's formulation is close to what is usually called *distributed knowledge*; see [Fagin *et al.*, 1995][p. 24]. The definition is this. (I assume a possible worlds semantics that interprets the individual modalities $\texttt{[K,a]}$ by means of a relation R_a over worlds.)

Definition 8. $\mathcal{M}, w \models \texttt{[K,G]}\phi$.
$\mathcal{M}, w \models \texttt{[K,G]}\phi$ iff for all w' such that $\langle w, w' \rangle \in \bigcap_{a \in G} R_a$, $\mathcal{M}, w' \models \phi$.

This amounts to saying that the group d-knows p if p follows from pooled or combined knowledge of the members of G. In fact, a corollary of Definition 8 is that $\texttt{[K,G]}\phi$ is true in a model if and only if there are formulas $\psi_1, \ldots, \psi_n$ and group members $a_1, \ldots, a_n$ such that $\texttt{[K,}a_i\texttt{]}\psi_i$ is true for each i, $1 \leq i \leq n$, and ϕ is a logical consequence of $\psi_1 \wedge \ldots \wedge \psi_n$. It does *not* assume, however, that the members of G can pool their knowledge freely; there may be constraints on their ability to communicate.

This definition eliminates an element of indeterminacy in DeRose's account. Also, I think, it is more faithful to the facts.

In the rather brief space available, I want to indicate how various hypotheses about 'might' could be explored by developing examples in which the reasoning activities of the participants are brought into the picture. The purpose is to make a plausible case for the fruitful use of such examples.

First, let's return to the Brown Corpus. Epistemic 'might' occurs rather frequently in this corpus (well over 100 instances). The following three examples are representative.

(2.1) Benington recalled that he once told Hartweger that he doubted Gordon would ever play much for him because he seemed to be lacking in all of the accepted basketball skills. After the coach listed all the boy's faults, Hartweger said, "Coach before I leave here, you'll get to like me". Mrs. Benington admired Gordon's spirit and did what she could to persuade her husband that the boy might help the team.

(2.2) The weather bureau has estimated that radioactive fallout from the test might arrive here next week.

(2.3) The heightened tension, in fact, had been a major factor in the President's change of view about the urgency of a meeting with the Soviet leader. He was not going to Vienna to negotiate—the simultaneous announcements in Washington and Moscow last week stressed that no formal negotiations were planned. But Mr. Kennedy had become convinced that a personal confrontation with Mr. Khrushchev might be the only way to prevent catastrophe.

In each of these examples, there is reasoning in the background, either by an individual or a group.

Mrs. B and Mr. B, as a group, were debating the worthiness of various candidates for the basketball team. At the beginning of the debate, we can assume, Mr. B claimed to know the boy would not help the team. In persuading him he was wrong (which in fact, the passage goes on to say, she does), she got him to admit that the boy might help the team. At the beginning of the debate, Mr. B denied MIGHT(p). At the end, he accepted MIGHT(p).

The weather bureau constitutes a deliberating group in Example (2.2). At the outset, there may have been many possibilities: perhaps the fallout will arrive this week, or any of the three subsequent weeks. They gather data and make projections. Week 1 and week 4 are ruled out, and (let us suppose) week 2 gets a less than 40% chance, while week 3 is deemed more likely. The newspaper reports the result of the deliberations using the earlier date, to simplify and perhaps to provide a more impressive warning.

President Kennedy and his advisors deliberated, with the heavy responsibility of avoiding a nuclear war without making the United States seem willing to be influenced by threats. They considered many options. One by one, they discovered powerful arguments against all but one of them. At that point, it seemed to them that there might be only one alternative.

These simple examples make it clear that the status of 'might' can be influenced in the course of a group deliberation. If 'might' is a group epistemic modality, it is one that changes and that sometimes—as in the case of Example (2.1)—can be explicitly negotiated in the course of a dialogue. I conclude that, in presenting examples as data for theories of 'might', we must be very careful to provide enough detail so that we know with exactly which point in the reasoning the 'might' is to be associated.

In Example (2.1), Mr. and Mrs. B begin their negotiation by disagreeing about whether Hartweger might help the team. At the end, Mr. B (let us say) is convinced he was wrong; he sincerely says "I was mistaken; actually, that boy might help the team." At an intermediate point, before he was persuaded but after he began to feel the force of Mrs. B's argument,

he may have been genuinely uncertain about whether Hartweger might help the team. This example shows, then, that the members of G may disagree about $\texttt{[K,G]}p$, and that they can be uncertain about whether $\texttt{[K,G]}p$. Such facts, of course constrain theories of the meaning of epistemic 'might'.

As a tool in exploring the contours of 'might', I suggest the use of dialogues involving the group solution of constraint satisfaction problems and, in particular, of word puzzles.

Imagine, for instance, a father who has already solved a crossword puzzle and who is helping his young daughter to a solution. The word at $\langle 1, \text{Down}\rangle$ has the clue "a large feline, found in Central America." There are six letters for this word. The girl says "It might be 'jaguar'." The father says "Yes, it might. But look at $\langle 1, \text{Across}\rangle$." The clue is "A small house, often used as a summer home." "Oops," says the girl. "That's 'cottage'—I know that. Are there cougars in Central America?" "Yes," says the father. "I had to check that." "Then the word can't be 'jaguar'," the girl says. "It must be 'cougar'."

Here, less cooperatively, the father could have said "No it couldn't be 'jaguar'. Did you look at $\langle 1, \text{Across}\rangle$?" In the former case, I would say that the relevant operator isn't always knowledge; the father is suspending his knowledge to participate as a partner in the daughter's problem-solving. Egan, Hawthorne and Weatherson consider a similar, but less elaborated example and conclude that the speaker may not always be a member of the group.

The systematic use of information-seeking conversations in relatively structured domains can yield useful evidence. Consider, for instance, the following dialogue.

Round 1.	M to A:	"The mystery word is a noun."
Round 1.	M to B:	"It has 3 letters."
Round 1.	A to all:	"It might be 'mother'."
Round 1.	B to all:	"It can't be 'mother' but it might be 'the'."
Round 1.	A to all:	"OK, it might be 'car'."
Round 1.	B to all:	"Yes, It might be 'car'."
Round 2.	M to A:	"It begins with 't'."
Round 2.	M to B:	"It ends with 'p'."
Round 2.	A to all:	"It might be 'time'."
Round 2.	B to all:	"No, but it might be 'cap'."
Round 2.	A to all:	"No, it can't be 'cap'. But it might be 'tap'."
Round 2.	B to all:	"Yes, it might be."
Round 3.	M to all:	"That's right, game's over."

Here, three agents—two guessers A and B, and a moderator M whose job is to reveal constraints—are playing a guessing game. The rules of the game prevent A and B from directly communicating constraints known to them. A can't overhear M's communications to B and B can't overhear M's communications to A. Dialogues of this sort can be used to address issues such as whether a distributed knowledge analysis is more faithful to the facts than DeRose's. But they can be a fruitful and relatively precise instrument, I think, for exploring many other issues having to do with epistemic modals.

5. Conclusion

Although I have been rather brief, I hope that I have managed to make a convincing case for my main points: (1) modeling methods from logic can be helpful in exploring issues in the philosophy and semantics of context and epistemic locutions, and (2) in combination with examples designed to bring the reasoning that 'might' is tracking into prominence,

these methods may be able to lead us to improved theories of this interesting and important cluster of phenomena.

References

[Belnap, Jr. *et al.*, 2001] Nuel D. Belnap, Jr., Michael Perloff, and Ming Xu. *Facing the Future: Agents and Choices in Our Indeterminist World*. Oxford University Press, Oxford, 2001.

[DeRose, 1991] Keith DeRose. Epistemic possibilities. *The Philosophical Review*, 107(2):581–605, 1991.

[Egan *et al.*, 2005] Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In Gerhard Preyer and Georg Peter, editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, pages 131–168. Oxford University Press, Oxford, 2005.

[Fagin *et al.*, 1995] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. The MIT Press, Cambridge, Massachusetts, 1995.

[Hawthorne, 2004] John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, Oxford, 2004.

[Kaplan, 1978] David Kaplan. On the logic of demonstratives. *Journal of Philosophical Logic*, 8:81–98, 1978.

[Lewis, 1979] David K. Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359, 1979.

[MacFarlane, 2002] John MacFarlane. Future contingents and relative truth. *The Philosophical Quarterly*, 53(212):321–336, 2002.

[MacFarlane, 2006] John MacFarlane. Epistemic modals are assessment-sensitive. Unpublished manuscript, University of California at Berkeley, 2006.

[Preyer and Peter, 2005a] Gerhard Preyer and Georg Peter, editors. *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press, Oxford, 2005.

[Preyer and Peter, 2005b] Gerhard Preyer and Georg Peter, editors. *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press, Oxford, 2005.

[Thomason and Gupta, 1980] Richmond Thomason and Anil Gupta. A theory of conditionals in the context of branching time. *The Philosophical Review*, 80:65–90, 1980.

[Thomason, 1970] Richmond H. Thomason. Indeterminist time and truth-value gaps. *Theoria*, 36:246–281, 1970.

[Thomason, 1987] Richmond H. Thomason. The multiplicity of belief and desire. In Michael P. Georgeff and Amy Lansky, editors, *Reasoning about Actions and Plans*, pages 341–360. Morgan Kaufmann, Los Altos, California, 1987.

[Thomason, 2007] Richmond H. Thomason. Truth after the fact in indeterminist tense logic. Unpublished manuscript, Philosophy Department, University of Michigan. Available from the author, 2007.

[van Fraassen, 1969] Bas C. van Fraassen. Presuppositions, supervaluations, and free logic. In Karel Lambert, editor, *The Logical Way of Doing Things*, pages 67–91. Yale University Press, New Haven, Connecticut, 1969.