

Dynamics of Epistemic Modality

DRAFT AS OF OCTOBER 1, 2009. DO NOT CITE WITHOUT PERMISSION.

Abstract

A dynamic semantics for epistemically modalised sentences is an attractive alternative to the orthodox view that our best theory of meaning ascribes to such sentences truth-conditions relative to what is known. I will demonstrate that a dynamic story about *might* and *must* offers elegant explanations of a range of puzzling observations about epistemic modals. The first part of the story offers a unifying treatment of disputes about epistemic modality and disputes about matters of fact while at the same time avoiding the complexities of other theories that are currently on the market. The second part of the story extends what has been said to cover some further relevant data, including retraction, the interaction between epistemic modality and tense, and embeddings of epistemically modalised sentences under attitude verbs. An in-depth comparison between the suggestion made in this article and current versions of the orthodoxy is provided.

1 Introduction

Orthodox semantics assigns to epistemically modalised sentences truth-conditions relative to what is known. It is a well-worn story that orthodoxy has to become quite extravagant if it wants to do justice to how such sentences are used in discourse. For what is known varies from speaker to speaker, and ordinary people have the nasty habit of evaluating present tense claims of epistemic modality by testing them against their own perspective. Consider the example from [von Fintel and Gillies \(2008b\)](#). Mary is looking for her keys, Alex is trying to help:

- (1) Mary: I can't find my keys.
Alex: They might be in the car.
Mary: No, they can't be in the car. I still had them with me when I came in.

The intuition is that Mary denies what Alex has asserted, namely that the keys might be in the car. But this is not the result we get if we interpret Alex and Mary as reporting on what is or is not compatible with what they (respectively) know.

One extraordinary version of the orthodoxy is that epistemically modalised sentences vary in truth-value across points of assessments (a judge or some other sequence which varies with the assessor). So Alex’s utterance is literally true when he is the judge and false when it is Mary (and vice versa for Mary’s utterance).¹ Another radical position is that we have to give up the idea of a unique contextually determined proposition expressed by a judgement of epistemic modality. Alex’s utterance has a solipsistic reading but in addition affects the discourse by ‘putting into play’ the proposition that the keys being in the car is compatible with what Mary knows. And once we have reinvented the pragmatics of assertion and denial, we can predict that this discourse effect legitimates Mary’s denial of Alex’s utterance.² These are evidently complex views, too complex to be efficiently dismantled here (see §5 for some criticisms, though).³ But they are also complex *enough* to make one wonder whether orthodoxy itself is the problem, and how much simpler life could be without it.

It is, to say the least, a bit surprising that the dispute in (1) is problematic. From a naive perspective, the case seems completely innocent. Mary rejects Alex’s judgement on the basis of what she knows. She knows that the keys cannot be in the car, and this is why she denies what Alex has asserted. From this perspective, the dispute about where the keys might be does not really differ from the following dispute about matters of fact:

- (2) Mary: Where are my keys?
Alex: They are in the car.
Mary: No, they are not. I still had them with me when I came in.

Here again one wants to say that Mary evaluates Alex’s claim against her own perspective. She knows that the keys are not in the car and, as a result, denies what he has asserted. Variation in what is known leads to different assessments of Alex’s claim—Alex thinks it is right, Mary thinks it is wrong—and this is just how things should be.

In the second scenario, Alex’s and Mary’s epistemic situations differ, and this explains why they disagree. What is important is that the difference only matters for how they *assess* the claim that the keys are in the car, not for what they *say* when they discuss the whereabouts of the keys. This is why Mary can evaluate Alex’s judgement against her own perspective without risk of misinterpreting what he said. The naive perspective assumes that differences in what is known play exactly the same role when it comes to judgements of epistemic modality. Everybody, I think, can agree that the naive conception is very attractive: it provides a simple and uniform perspective on disputes about epistemic modality and disputes about matters of fact and creates no need for relativism or new age contextualism.

¹See Egan et al. (2005); Egan (2007), Yalcin (2007), Stephenson (2007, 2008), MacFarlane (2009a).

²See von Stechow and Gillies (2009).

³I have omitted a view which has served as the whipping boy in modern debates about epistemic modality, namely that epistemically modalised sentences are always used to describe what is known by a contextually salient group of individuals (in most cases involving the speaker). For critical discussion see Egan et al. (2005), §2 and MacFarlane (2009a), §3.

Highlighting the naive perspective on disputes about epistemic modality is important because it helps us appreciate *why* the orthodox view leads to so many complexities. On the orthodox view, the differences in what is known matter for what Alex and Mary *say* when they discuss where the keys might be. What Alex says is true just in case the keys being in the car is compatible with what *he* knows: Mary seems to misinterpret what Alex just said when she rejects his claim on the basis of what *she* knows. It now becomes a nontrivial task to get the facts about modal disagreement straight. Relativists isolate a content common to Alex's and Mary's judgement which is true given Alex's epistemic situation and false from Mary's perspective. New age contextualists enrich what Alex has said so that he also makes a wrong judgement about Mary's epistemic situation, thus legitimatising Mary's rejection of Alex's claim on the basis of what she knows. But nothing of this would be necessary if we were to accept the naive view that one's epistemic situation only matters for how one assesses a judgement of epistemic modality, and not for what one says when one makes such a judgement.

What I have done so far is to contrast two possible views about the role of one's epistemic situation for judgements of epistemic modality: the naive conception which holds that one's epistemic situation only matters for how one assesses a judgement of epistemic modality, and the orthodox conception which holds that one's epistemic situation is also relevant for what one says when one makes such a judgement. I have also outlined some initial motivations for adopting the naive conception. In the remainder of this article I will show that the naive conception is indeed a viable position. My suggestion then is to throw orthodox semantics for epistemically modalised sentences overboard and see how we can live without it. True, no one knows how to give sane truth-conditions for epistemically modalised sentences without signing up for the orthodoxy, but this just means that we should not give them truth-conditions in the first place. Our semantic story should run under the slogan that epistemically modalised sentences have content, but not truth-conditional content. I will demonstrate that such a story can be told.

My plan is as follows. In §2 I will tell just enough of my story to offer a simple and uniform perspective on our two cases and to dispel some old prejudices against non-truth-conditional approaches to the semantics of epistemic modals. In §3 I throw in some more data every adequate theory of epistemically modalised sentences has to take care of. What has been said in §2 will readily take care of some of the issues, but minor extensions will be necessary. The remaining work is done in §4. Finally, in §5, I offer a more in-depth comparison between the framework developed in this article and the orthodox perspective.

2 Basics

The protagonists of the story I will tell in this section are epistemically modalised sentences, and since we do not care too much about sentential structure, the initial language is confined to classical propositional language, but extended with the epistemic possibility operator (\diamond) and the epistemic necessity operator (\square):

Definition 1 (Base Language) \mathcal{L} is the smallest set that contains any sentential atoms $\mathcal{A} = \{p, q, \dots\}$ and is closed under negation (\neg), conjunction (\wedge), and the epistemic modal *might* (\diamond). \mathcal{L}_0 is defined as the non-modal fragment of \mathcal{L} . Disjunction (\vee), the material conditional (\supset), and the epistemic modal *must* (\square) are defined in the usual way.

The framework will later be extended to cover various embeddings of epistemically modalised sentences, including tense operators and belief ascriptions.

The semantics I intend to develop is dynamic in that it understands the semantic value of a sentence as its Context Change Potential (CCP).⁴ One way of motivating a dynamic perspective on meaning and communication starts with a familiar picture about context-content-interaction. Take the truisms about assertions from [Stalnaker \(1978\)](#): assertions express propositions and are made in a context; in fact, context and what is said frequently affect each other. Since language has context-sensitive expressions, which proposition the assertion expresses may very well depend on the context. On the other hand, assertions in turn affect the context, and they do so by adding the proposition expressed by that assertion to the context.

In Stalnaker’s picture all context change is mediated by propositional content, and so we may happily maintain that the primary task of a semantic theory consists in assigning a truth-condition determining proposition to each declarative sentence of a given language. But the picture also suggests a *change* of perspective: instead of being all about truth-conditions, a semantics may instead be all about how an utterance relates an input context (the context in which it is made) to an output context (the context posterior to the utterance). Meanings then become fully relational: they are relations between contexts.⁵ Of course, and hopefully to the traditionalists’ relief, *some* context change may be mediated by propositional content, but there is no commitment to the claim that *all* context change is thus mediated. It is this lack of commitment that I intend to exploit in my story about epistemic modality. Epistemically modalised sentences have CCP, but the context change is not mediated via propositional content. So it is in this sense that claims of epistemic modality have content, but no truth-conditional content.

With so much focus on contexts, it is legitimate to ask what exactly they are supposed to be. There is no uniform answer to that, but [Veltman \(1985, 1996\)](#) is most naturally understood as treating contexts as information states. I will follow suit, not least because my proposal bears some resemblance to his Update Semantics.⁶ Such a dynamic approach is not only interesting due to its lack of

⁴Some popular dynamic semantics: Discourse Representation Theory ([Kamp \(1981\)](#); [Kamp and Reyle \(1993\)](#); [Kamp et al. \(2009\)](#)), File Change Semantics ([Heim \(1982\)](#)), Update Semantics ([Veltman \(1985, 1996\)](#)) Dynamic Predicate Logic ([Groenendijk and Stokhof \(1991\)](#)).

⁵This way of motivating dynamic semantics – by moving from a familiar picture about context-content-interaction to a purely relational view about semantic values – is, plus minus a bit, the route taken by [von Stechow and Gillies \(2008a\)](#), §6 or [Dever \(2009\)](#), §1.

⁶The differences between the upcoming proposal and Update Semantics are substantial, primarily because the latter cannot explain why utterances of epistemically modalised sentences can be informative for the audience. So you may wish to take a look at Veltman’s work in case you are interested in the original. See also [Groenendijk et al. \(1996\)](#), [van der Does et al. \(1997\)](#), [Beaver \(2001\)](#), and [Gillies \(2001\)](#) for some applications and extensions of Veltman’s proposal.

commitment to ascribing truth-conditions to every sentence of our toy language. It also offers an appealing perspective on our observation about how utterances of epistemic modality are assessed by different people: since the effect of an utterance may vary from information state to information state, we should expect that speakers may have different reactions to one and the same utterance, let it be factual or modal. Thus the story I wish to tell will be in the position to offer a simple and uniform account of how speakers assess utterances in discourse.

It is common to model information states as sets of possible worlds compatible with what is known by the agent. This view has a problem, but it is not the problem you think it has. The problem of logical omniscience – that if an agent knows ϕ , then the agent also knows all the logical consequences of ϕ – is deep and deserves further analysis, but not here. The problem I have in mind is that if we just consider the set of possible worlds compatible with what is known and call those the agent’s epistemic possibilities, then all epistemic possibilities are created equal. This may often be unproblematic, but it is inappropriate for current purposes. Intuitively, an agent who learns that so-and-so might be the case becomes aware of a certain epistemic possibility – something which has been previously merely compatible with what is known has become a ‘live’ epistemic possibility. To capture this intuition, our information states need to be fine-grained enough so that we can distinguish between what is merely compatible with the agent’s information and what is a live epistemic possibility – a possibility which the agent is aware of.

There are different ways to model the distinction between mere epistemic possibilities and live epistemic possibilities, but not all of them will serve our purpose. Consider, for example, a view which postulates a single set of possible worlds, but adds a partition of this set as a bit of additional structure.⁷ From a dynamic perspective, the role of claims of epistemic possibility would then be similar to the role of questions in that they raise issues. Such operations do not embed very well, however, and so I will choose a different approach here.⁸ The idea is that information states are reconstructed as sets of sets of possible worlds. p is an epistemic possibility in an information state Σ if there is at least one set of possible worlds in Σ containing a p -world (a world at which p is true). p is a settled or live epistemic possibility in an information state Σ if and only if every set of possible worlds in Σ has a p -world in it, provided there is at least one set of possible worlds in Σ . So we have epistemic possibilities that are not live epistemic possibilities, and on top of that preserve the simple and elegant structure of a possible worlds model. This will bring us much happiness once we take care of the technical details.

This proposal is also the one favoured by [von Fintel and Gillies \(2008a\)](#), but does not play a role in their later work on epistemic modality.

⁷This is the approach taken by [Yalcin \(2009\)](#) and good enough for his purposes. But these purposes do not include telling a dynamic story about *might* and *must*, and so what works for him may not work for me.

⁸Considerable progress has been recently made towards a semantics which allows for embeddings of questions, e.g. by [Asher \(2007\)](#), [Groenendijk \(2008\)](#), and [Asher and Lascarides \(2009\)](#). However, Groenendijk’s framework does not allow for simple negations of questions, and on the other proposals a question like ‘Are the keys in the car?’ relates to its negation ‘Aren’t the keys in the car?’ in a way which is quite different from the way ‘The keys might be in the car’ relates to ‘The keys can’t be in the car’. Accordingly, it is hard to see how any of these proposals could be of help when we try to give a semantics for epistemically modalised sentences.

Definition 2 (Possible Worlds, States, Information States) w is a *possible world* iff. $w: \mathcal{A} \mapsto \{0, 1\}$. W is the set of such w 's. σ is a *state* iff $\sigma \subseteq W$. S is the set of such σ 's. Σ is an *information state* iff $\Sigma \subseteq (S \setminus \emptyset)$, i.e. an information state is a (possibly empty) set of non-empty sets of possible worlds. I is the set of such Σ 's. The *initial* information state Σ_0 is identical with $(S \setminus \emptyset)$, the *absurd* information state Σ_\emptyset with \emptyset .

All of this is based on a quite familiar idea: start with the classical conception of an information state as a set of possible worlds. A possible world is a complete answer to the question ‘What is the case?’, and one often says that p is settled in an information state just in case every world in it is p -world. Similarly, we can think of a set of possible worlds as a complete answer to the question ‘What might be the case?’, and understand an information state as a set of differing answers to that question. And of course, for this picture to make sense there need not be a fact of the matter about what might or might not be the case – there is no need for some piece of reality to be captured by a set of possible worlds. But that does not mean that information states do not represent at all. Given an information state Σ , we may legitimately ask whether the actual world w_a is an epistemic possibility, and this is just to ask whether $w_a \in \bigcup \{\sigma: \sigma \in \Sigma\}$. If it is, then Σ is *truthful*, as the information encoded in it does not rule out the actual world; otherwise we can say that Σ *misrepresents* reality. In other words, the present conception of information states as sets of sets of possible worlds allows us to say everything the classical conception allows us to say, and a bit more.

So intuitively speaking, p is a live epistemic possibility for an agent S who is in an information state Σ if and only if, whenever the agent considers how things might completely be, S will include p . And when this is not so yet S has no information ruling out p , then p is just a plain epistemic possibility. So much for the informal characterisation. Here is the more formal description of the two ways in which p may be an epistemic possibility for an agent:

Definition 3 (Possibilities (Special Case)) Consider any $\Sigma \in I$ and $p \in \mathcal{A}$:

1. p is an *epistemic possibility* in Σ iff $\exists w \in \bigcup \Sigma: w(p) = 1$
2. p is a *live epistemic possibility* in Σ iff $\exists w \in \bigcup \Sigma: w(p) = 1$ and
 $\forall \sigma \in \Sigma \exists w \in \sigma: w(p) = 1$

There are information states in which p is an epistemic possibility – p is not excluded by the agent’s information – but p is not a live epistemic possibility. In such an information state there is a state without a p -world. Clearly, in Σ_0 every sentential atom is a possibility, but no sentential atom is a live epistemic possibility. On the other hand, Σ_\emptyset does not encode any epistemic possibilities at all – we have simply excluded too much. To give a case in more detail, consider a simple language restricted to the sentential atom p . Then there are two possible worlds $w_1(p) = 1$ and $w_2(p) = 0$. Consider the information state $\Sigma = \{\{w_1\}, \{w_2\}, \{w_1, w_2\}\}$. Then p is a possibility – the information state does not exclude p – but it is not a

live epistemic possibility. This distinction is the source of the informativity of judgements of epistemic modality. If an agent learns that p might be the case, the agent will exclude all those states in which there is no possible world verifying p . Going back to our toy example, the agent excludes $\{w_2\}$ from Σ , resulting in an information state in which p is a live epistemic possibility. In other words, the role of judgements of epistemic possibility is to raise a certain epistemic possibility into salience. On the other hand, if an agent learns that p cannot be the case, the agent excludes all those states in which there is a possible world verifying p , i.e. $\{w_1\}$ and $\{w_1, w_2\}$. Evidently, the result is just the complement of the result of learning that p is an epistemic possibility. This will allow us to capture the intuitions involving utterances of epistemically modalised sentences in a very elegant way.

As a preparation for how information states are modified in the light of new information, we lay out how formulas in \mathcal{L} affect the elements of information states. Once this is in place, we will define updates of information states in terms of updates on their elements.

Definition 4 (Updates on States) Consider any $\sigma \in S$, $p \in \mathcal{A}$ and $\phi, \psi \in \mathcal{L}$. An update on a state is a function $\uparrow: S \mapsto S$ defined by the following recursion:⁹

- (1) $\sigma \uparrow p = \{w \in \sigma: w(p) = 1\}$
- (2) $\sigma \uparrow \neg\phi = \sigma \setminus (\sigma \uparrow \phi)$
- (3) $\sigma \uparrow \phi \wedge \psi = (\sigma \uparrow \phi) \uparrow \psi$
- (4) $\sigma \uparrow \diamond\phi = \{w \in \sigma: \sigma \uparrow \phi \neq \emptyset\}$

The clause in (1) requires that updating a state σ with an atom p eliminates all possibilities from σ in which p is false. According to clause (2), updating with $\neg\phi$ comes down to leaving everything in σ which gets eliminated by an update with ϕ . To update a state with a conjunction, update with the first conjunct and then update the resulting state with the second conjunct (cf. (3)). Clause (4) captures a *test*-conception of claims of epistemic modality. Updating of σ with a formula of the form $\diamond\phi$ is to run a test on the state: if updating σ with ϕ does not return the empty set, the state passes the test. Otherwise, we get back the empty set, and we do not need to consider the state at a later stage. For example, updating σ with $\diamond p$ returns σ if there is at least one p -world in σ , and returns the empty set otherwise.¹⁰

We can now describe which impacts formulas of \mathcal{L} have on an information state. In the spirit of the present proposal, we will call a full description of the impacts of a formula ϕ on information states its meaning. The purpose of a semantics is then not to assign a truth-condition determining proposition to each formula in \mathcal{L} . Instead, semantic values are fully relational: they are relations between information states.

⁹If necessary, the update effects of elements of \mathcal{L} can be relativised to contexts of utterance. But it is not necessary now, as epistemic modals are not context-sensitive. Context-sensitivity will become an issue once tense enters the stage in §4.

¹⁰These rules, and in particular the test-conception of epistemically modalised sentences, lie at the heart of Veltman's Update Semantics. But as we will see, how updates on states work is only one part among many others in a complete dynamic story about *might* and *must*.

Definition 5 (Updates on Information States) Consider any $\Sigma \in I$ and $\phi \in \mathcal{L}$. An update on an information state is a function $[\cdot] : I \mapsto I$ defined as follows:

$$\Sigma[\phi] = \{\sigma : \sigma \neq \emptyset \wedge \exists \sigma' \in \Sigma : \sigma' \uparrow \phi = \sigma\}$$

Update of an information state Σ with a formula ϕ thus comes down to the following procedure: first update every element of Σ with ϕ ; then gather all the resulting states together, leaving out the empty state. This gives you the output information state. We can now generalise the notion of a (live) epistemic possibility as introduced in the third definition:

Definition 6 (Possibilities (General Case)) Consider any $\Sigma \in I$ and $\phi \in \mathcal{A}$:

1. ϕ is an *epistemic possibility* in Σ iff $\Sigma[\phi] \neq \emptyset$
2. ϕ is a *live epistemic possibility* in Σ iff $\Sigma[\phi] \neq \emptyset$ and $\forall \sigma \in \Sigma : \sigma \uparrow \phi \neq \emptyset$

The following notions will turn out useful in later discussion.

Definition 7 (Settledness, Admission, Entailment) Let Σ be an information state and ϕ, ψ be formulas:

1. Σ *supports* ϕ , ϕ is *settled* in Σ , $\Sigma \models \phi$, iff $\Sigma[\phi] = \Sigma$
2. Σ *admits* ϕ , $\Sigma \triangleright \phi$, iff $\Sigma \not\models \phi$ and $\Sigma \not\models \neg\phi$
3. ϕ *entails* ψ , $\phi \models \psi$, iff $\forall \Sigma : \Sigma[\phi] \models \psi$

There are three possible relations between a $\Sigma \in I$ and $\phi \in \mathcal{L}$: $\Sigma \models \phi$, $\Sigma \triangleright \phi$, or, the unfortunate case, $\Sigma[\phi] = \emptyset$.

We are now in the position to lay out a couple of important features of the current framework. First, we find that a non-absurd information state supports $\ulcorner \diamond\phi \urcorner$ just in case ϕ is a live epistemic possibility:

Fact 1 Given non-absurd Σ , $\Sigma \models \diamond\phi$ iff ϕ is a live epistemic possibility in Σ

Since $\Sigma[\diamond\phi] \models \diamond\phi$, we get the desired result that an admissible update with an epistemically modalised sentence raises the prejacent to a live epistemic possibility. Specifically, any admissible update with $\ulcorner \diamond\phi \urcorner$ raises ϕ from something merely compatible with what is known to a live epistemic possibility.

Judgements of epistemic modality and factual judgements stand in the right logical relations with each other. On the one hand, once one has accepted some ϕ from our non-modal fragment \mathcal{L}_0 of \mathcal{L} , one is also committed to the corresponding judgement of epistemic necessity:

Fact 2 For all $\phi \in \mathcal{L}_0$: $\phi \models \Box\phi$

As an immediate consequence, once $\phi \in \mathcal{L}_0$ is settled in Σ , there cannot be any doubt about ϕ , i.e. the possibility of ϕ being false is excluded from Σ .

On the other hand, an agent for whom ϕ is a live epistemic possibility is in no way committed to ϕ being true:

Fact 3 $\diamond\phi \not\models \phi$

If ϕ is a live epistemic possibility in Σ , Σ won't admit the judgement that ϕ *must* be false. But it is not excluded that ϕ turns out to be false. *Might* is non-factive. And since *might* and *must* are duals by design, it follows immediately that Contraposition fails. That is the right result when we have a dynamic consequence relation: once your information state supports ϕ it also supports $\lceil \Box\phi \rceil$; but it does not follow that $\lceil \neg\phi \rceil$ is settled whenever the falsity of ϕ is a live epistemic possibility.

Finally, we find that the current framework validates the characteristic axiom of S5:

Fact 4 $\diamond\phi \models \Box\diamond\phi$

Thus any combination of boxes and diamonds reduces to the innermost operator. This is an acceptable result, I think, as it is controversial that embedding an epistemic modal under another epistemic modal has any interesting semantic effects.

This is all I have to say about the semantics of \mathcal{L} , so I might as well move on to the pragmatics.¹¹ Things are appealingly simple at this point. In ordinary circumstances, an assertion of ϕ is made with the intention that the hearers update their information states with the content, i.e. the CCP, of ϕ . How hearers assess an utterance of ϕ will then depend on the update effects of ϕ on their information states. Thus we have a two-stage picture: first determine how well ϕ plays with the current information state and then assess the assertion of ϕ on that basis.

Assessment Let $\phi \in \mathcal{L}$ and consider a subject A with information state Σ_A . Then A will by default assess an utterance of ϕ as follows:

- **Agree** in case $\Sigma_A \models \phi$
- **Admit** in case $\Sigma_A \triangleright \phi$
- **Reject** in case $\Sigma_A[\phi] = \emptyset$

In words, A will agree with ϕ if A 's information state already encodes the information encoded in ϕ . If A 's information is incompatible with ϕ , then we

¹¹At least that is all I have to say without digressing too much. One thing that certainly deserves further discussion – especially when one is in a dynamic mood – is the assumption that *might* and *must* are duals. The fact that *might* and *must* behave differently when it comes to modal subordination (Roberts (1987, 1989)) speaks heavily against this view. But all this is irrelevant for current purposes, and the story could be extended to cover modal subordination, adopting some of the things said by Asher and McCready (2007).

expect A to reject an assertion of ϕ . And finally, if A is agnostic about ϕ , then A might as well admit that ϕ is the case.

I have promised a simple and uniform perspective on Alex’s and Mary’s disputes about the keys, and I can now deliver. First example: Mary denies what Alex has asserted, as the content of her assertion – that the keys can’t be in the car – is just the negation of what Alex has said. The difference in what is known does not matter for *what* they say when they make their utterances, but for *why* they say it. Let p be short for ‘The keys are in the car’, and let Σ_A and Σ_M be Alex’s and Mary’s information states, respectively. The crucial difference between Alex and Mary is that $\Sigma_A \models \Diamond p$ but $\Sigma_M \models \neg p$ and thus, by **Fact 2**, $\Sigma_M[\Diamond p] = \emptyset$. So variation in what is known fully explains why Mary denies what Alex has asserted, and does nothing to threaten our ability to predict that there is a dispute in the first place. And we can say the same about the second example, the only difference being that here $\Sigma_A \models p$, i.e. Alex accepts that the keys are in the car. Again, Mary denies what Alex has asserted, and this is because $\Sigma_A \models p$ but $\Sigma_M[p] = \emptyset$.

What makes all of this possible is that variation in what is known leads to variation in how statements of epistemic modality are assessed, and not to variation in what people say when they utter an epistemically modalised sentence. And we can hold this view because we have managed to develop a semantics that does not ascribe truth-conditions to such sentences. Instead, the content of epistemically modalised sentences is described in terms of how they affect contexts, here understood as information states. But we did not stray too far from classical truth-conditional semantics. The basic picture about context-content interaction is perfectly compatible with the view that semantics is, at its core, truth-conditional semantics. In fact, consider again the non-modal fragment \mathcal{L}_0 of \mathcal{L} . For each element of \mathcal{L}_0 we can define a set of indices at which it is true.

Definition 8 (Propositions) The function $\llbracket \cdot \rrbracket$ assigns to each $\phi \in \mathcal{L}_0$ a proposition, understood as a subset of W , as follows:

- (1) $\llbracket p \rrbracket = \{w \in W : w(p) = 1\}$
- (2) $\llbracket \neg \phi \rrbracket = W \setminus \llbracket \phi \rrbracket$
- (3) $\llbracket \phi \wedge \psi \rrbracket = \llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket$

It is now easy to verify that updating a state σ with $\phi \in \mathcal{L}_0$ comes down to adding the proposition associated with ϕ to σ .

Fact 5 For all $\phi \in \mathcal{L}_0$, $\sigma \in S$: $\sigma \uparrow \phi = \sigma \cap \llbracket \phi \rrbracket$

It is just that when we throw epistemically modalised sentences into the game context update is not always mediated by propositional content. And this is because we cannot assign to an epistemically modalised sentence a set of points at which it is true. If it weren’t for these constructions, then, there would be no departure from the foundations of truth-conditional semantics.

So here is the epilogue to the story. It is often said that a non-truth-conditional semantics for epistemically modalised sentences is not viable since such sentences

embed (though not as freely as factual sentences).¹² And this is certainly a fair objection to the popular view that epistemic modals are used to express some kind of ‘comment’, indicating the degree or source of the speaker’s commitment to the embedded proposition.¹³ Such theories are clearly in bad shape when it comes to handling embeddings. But it should be already obvious from what has been told so far that the framework laid out here does not suffer from this problem – it gives clear answers to what happens when epistemic modals are embedded under negation, conjunction, etc. So there is no reason at all to believe that there is any embedding problem for the framework developed so far. Some other challenges will come up in the next section.

3 More Data

The framework developed in the previous section allows for a simple and uniform view on the two disputes between Alex and Mary. This is an achievement, for no other semantic theory currently on the market does the same. But there are other data we need to take care of. Some of them will turn out to be already covered by what has been said so far, but others will demand minor extensions.

3.1 Ignorance

It is a rarely discussed fact that agents are sometimes agnostic about a certain epistemic possibility, even though the prejacent is a question under discussion. As [DeRose \(1991\)](#) points out, the following discourse sounds perfectly fine:

- (3) Mary: I’ve heard that John is sick. Might it be cancer?
Alex: I don’t know whether it might be cancer. The tests will be in tomorrow.

This is a rather surprising observation if one interprets ‘It might be that ϕ ’ in the mouth of a speaker S as ‘It is compatible with what S knows that ϕ ’, as neither Alex nor Mary should have any doubts as to whether it is compatible with what they know that John has cancer. John’s having cancer is compatible with what Mary and Alex know but, as we shall say, it is not a live epistemic possibility for them. It follows immediately from this that one may learn that so-and-so might be the case, in the same sense in which one may learn that so-and-so is the case. This conclusion is also supported by the following example:

¹²See, e.g., [MacFarlane \(2009a\)](#), §4.

¹³This approach to epistemic modality operators as comments on embedded propositions is certainly the standard non-truth-conditional view. See, among many others, [Kant \(1781\)](#), [Frege \(1879\)](#), [Price \(1983\)](#).

- (4) Mary: I can't find John. Do you know where he is?
Alex: He might be at home.
Mary: Oh, OK, I call him and check.

Even if it has already been compatible with what Mary knows that John is at home, there is a sense in which Alex has provided her with some non-trivial information. We can say that she has become aware of a certain epistemic possibility, and that this change in her information state allows her to coordinate her efforts to contact John.

The distinction between mere compatibility of ϕ with what is known and ϕ being a live epistemic possibility is also of relevance for the following observations about utterances of epistemic necessity. Many discussions in the semantic literature on epistemic modality are driven by the observation that speakers reliably reject a claim of epistemic possibility if the truth of the prejacent is incompatible with what they know. One should then expect that they reject an utterance of epistemic necessity if the falsity of the prejacent is compatible with what they know. But this is not always the case:

- (5) Mary: I can't find Colin. Do you know where he is?
Alex: He must be in Chicago right now.
Mary: Oh, OK. What is he doing in Chicago?

Clearly, Mary does not know where Colin is, so it is compatible with what Mary knows that Colin is not in Chicago. Nevertheless, it is natural for Mary not to reject Alex's utterance and instead to uptake the encoded information. In fact, if Mary has no clue where Colin is, it would be plain weird for her to reject Alex's utterance by pointing out that, for all she knows, Colin might very well be in Boston, Madagascar or Siberia.

In contrast, consider a case in which an assessor has good grounds for believing that the speaker has overlooked a certain possibility. Alex and Mary are passing by Colin's house. They cannot see Colin, but the lights are on:

- (6) Alex: Colin must be at home.
Mary: No, he might be out – maybe he just forgot to turn off the lights.

Here Mary points out that the available evidence does not eliminate the possibility that Colin is out – he simply might have forgotten to turn off the lights when leaving the house. Notice, again, that Mary's reaction to Alex's judgement is based on her very own information state. Alex may very well be in the position to rule out the possibility that Mary is concerned with. For example, the discourse in (6) may continue as follows:

- (7) Alex: No, he never forgets to do that.
Mary: Oh, OK.

Of course, conversations do not always proceed as in our toy examples. We expect there to be situations in which an agent will simply go along with an utterance of epistemic necessity, even though the recipient is aware of a possibility in which the prejacent is not true. This is not a special feature of epistemic claims: we very often revise our current beliefs in order to accommodate new information, especially when it comes from a source that we think knows much more about the issue than we do. But such subtle issues as trust and authority need not detain us here: even if we keep the degree of Mary’s trust in Alex’s reliability fixed, there is a clear distinction between the situation in which it is legitimate for her to reject Alex’s claim of epistemic necessity and the one in which she has no grounds whatever to reject. The former is a situation in which Mary has grounds for believing that the prejacent might be false, the latter is a case in which the falsity of the prejacent is merely compatible with what she knows. This, I submit, is a distinction that a semantics for judgements of epistemic modality should be able to draw.

We already know that the semantics developed in §2 can in fact draw the distinction. There I said that learning that ϕ might be the case is best understood as a transition from ϕ being merely compatible with what is known to ϕ being a live epistemic possibility. To make sense of this, information states have to be fine-grained enough so that we can distinguish between what is merely compatible with what is known and a live epistemic possibility. Hence I chose to model information states as sets of sets of possible worlds. One result of §2 was that the distinction between ϕ being compatible with what is settled in Σ and ϕ being a live epistemic possibility in Σ corresponds to the distinction between $\Sigma \triangleright \diamond\phi$ and $\Sigma \models \diamond\phi$. And the good news is that this is all we need to make sense of the recently made observations. Let’s first get clear on what it means for an update to be informative.

Definition 9 (Informativity) Consider any $\phi \in \mathcal{L}$ and $\Sigma \in I$:

ϕ is *informative with respect to* Σ iff $\Sigma[\phi] \neq \Sigma$ and $\Sigma[\phi] \neq \emptyset$

In words, [Stalnaker \(1978\)](#) basically got it right. Adding the information encoded in ϕ to Σ should eliminate some but not all elements of Σ . Anything else would miss the point that what is said should be informative.

An agent who is in a ‘ $\diamond\phi$ ’-admitting information state is agnostic about that sentence – the information the agent possesses neither entails ‘ $\diamond\phi$ ’ nor ‘ $\neg\diamond\phi$ ’. We can make perfect sense of dialogues like the one in (3): it is compatible with what Alex and Mary know that John has cancer, but that alone does not make it a live epistemic possibility. That is, the possibility is admitted, but not supported, by the relevant information states.

The existence of information states that admit epistemically modalised sentences also guarantees that update with such sentences can be informative, as observed in (4). The following is a trivial consequence of the definitions of admission and support:

Fact 6 For all $\phi \in \mathcal{L}$, $\Sigma \in I$: If $\Sigma \triangleright \phi$, then $\Sigma[\phi] \neq \Sigma$ and $\Sigma[\phi] \neq \emptyset$

Thus whenever we have a $\ulcorner \diamond \phi \urcorner$ -admitting information state, $\ulcorner \diamond \phi \urcorner$ will be informative with respect to that information state.

And here is how the story connects with *must*. Whenever Σ admits $\ulcorner \diamond \phi \urcorner$, it also admits $\ulcorner \square \neg \phi \urcorner$:

Fact 7 For all $\phi \in \mathcal{L}$, $\Sigma \in I$: If $\Sigma \triangleright \diamond \phi$, then $\Sigma \triangleright \square \neg \phi$

But when Σ supports $\ulcorner \diamond \phi \urcorner$, $\Sigma[\square \neg \phi] = \emptyset$ – this is trivial since we treat *might* and *must* as duals.

Thus the contrast between admitting and supporting a certain epistemic possibility also allows us to capture the observations we made in connection with epistemic *must*. If an agent is in a $\ulcorner \diamond \phi \urcorner$ -admitting information state, what the agent knows is compatible with ϕ , but does not encode any grounds for believing that ϕ is a real epistemic possibility. Accordingly, the agent has no grounds for disbelieving $\ulcorner \square \neg \phi \urcorner$. Going back to the discourse in (5), it is merely compatible with what Mary knows that Colin is not in Chicago. Thus Mary’s information admits Alex’s claim that Colin must be in Chicago, and we predict that she does not – and in fact should not – reject Alex’s utterance.

On the other hand, Mary’s information in (6) supports the epistemic possibility that Colin is out, the fact that the lights are on notwithstanding. So if we let q stand short for ‘Colin is at home’, then Mary’s information state Σ_M is such that $\Sigma_M \models \diamond \neg q$ and, accordingly, $\Sigma_M[\square q] = \emptyset$. Thus we predict that Mary may legitimately reject Alex’s claim of epistemic necessity – Colin *might* be out, and so he does not *have* to be at home.

Moving away from standard approaches by modelling information states as sets of sets of possible worlds does some genuine work. It leaves room for speakers to be agnostic about certain epistemic possibilities. This, I submit, is the right diagnosis for the observations we made at the beginning of this section. Utterances of epistemically modalised sentences can be informative, and agents admit or reject claims of epistemic necessity depending on whether or not they have good reason to believe that the prejacent might be false. All of this makes perfect sense once we realise that ‘It might be the case that ϕ ’ is very different from ‘ ϕ is compatible with what is known’.

3.2 Retrospective Assessment

It is often observed that speakers tend to retract previously made claims of epistemic modality in the light of new information:

- (8) Mary: I can’t find my keys.
 Alex: They might be in the car.
 Mary: No, they can’t be in the car. I still had them with me when I came in.
 Alex: Oh, so I guess I was wrong.

But as [von Fintel and Gillies \(2008b\)](#) point out, this is not the only possible reaction. When it comes to statements of what might or might not be the case, speakers can quite often resist the invitation to retract even if they have become better informed. Thus Alex is free to react in the following manner:

- (9) Mary: I can't find my keys.
Alex: They might be in the car.
Mary: No, they can't be in the car. I still had them with me when I came in. Why would you say that?
Alex: Look, I didn't say that they were in the car. I said that they might have been in the car.¹⁴ And they might have been.

Here Alex sticks to his guns, and the question why this is a legitimate move is as important as the question why it is also sometimes natural to withdraw a commitment to a certain epistemic possibility in the light of new information.

One part of the puzzle is that Alex may admit that the keys weren't in the car yet insist that they might have been there. But even this piece of datum is anything but innocent. To say that the keys might have been in the car is to either talk about a past possibility or a possible past.¹⁵ Thus consider:

- (10) a. The keys weren't in the car, but they might have been in the car.
b. The keys might have been in the car, or they might have been in the drawer. I don't remember.

What (10a) says is that the keys being in the car was once an epistemic possibility, even though it is not any longer. What the first disjunct in (10b) says is that it is now an epistemic possibility that the keys were in the car. An alternative way of putting (10b) is the following:

- c. Maybe the keys were in the car, or maybe they were in the drawer. I don't remember.

The reason why all this matters is that when Alex admits that the keys weren't in the car yet insists that they might have been there, he must be concerned with a past possibility, not a possible past. To see this more clearly, contrast (10a) with the following strange example:

- d. # The keys weren't in the car, but maybe they were.

If Alex learns that the keys weren't in the car, he may hold on to the past possibility of the keys being in the car, but not to the possibility that the keys

¹⁴Speakers' intuitions are not uniform when it comes to how Alex should make his point here. Some speakers, including me, prefer the backshift of *might*, while others, including von Fintel and Gillies, avoid the backshift. Such minor disputes need not be resolved here.

¹⁵See also [Stephenson \(2007, 2008\)](#) for this distinction.

were in the car. And this is just another observation that requires an explanation.

What we have seen is that there are two interesting observations about retrospective assessment. First, learning that the prejacent of an earlier *might*-statement is false may trigger a positive or negative assessment: the speaker may admit that something went wrong or insist that the new information does not change the fact that the truth of the prejacent was an epistemic possibility. This is just the moral drawn from the examples in (8) and (9). Second, saying that the truth of the prejacent was an epistemic possibility is importantly different from saying that it is an epistemic possibility that the prejacent was true: the former goes quite well together with admitting that the prejacent was in fact wrong, but the latter does not. This is just the moral drawn from the comparison of (10a) and (10d).

The cases under consideration involve agents whose information changes over time, and rest on the interaction between epistemic modality and tense. The framework developed so far has nothing to say about this interaction, as it says nothing about tense. So we hit a point at which an extension of the story told so far is in order.

3.3 More on Embeddings

As Yalcin (2007, 2009) observes, there is an interesting logical connection between factual judgements and judgements of epistemic modality. Thus the following sentence sounds terrible, and cannot even be supposed to be true:

- (11) a. # It is raining and it might not be raining.
b. # Suppose it is raining and it might not be raining.

This is anything but surprising from a dynamic perspective on epistemic modality. It was already observed in the previous section that according to the story told here, ‘It is raining’ entails ‘It must be raining’. And since *might* and *must* are duals, it follows that updating with ‘It is raining and it might not be raining’ always results in the absurd state.¹⁶ So an attempt at updating one’s information state with (11a) will have the same outcome as an attempt at updating one’s information state with a more familiar contradiction like ‘It is raining and it is not raining’. Thus (11a) cannot even be supposed to be true, any more than a plain classical contradiction. And given the more than reasonable assumption that one cannot assert what cannot at least be supposed, it follows that a plain assertion of (11a) is pragmatically infelicitous.

The framework under consideration allows epistemically modalised sentences to embed under negation and conjunction, and it does so in a way that takes care of some rather tricky data. But we also know that such sentences can occur in conditionals and under the scope of attribute ascriptions. Thus consider:

¹⁶But remember that Contraposition fails: even though ‘It is raining’ entails ‘It must be raining’ and *might* and *must* are duals, ‘It might not be raining’ does not entail ‘It is not raining’.

- (12) a. If it might be raining, I stay at home.
 b. If we do not invite John, he might be disappointed.
 c. Alex believes that the keys might be in the car.
 d. Mary does not believe that Colin must be out.

Since we have not said anything about attitude ascriptions and natural language conditionals, handling such embeddings requires an extension of the framework developed so far.¹⁷ One often hears it said that the possibility of such embeddings poses insurmountable problems to any semantics for epistemically modalised sentences that is not truth-conditional, but it does not take much to figure out how to accommodate the new data in our dynamic story. To ascribe a belief that ϕ to an agent A is to claim that A is in a ϕ -supporting information state. And processing a conditional comes down to eliminating all those states that do not support the consequent after being updated with the antecedent. The details of these suggestions will be elaborated in the next section.

4 Extensions

Part of the exercise is to provide a semantics for the following extension of our base language \mathcal{L} :

Definition 10 (Full Language) \mathcal{L}^+ is the smallest set that contains any sentential atoms $\mathcal{A} = \{p, q, \dots\}$ and is closed under negation, conjunction, the epistemic modal *might*, the natural language conditional (\Rightarrow) and the temporal operators for *past* (P) and *future* (F). \mathcal{L}_0^+ is defined as the smallest set that contains the sentential atoms of \mathcal{A} and is closed under negation, conjunction, and the temporal operators for *past* and *future*, i.e. as the extension of the non-model fragment of \mathcal{L}^+ with our temporal operators P and F . Disjunction, the material conditional, and the epistemic modal *must* (\Box) are defined in the usual way.

Everything that has been said about the \mathcal{L} remains valid, so all we have to deal with are the newly introduced operators. The final point on the to-do list is then to indicate how formulas of \mathcal{L}^+ interact with attitude ascriptions.

4.1 Conditionals

Ramsey famously suggested that in order to evaluate a conditional, one should check whether one hypothetically accepts the consequent under the assumption that the antecedent holds. He also held that conditionals do not bear truth-values,

¹⁷I assume here without further discussion that ‘if ... then’ cannot be semantically analysed in terms of the material conditional. This assumption should be unsuspecting as it means, in the end, more work for me.

as one can agree on all the facts yet disagree on a conditional.¹⁸ Thus if Ramsey’s suggestion is on the right track – and I think it is – then conditionals are very similar to epistemically modalised sentences. They have content, but not truth-conditional content. Specifically, the CCP of a conditional is similar to the one of an epistemically modalised sentence in that it eliminates states that do not pass a certain test. And the test is just that the consequent holds throughout the whole state once it has been updated with the antecedent. Here is the more precise version:

Definition 11 (Conditionals) Consider any $\sigma \in S$, $\phi, \psi \in \mathcal{L}^+$. An update on a state with $\lceil \phi \Rightarrow \psi \rceil$ is defined as follows:

$$\sigma \uparrow \phi \Rightarrow \psi = \{w \in \sigma : (\sigma \uparrow \phi) \uparrow \psi = \sigma \uparrow \phi\}$$

In words, very much like a claim of epistemic modality, a conditional invites the agent to run a *test* on each state. If updating a state σ with ϕ and then with ψ amounts to nothing more than updating σ with ϕ , then the state passes the test. Otherwise, we get back the empty set and do not need to consider the state at a later stage. For example, updating σ with $\lceil p \Rightarrow q \rceil$ returns σ if all p -worlds in σ are q -worlds, and returns the empty set otherwise.

There are a lot of positive things to say about this semantics for conditionals, but I will limit myself to the points of immediate importance.¹⁹ What we needed to explain is how epistemically modalised sentences can occur in conditionals. So consider the examples in (12a) and (12b) again (repeated):

- (12) a. If it might be raining, I stay at home.
b. If we do not invite John, he might be disappointed.

Depending on whether the state contains at least one rainy world, updating it with the antecedent of (12a) will either return the original state or the empty set. Since the test is trivially passed in the latter case, we may as well leave those states alone and only check at the remaining states whether I stay at home in all possible worlds. If I do, the original state is returned; otherwise we get back the empty set. Thus updating an information state with (12a) will eliminate all those states in which (i) there is at least one possible world in which it rains and (ii) I do not stay at home in all possible worlds. In processing (12b), we first eliminate all worlds in which we invite John, and then check whether updating with the epistemically modalised sentence induces any change on the resulting state. This will not be the case only if the resulting state is already empty or there is at least one possible world in which John is disappointed. Thus updating an information state with (12b) will eliminate all those states in which (i) there is a possible world in which we do not invite John and (ii) John fails to be disappointed in any of those possible worlds in which we do not invite him. Embedding *might* into conditionals thus

¹⁸All this in his “General Propositions and Causality”; consult Ramsey (1990) for details.

¹⁹To fully appreciate the glory of a dynamic semantics for conditionals, I recommend that you take a look at what Gillies (2004, 2009a,b) has to say about this issue.

poses no problems, and nor does, as it is straightforward for the reader to verify, embedding *must*. Let me now tell you how some basic tense operators should be integrated into the system, which will allow us to say something reasonable about the way speakers tend to evaluate previously made judgements of epistemic modality in the light of new information.

4.2 Tense

Reality changes over time, and tense operators are frequently used to say how things were or will be. But information states change over time as well, and so we should expect that tense operators interact with epistemic modals to report on past or future features of a given information state. This is crucial for our observations about retrospective assessment in §3.2. There we saw that Alex may insist that the keys might have been in the car, even though they weren't. What he correctly points out is that there once was the epistemic possibility of the keys being in the car, though this possibility does not exist any longer. The task now is to show how this is possible.

What we said before does not allow us to understand possible worlds or information states as extended over time, so the first thing that needs to be done is to change the relevant definitions accordingly. The obvious steps are to understand possible worlds as assigning elements of \mathcal{A} truth-values relative to points of time, and to think of states as functions from points of time to sets of such possible worlds. More precisely:

Definition 12 (Tense) Let $Temp$ be any nonempty set, $<$ be a transitive relation on $Temp$ which is also irreflexive and linear. v is a *temporally extended possible world* iff $v:Temp \mapsto (\mathcal{A} \mapsto \{0,1\})$. V is the set of such v 's. τ is a *temporal state* iff $\tau:Temp \mapsto (V \mapsto \{0,1\})$. S^e is the set of such τ 's. T is a *temporal information state* iff $T \subseteq (S^e \setminus \emptyset)$, i.e. a temporal information state is a (possibly empty) set of non-empty sets of temporally extended possible worlds. I^e is the set of such T 's. The *initial* temporal information state T_0 is identical with $(S^e \setminus \emptyset)$, the *absurd* temporal information state T_\emptyset with \emptyset .

Earlier we understood updates of information states in terms of updates of their elements, and there is no reason to change this once we move on to temporal information states. Thus the most important task is to determine how temporal states are modified. Temporal states change over time and are modeled as functions from points of time to sets of possible worlds. So the natural thing to do is to understand updates on temporal states as resetting the values assigned to temporal indices. For this proposal to make sense, one needs to say how the values assigned to temporal indices are determined by updates and what it means to reset the values assigned to temporal indices. Let me address these issues in turn.

Definition 13 (Updates on Temporal States) Consider any $\tau \in S^e$, $t \in Temp$, $p \in \mathcal{A}$ and $\phi, \psi \in \mathcal{L}^+$. An update on a temporal state is a function $\uparrow: \wp(V) \mapsto \wp(V)$ defined by the following recursion:

- (1) $\tau(t) \uparrow p = \{v \in \tau(t) : v(t)(p) = 1\}$
- (2) $\tau(t) \uparrow \neg\phi = \tau(t) \setminus (\tau(t) \uparrow \phi)$
- (3) $\tau(t) \uparrow \phi \wedge \psi = (\tau(t) \uparrow \phi) \uparrow \psi$
- (4) $\tau(t) \uparrow \diamond\phi = \{v \in \tau(t) : \tau(t) \uparrow \phi \neq \emptyset\}$
- (5) $\tau(t) \uparrow \phi \Rightarrow \psi = \{v \in \tau(t) : (\tau(t) \uparrow \phi) \uparrow \psi = \tau(t) \uparrow \phi\}$
- (6) $\tau(t) \uparrow P\phi = \{v \in \tau(t) : \exists t' < t : v \in \tau(t') \uparrow \phi\}$
- (7) $\tau(t) \uparrow F\phi = \{v \in \tau(t) : \exists t' > t : v \in \tau(t') \uparrow \phi\}$

As before, the update function \uparrow operates on sets of possible worlds, which are now understood as assigning truth-values to sentential atoms relative to points in time. Which set of possible worlds in a state gets selected for update depends on a temporal index, which in turn will ultimately depend on the context in which the sentence is uttered (see below). So consider the set of worlds a state assigns to time t . The clause in (1) requires that updating that set with an atom p eliminates all possibilities in which p is false at t . Clauses (2) – (5) should be obvious from what has been said earlier. Clauses (6) and (7) treat tense operators in \mathcal{L}^+ as shift: in order to determine the update effects of a temporally modified formula on a set of worlds $\tau(t)$, it is required to check the update effect of the prejacent on a different set of worlds $\tau(t')$. So for instance, update of $\tau(t)$ with $'P\phi'$ leaves those worlds in $\tau(t)$ which would have survived update of some previous $\tau(t')$ with ϕ . Specifically, v remains in $\tau(t)$ as updated with Pp just in case there is some earlier time t' such that v remains in $\tau(t')$ as updated with p . And that just means that v remains in $\tau(t)$ as updated with Pp just in case there is some earlier time t' such that $v(t')(p) = 1$. That said, the update effects of formulas involving the future tense operator should be obvious.

Temporal states change over time, and we agreed that since temporal states are functions from points of time to sets of possible worlds, this should be understood as a resetting of such functions.

Definition 14 (Reset) Consider any $\tau \in S^e$, $t \in Temp$, $\phi \in \mathcal{L}^+$:

1. $\tau^{\tau(t) \uparrow \phi}$ is the temporal state which is just like τ except that for all $t' \not\prec t$, $\tau^{\tau(t) \uparrow \phi}(t') = \tau(t) \uparrow \phi$
2. $\tau \uparrow^t \phi = \tau^{\tau(t) \uparrow \phi}$

A temporal state τ is reset relative to a formula ϕ and a time t . For all $t' < t$, we leave the value of τ untouched. For all other temporal indices, we reset the value of τ to $\tau(t) \uparrow \phi$, i.e. the result of updating the image of t under τ with ϕ .

Updates of temporal information states are now defined as indicated before.

Definition 15 (Updates on Temporal Information States) Consider any $T \in I^e$, $t \in Temp$, $\phi \in \mathcal{L}^+$. An update on a temporal information state at t is a function $[\cdot]^t : I^e \mapsto I^e$ defined as follows:

$$T[\phi]^t = \{\tau : \tau(t) \neq \emptyset \wedge \exists \tau' \in T : \tau' \uparrow^t \phi = \tau\}$$

Update of a temporal information state T with a formula ϕ at time t comes down to the following procedure: first reset each element of T with respect to ϕ and t ; then gather all the resulting temporal states together, leaving out those which assign to t the empty set. This gives you the output temporal information state. A more informal characterisation is this: agents are often agnostic about past, present or future facts, and we capture this, as usual, in terms of possible worlds which disagree on the past, present or future. As time proceeds, an agent acquires information, removing epistemic uncertainty. The moment in which the information is acquired we update each set of worlds and set the result as the current value at each temporal state. The set of worlds remains in place until some new information is processed. But of course, in doing so the agent does not forget what her previous information states looked like. This is why we only reset states from the moment at which the information is received.²⁰

It is not necessary to write down all the obvious modifications of our previous definitions, but the following will facilitate the discussion:

Definition 16 (Temporal Settledness, Temporal Admission, Entailment)

Let T be a temporal information state, t be a time and ϕ, ψ be formulas:

1. T supports ϕ at t , ϕ is settled in T at t , $T \models_t \phi$, iff $T[\phi]^t = T$
2. T admits ϕ at t , $T \triangleright_t \phi$, iff $T \not\models_t \phi$ and $T \not\models_t \neg\phi$
3. ϕ entails ψ , $\phi \models \psi$, iff $\forall T, t : T[\phi]^t \models_t \psi$

This is then how I suggest integrating the basics of a semantics for tense into our dynamic framework. Without doubt, it is nothing more than a very small step towards a full theory of tense and its interaction with modality. But it is substantial enough to cover the main issues about tense and epistemic modality.

One observation is that the interaction between tense and modality gives rise to scope issues, and this is all we need to explain our earlier observation in §3.2 that one can talk about a possible past or a past possibility. Remember our earlier examples:

²⁰And of course, we idealise a bit insofar as agents never forget acquired information and always know what time it is.

- (12) a. The keys weren't in the car, but they might have been in the car.
 b. The keys might have been in the car, or they might have been in the drawer. I don't remember.
 c. Maybe the keys were in the car, or maybe they were in the drawer. I don't remember.

What (12a) is about is a past possibility, and thus we expect the epistemically modalised sentence to be of the form ' $P\Diamond p$ '. On the other hand, the first conjuncts in (12b) and (12c) are concerned with a possible past, and that means that they are of the logical form ' $\Diamond Pp$ '. This treatment connects nicely with our observation that when Alex admits that the keys weren't in the car yet insists that they might have been there, he must be concerned with a past possibility, not a possible past. What justified this verdict was that the sentence in (12d) sounds terrible (repeated):

- d. # The keys weren't in the car, but maybe they were.

And the infelicity of (12d) is not surprising, any more than the infelicity of Yalcin's 'It is raining and it might not be raining'. An update with 'The keys weren't in the car' eliminates all possible worlds in which the keys were in the car, and thus there remains no doubt that the keys weren't in the car. Accordingly, updating with the second conjunct results in the absurd state. So updating with (12d) has the same outcome as updating with a classical contradiction like 'The keys weren't in the car and the keys were in the car'. What remains to be explained is why Alex may hold on to the past possibility of the keys being in the car after learning that they weren't in the car. And this brings us back to the complicated data about retrospective assessment. The really good news is that the current framework can explain all this as well. Let us go through the details.

The key premise of our framework is that the meaning of a sentence ϕ is to be understood in terms of its update effects on (temporal) information states. So far we have focussed on the commitments which *arise* from an update with ϕ , i.e. on what becomes settled in an information state as a result of updating with ϕ . And this was the right thing to do because we wanted to know what is entailed by the sentences of our toy language. But we may also ask which commitments *survive* an update with ϕ , i.e. which commitments can be rationally maintained once it is learnt that ϕ is the case. And this, I submit, is the right thing to ask when we want to explain how agents assess previously endorsed judgements of epistemic modality after learning something new.

What we want to know is to what degree updating with sentences of \mathcal{L}^+ *preserves* already existing commitments:

Definition 15 (Preservative Update) Consider any $\phi, \psi \in \mathcal{L}^+$. Updating with ψ *preserves* ϕ iff for all $T \in I^e$, $t \in Temp$: If $T \models_t \phi$, then $T[\psi]^t \models_t \phi$. We say that updating *preserves* ϕ iff for every sentence ψ in \mathcal{L}^+ , updating with ψ preserves ϕ .

Take any T which settles ϕ at t . If updating with ψ preserves ϕ , then updating T with ψ at t does not remove the agent’s commitment to ϕ . And updating preserves ϕ just in case no update whatever could remove the agent’s commitment to ϕ .

Preservation (or lack thereof) matters since it determines how agents assess a commitment to ϕ after learning something new.

Post-Update Assessment Let $\psi \in \mathcal{L}^+$ and consider a subject A with temporal information state T_A such that $T_A \models_t \phi$. Then A will by default assess his/her commitment to ϕ at t as follows after updating with ψ :

- **Retain** in case $T_A[\psi]^t \models_t \phi$
- **Withdraw** in case $T_A[\psi]^t \not\models_t \phi$

Thus if updating with ψ preserves ϕ , we expect that agents will always retain their commitment to ϕ in the face of learning that ψ . On the other hand, if updating with ψ fails to preserve ϕ , then we expect there to be situations in which learning that ψ may force an agent to withdraw a previous commitment to ϕ . It is in these situations that we expect an agent who learns that ψ is the case to admit that s/he was wrong in believing that ϕ .

Preservation is a big issue in the literature on belief revision, and some people claim that unless you already believe $\lceil \neg\phi \rceil$ in a prior state, revising that state with ϕ should result in a posterior state that carries at least as many commitments as the prior state.²¹ It is a well-worn story that so much doxastic conservatism leads to trouble when you also have beliefs about what might or might not be the case: you may believe that $\lceil \Diamond\neg\phi \rceil$ without believing that $\lceil \neg\phi \rceil$, and still a commitment to $\lceil \Diamond\neg\phi \rceil$ needs to go once you learn that ϕ is the case.²² As Gillies (2006) has shown, dynamic semantics gives us independent reason to reject the idea that updating is always preservative updating, and it is this feature that I intend to exploit in my story about retrospective assessment. The crucial point is that updating temporal information states fails to preserve previous commitments the right way. Let me explain.

It is straightforward to verify that updating preserves all sentences of \mathcal{L}_0^+ :

Fact 8 For all $\phi \in L_0^+$: Updating preserves ϕ .

The reason is simple: for every $\phi \in L_0^+$, we can find a set of worlds $[\![\phi]\!]$ at which ϕ is true. T settles ϕ at t just in case $\bigcup\{\tau(t) : \tau \in T\} \subseteq [\![\phi]\!]$. And since updating a temporal state with some ψ at t is always *eliminative* – we never add worlds to temporal states – it follows that $\bigcup\{\tau(t) : \tau \in T[\psi]^t\} \subseteq [\![\phi]\!]$. In informal terms: an agent is committed to $\phi \in L_0^+$ just in case ϕ holds in all of the agent’s live epistemic possibilities. And since updating only removes epistemic uncertainty, it follows that an existing commitment to $\phi \in L_0^+$ survives any update whatever.

But what really interests us are commitments to epistemically modalised sentences, and the first part of the story is that updating preserves simple past tensed

²¹See Alchourrón et al. (1985) and Gärdenfors (1988).

²²See Levi (1988), Fuhrmann (1989), and Rott (1989).

sentences:

Fact 9 For all $\phi \in L$: Updating preserves $P\phi$.

This differs from **Fact 8** insofar as ϕ may very well be a test. If ϕ is a test, T supports ' $P\phi$ ' at t just in case all temporal states in T pass the test at some earlier time t' . But no update at t changes the image of t' under some $\tau \in T$, and thus it is guaranteed that T continues to support ' $P\phi$ ' after the update. To go through an example, consider an agent whose temporal information state T supports ' $P\Diamond p$ ' at time t . We know that $T \models_t P\Diamond p$ just in case for all $\tau \in T$, $\tau(t) \uparrow P\Diamond p = \tau(t)$. And this holds just in case for every $\tau \in T$, there is some $t' < t$ such that $\tau(t') \uparrow p \neq \emptyset$, i.e. such that the image of t' under τ contains at least one world v at which p is true at t' . But whatever happens to τ at t , we never reset the values τ assigns to any $t' < t$. Thus the relevant v is guaranteed to remain an element of the image of t' under τ after any update. It follows that for every $\psi \in L^+$, $T[\psi]^t \models_t P\Diamond p$. What underlies all this is that learning something new only changes the agent's current information state, but not the agent's information states at some earlier time.

According to **Fact 9**, past possibilities never get eliminated, and thus updating preserves commitments to past possibilities. But preservation has to stop somewhere, and what matters is that it stops when it comes to commitments to current epistemic possibilities. Specifically, updating with ' $\neg\phi$ ' does not preserve ' $\Diamond\phi$ ' unless ϕ is a tautology. Thus we get:

Fact 10 For all non-trivial $\phi \in \mathcal{L}^+$: Updating with $\neg\phi$ does *not* preserve $\Diamond\phi$.

And this should not surprise anyone: one's commitment to the possibility of ϕ being the case has to go once it is learnt that ϕ is not the case. In fact, the right kind of commitment now is that ϕ cannot be the case. So while past epistemic possibilities never die, current epistemic possibilities may very well fall victim to new information.

We now have everything we need to explain the complicated data about how agents assess their judgements of epistemic possibility after learning that the pre-jacent in fact is false. The puzzle is that there are two appropriate reactions for Alex after learning that the keys are not in the car: insisting that the keys might have been in the car, or admitting that something went wrong. What underlies all this is that Alex may either flag that the new information allows him to retain his belief that the keys might have been in the car or indicate that the new information forces him to withdraw the belief that the keys might be in the car. Thus after learning that the keys are not in the car, Alex may respond in two ways:

- (13) Alex: a. Fair enough. But they might have been in the car.
b. Oh, so I guess I was wrong. They can't be in the car.

What makes all this possible is that updating with ' $\neg\phi$ ' preserves ' $P\Diamond\phi$ ' but not in general ' $\Diamond\phi$ '. Specifically, it follows from **Fact 9** and **Fact 10** that updating with 'The keys are not in the car' preserves 'The keys might have been in the car'

but not ‘The keys might be in the car’. And Alex, we can say, is committed to two very different things in the moment he learns that the keys are not in the car: the past possibility of the keys being in the car, and the current possibility of the keys being in the car. Learning that the keys are not in the car does nothing to remove the past possibility, but it eliminates the current possibility. Thus we expect there to be retention as well as withdrawal: Alex may insist that the keys might have been in the car, the fact that they aren’t notwithstanding. But he may also admit that the new information forces him to give up his commitment to the current epistemic possibility of the keys being in the car. In other words, Alex still accepts ‘ $P\Diamond p$ ’ but no longer accepts ‘ $\Diamond p$ ’ (where p is short for ‘The keys are in the car’). Accordingly, (13a) as well as (13b) are perfectly legitimate reactions.

This is another bonus of the dynamic perspective on epistemic modality. Updating fails to be preservative when it comes to what the agent currently takes to be a live epistemic possibility, and this is how it should be. Current commitments to epistemic possibilities may need to go even though the agent has no reason to think that the possibilities never existed. To put a shiny label on our theory: commitments to past possibilities are static, commitments to current possibilities are dynamic. And together with our principle about post-update assessment, we can put this theory to good work when we want to explain how agents assess their commitments to epistemic to epistemic possibilities after learning something new.

What has happened so far should make us confident that a successful semantics for epistemically modalised sentences does not have to be a truth-conditional semantics. Accordingly, there is no need to buy into one of the extravagant consequences that orthodoxy brings in its wake. Since the full semantics for \mathcal{L}^+ is now in place, it is time to finish the discussion with a quick outlook on the interaction between epistemically modalised sentences and attitude verbs.

4.3 Attitude Ascriptions

The fact that epistemically modalised sentences may appear under the scope of attitude verbs is often perceived as yet another embedding problem for a non-truth-conditional semantics of such constructions. What is true is that if such verbs denote relations between individuals and propositions, then every semantics will be in trouble which does not assign to epistemically modalised sentences truth-condition determining propositions. But we may as well understand attitude verbs as denoting relations between individuals and CCPs.²³ More precisely, let s range over indices of evaluation, e range over individual objects, t range over truth-values. If attitude verbs denote relations between individuals and propositions, they are of type $\langle\langle s, t \rangle, \langle e, t \rangle\rangle$. Things look only slightly different from the alternative perspective. States are, like propositions, of type $\langle s, t \rangle$, information states are of type $\langle\langle s, t \rangle, t \rangle$, and CCPs are of type $\langle\langle\langle s, t \rangle, t \rangle, \langle\langle\langle s, t \rangle, t \rangle, t \rangle\rangle$. Thus if attitude verbs denote relations between individuals and CCPs, they are of type $\langle\langle\langle\langle s, t \rangle, t \rangle, \langle\langle\langle s, t \rangle, t \rangle, t \rangle\rangle, \langle e, t \rangle\rangle$. The only notable

²³If you think that belief is a relation between individuals and sentences, be my guest – that is something a fan of the dynamic picture could say as well.

difference between the two proposals is the type of the first relatum. From a structural perspective, the proposals are alike: to say that ‘ S believes that ϕ ’ is to claim that S stands in the belief-relation to the meaning of ϕ , i.e. a proposition or – if one prefers the dynamic view – a CCP. Thus whatever option we take, we capture the usual inferential features of attitude ascriptions that frequently motivate a relational analysis. The fact that one may believe that so-and-so might or must be the case does not entail that epistemically modalised sentences express truth-condition determining propositions.

So far I have treated the belief relation as basic, but there is an alternative. Instead we could say that an agent believes that ϕ just in case the agent is in a ϕ -supporting information state. Thus our lexical semantics would associate the following denotation with *to believe*:

$$(14) \quad \llbracket \text{believe} \rrbracket^i = \lambda R. [\lambda x. \exists \Sigma: x \text{ is in } \Sigma \text{ at } i \text{ and } \langle \Sigma, \Sigma \rangle \in R]$$

In other words, a sentence of the form ‘ S believes that ϕ ’ is true at an index i if and only if S is in a ϕ -supporting information state at i . The semantics of such sentences is then given in terms of their CCPs: they eliminate, from each information state, those worlds in which the subject is not in an appropriate information state. This is just another nice feature of the new semantic framework: though it is compatible with it to take the belief-relation as basic, it is also straightforward to provide a non-trivial analysis in terms of information states and CCPs.

One may have beliefs about what might or must be the case, but these beliefs are not true or false. Nevertheless each such belief belongs to an information state for which the issue of truthfulness can be raised (cf. §2). Let v_a be the (temporally extended) actual world and T be a (temporal) information state. Then we can ask (i) whether $v_a \in \bigcap \{\tau(t): \tau \in T\}$ and (ii) whether $v_a \in \bigcup \{\tau(t): \tau \in T\}$. If (i) holds the actual world is a live epistemic possibility; if (i) fails but (ii) holds the actual world is at least not excluded by the information encoded in T . If the answers to (i) and (ii) are both negative, then the information state misrepresents reality. Beliefs about epistemic modality do not (mis)represent reality, but they are incorporated into a system which, we may say, aims at representing the actual world and which may or may not be truthful.

It would be too much to try to give a full analysis of attitude ascriptions here. The purpose of this exercise was to show that when it comes to attitude ascriptions, the current framework for epistemic modality is at least as good as any other current semantics for epistemic modality. The claim that attitude ascriptions pose a particular problem to a non-truth-conditional approach to epistemically modalised sentences is unfounded.

I have shown how to extend the semantic framework developed in §2 to cover some phenomena involving conditionals, tense, and attitude ascriptions. Some of these phenomena are quite puzzling and have proven difficult for alternative frameworks. The fact that the story I have told deals with the data fairly nicely should make us confident that we have an elegant and powerful analysis of epistemic modality in our hands. Let me now proceed to a comparison between my story and some other theories currently on the market.

5 Comparisons

Part of the exercise was to look for an alternative to the orthodox view that epistemically modalised sentences have truth-values relative to what is known, where the relevant knowledge is either determined by the context of utterance (contextualism) or a point of assessment (relativism). The canonical definitions:²⁴

Definition 16 (Contextualism) Consider any context of utterance c and index of evaluation i :

- (1) $\llbracket \text{might}(B)(\phi) \rrbracket^{c,i} = 1$ iff. $\exists v \in \llbracket B \rrbracket^{c,i}: \llbracket \phi \rrbracket^{c,\langle v,t_i \rangle} = 1$
- (2) $\llbracket B \rrbracket^{c,i} = \{v: v \text{ is compatible with the } c\text{-relevant knowledge in } i\}$

Definition 17 (Relativism) Consider any context of utterance c , index of evaluation i , and point of assessment a :

- (1) $\llbracket \text{might}(B)(\phi) \rrbracket^{c,i,a} = 1$ iff. $\exists v \in \llbracket B \rrbracket^{c,i,a}: \llbracket \phi \rrbracket^{c,\langle v,t_i \rangle,a} = 1$
- (2) $\llbracket B \rrbracket^{c,i,a} = \{v: v \text{ is compatible with what } j_a \text{ knows at } t_a \text{ in } w_a\}$

The key reason to be sceptical of the orthodox view is the intuition that variation in what is known should lead to variation in how utterances of epistemically modalised sentences are assessed, not to variation in what has to be the case for the sentences to be true. This is the role that variation in what is known plays in the story I have told, and as a result this story offers a simple and elegant perspective on disputes about epistemic modality, one that brings such disputes in line with other everyday disputes about matters of fact. In this respect the story I have told is superior to current versions of the orthodoxy, and it has a range of further advantages. Let me explain.

One very nice feature of my story is that there is no need to stipulate an unappealing multitude of potentially relevant knowers. Sooner or later every version of the orthodoxy has to assume that claims of epistemic modality are sometimes concerned with the speaker's knowledge, sometimes with the hearer's, sometimes with the knowledge distributed among the speaker and the hearer, and so on. There is, as far as I can tell, no justification for this assumption. That the relevant knowledge is often the speaker's is uncontroversial – otherwise we could not explain why speakers frequently make claims of epistemic possibility if the prejacent is compatible with what they know. But the assumption that speakers sometimes use epistemic modals with the intention to report on what is or is not compatible with what someone else knows is less obvious. Without it, however, orthodoxy has

²⁴Kratzer (1977, 1981, 1991) is the *locus classicus* of contextualism. See also Hacking (1967), Teller (1972), and DeRose (1991) for some letters of support for this view. See fn. 1 for some proponents of relativism. The denotation of B is often called the *modal base*. It is common to further restrict the domain of quantification to the *closest* worlds in the modal base, to be determined by an ordering source. For current purposes, it is harmless to set this complication aside. These are indexical versions of relativism and contextualism insofar as the truth-conditions vary with the context of utterance or point of assessment. But this is a negotiable feature of the views under discussion, as the context of evaluation may vary instead of truth-conditional content. Again, this issue is orthogonal to what we are after here. See MacFarlane (2009b) for discussion.

a hard time explaining how such utterances of epistemically modalised sentences behave in discourse.

The point becomes immediately clear when we consider what contextualists tend to say about the dispute between Alex and Mary in (1). It is hard to see how there could be any dispute between the discourse participants if all they did was to report on what is or is not compatible with what they (respectively) know, and one response on behalf of contextualism is that epistemic modals are ambiguous between a solipsistic reading and a range of non-solipsistic readings. This is why in stating that the keys might be in the car, Alex ‘puts into play’ the proposition that the keys being in the car is compatible with what is known by each discourse participant, and it is this discourse effect that legitimates Mary’s denial of Alex’s utterance.²⁵

This position is dubious for at least two reasons. First, it is hard to find compelling evidence for the claim that epistemically modalised sentences are ambiguous. Common ambiguity tests fail to deliver the desired support.²⁶ Consider, for instance, the contradiction test, which works as follows: if ϕ has two or more readings, then speakers should agree that an utterance of a sentence of the form ‘ $\phi \wedge \neg\phi$ ’ can be true. But now take the example in (15):

(15) # Colin might be in Chicago and he can’t be in Chicago.

(15) strikes the ear as a glaring contradiction,²⁷ and this would be rather surprising if the epistemically modalised sentence had the multiple readings suggested by some contextualists. Of course, ambiguity tests have to be handled with care, but the fact that they do *not* support the ambiguity thesis puts the proposal under consideration on rather shaky grounds. Second, even if there were an ambiguity, it would be hard to find a reasonable semantic or pragmatic principle which predicts that *all* available readings are put into the discourse and become legitimate objects of acceptance or denial. Quite to the contrary, no one would expect that using an ambiguous phrase like ‘bank’ in discourse would put into play all available readings, which are then subject to legitimate rejection by the hearer. In fact, what cooperative hearers do when confronted with an ambiguous expression is to assume the most reasonable reading or to find out what the speaker really intended to say. And even if we abstract away from all these difficulties, it remains a strong point in favour of the account developed here that it does not require the assumption that claims of epistemic modality are often, if not always, concerned with what is known by someone other than the speaker.

A quite similar objection works against relativism, though the issue is a bit more subtle. To make sense of simple disputes like (1), it is not necessary for the relativist to depart from the fairly straightforward view that it is always the assessor’s knowledge that determines the truth-value of a given claim of epistemic modality. In the first example, it is Alex’s knowledge against which his claim is

²⁵See von Fintel and Gillies (2009).

²⁶See Sadock and Zwicky (1975) for a discussion of a whole range of such tests.

²⁷If you think that this is just a pragmatic issue, notice that (15) remains more than bad even if it is just *supposed*.

evaluated when he is the assessor, and it is Mary's when she is the judge. And this is all we need to predict that Alex asserts that the keys might be in the car, while Mary denies that the keys might be in the car. But the story cannot stay that simple once we move on to claims of epistemic necessity. We know that there are cases in which a hearer does not reject a claim of epistemic necessity even though the prejacent does not follow from what the agent knows. If it is always the assessor's knowledge that counts, Alex's claim in (5) is false from Mary's point of assessment, leaving it unexplained why she does not reject what Alex said.²⁸

Of course, all the relativist says is that the relevant knowledge may vary across points of assessment, not that it is always the assessor's knowledge. So it is open to the relativist to maintain that Alex's claim is true from Mary's point of assessment, since (i) the relevant knowledge from this point is the one distributed among Alex and Mary (or maybe just Alex's) and (ii) Alex knows that Colin is in Chicago. But relativism loses a lot of appeal once this step is made. Very much as before, we lack any plausible principle which explains why it is sometimes the assessor's knowledge against which a claim of epistemic modality is evaluated, and sometimes the knowledge distributed among the assessor and the speaker. This is particularly worrying since a claim of epistemic necessity is not always assessed against the knowledge distributed among the assessor and the speaker, as the dialogue in (6) shows. For here it seems that Mary rejects Alex's claim because she thinks that there is a possibility which he has overlooked (and regardless of whether he in fact did). So Mary's point of assessment selects her knowledge as the standard of evaluation, not the one distributed among her and Alex.

Thus it seems that when it comes to the more complex cases, relativists cannot offer a non-arbitrary way of selecting the modal base over which *might* and *must* ranges, at least as long this way of selecting the modal base is supposed to yield truth-values which do justice to how the relevant claim of epistemic modality is assessed. In contrast, since the framework developed in this article does not care that much about assigning truth-conditions to epistemically modalised sentences, we can explain the different reactions in (5) and (6) by distinguishing between information states which admit ' $\Diamond\neg\phi$ ' (and thus ' $\Box\phi$ ') and those which support ' $\Diamond\neg\phi$ ' (and thus ' $\neg\Box\phi$ '). Accordingly, it is always the assessor's knowledge against which an epistemically modalised sentence is evaluated, and thus the current framework simply avoids the *ad hoc* decisions relativists have to make at some point.

Another notable feature of the semantic framework developed in this paper is the integration of basic tense operators into a dynamic framework for epistemic modality. This is an achievement, as it captures how epistemic modal operators and tense interact with the development of information states over time. The highlight is that updating fails to preserve epistemic commitment, and just in the right way: it leaves past epistemic possibilities untouched, but may eliminate current epistemic possibilities. This explains why learning something new may trigger withdrawal as well as retention, and thus why Alex may either admit that he was wrong or

²⁸At least the claim is not true from Mary's point of assessment. I assume, however, that relativists are not that interested in introducing non-classical truth-values, as it is a key motivation for relativism that it allows to capture some puzzling cases involving epistemic modality without digressing too much from classical truth-conditional semantics.

point to a commitment (that the keys might have been in the car) which remains untouched by the new information. To my knowledge, no other account is in the position to offer an equally straightforward explanation of the complicated data about post-update assessment.²⁹ Since the story I have told has at least as good things to say about other embeddings as the orthodox alternatives – contrary to what is so often said about non-truth-conditional approaches – I conclude that it is on the whole superior to the orthodox alternatives currently on the market.³⁰

References

- Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). “On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision”. *Journal of Symbolic Logic*, 50 (2): 510–530.
- Asher, N. (2007). “Dynamic Discourse Semantics for Embedded Speech Acts”. In Tsohatzidis, S. L., editor, *John Searle’s Philosophy of Language*, pages 211–244. Cambridge University Press, New York, NY.
- Asher, N. and Lascarides, A. (2009). “Another Look at Questions in Dialogue”. Forthcoming in Riester, A. and Solstad, T. editors, *Proceedings of SuB13*, Stuttgart.
- Asher, N. and McCreedy, E. (2007). “Were, Would, Might, and a Compositional Account of Counterfactuals”. *Journal of Semantics*, 24 (2): 93–129.
- Beaver, D. I. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford, CA.
- DeRose, K. (1991). “Epistemic Possibilities”. *Philosophical Review*, 100 (4): 581–605.
- Dever, J. (2009). “Living the Life Aquatic: Does Presupposition Accommodation Mandate Dynamic Semantics?”. <https://webpace.utexas.edu/deverj/personal/papers/aquatic.pdf>. Unpublished manuscript.
- van der Does, J., Groeneveld, W., and Veltman, F. (1997). “An Update on Might”. *Journal of Logic, Language and Information*, 6 (4): 361–380.
- Egan, A. (2007). “Epistemic Modals, Relativism, and Assertion”. *Philosophical Studies*, 133 (1): 1–22.

²⁹Stephenson (2007, 2008) can explain why an agent who now believes that ‘ $\neg\phi$ ’ may nevertheless insist that ‘ $P\Diamond\phi$ ’ is the case, since she allows for variation of the temporal index independently of the point of assessment. But there is no equally straightforward explanation for why it is also natural for the agent to withdraw a commitment, as her framework does not consider the dynamic development of information states over time and thus lacks the natural failure of Preservation which is so central to the story told here.

³⁰Thanks to the audience at the 2009 NYU/Columbia Graduate Student Philosophy Conference for comments and discussion. Special thanks to Nicholas Asher, Josh Dever, Hans Kamp, and Mark Sainsbury for their very useful comments.

- Egan, A., Hawthorne, J., and Weatherson, B. (2005). “Epistemic Modals in Context”. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy*. Oxford University Press, New York.
- von Fintel, K. and Gillies, A. (2008a). “An Opinionated Guide to Epistemic Modality”. In Gendler, T. and Hawthorne, J., editors, *Oxford Studies in Epistemology* 2. Oxford University Press, Oxford.
- von Fintel, K. and Gillies, A. (2008b). “CIA Leaks”. *Philosophical Review*, 117 (1): 77 – 98.
- von Fintel, K. and Gillies, A. (2009). “‘Might’ Made Right”. <http://web.mit.edu/fintel/fintel-gillies-2008-mm.pdf>. Forthcoming in Egan, A. and Weatherson, B., editors, *Epistemic Modality*, Oxford University Press.
- Frege, G. (1879). *Begriffsschrift: Eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Louis Nebert, Halle.
- Fuhrmann, A. (1989). “Reflective Modalities and Theory Change”. *Synthese*, 81 (1): 115–134.
- Gärdenfors, P. (1988). *Knowledge in Flux*. MIT Press, Cambridge, MA.
- Gillies, A. (2001). “A New Solution to Moore’s Paradox”. *Philosophical Studies*, 105 (3): 237–250.
- Gillies, A. (2004). “Epistemic Conditionals and Conditional Epistemics”. *Noûs*, 38 (4): 585–616.
- Gillies, A. (2006). “What Might be the Case after a Change in View”. *Journal of Philosophical Logic*, 35 (2): 117–145.
- Gillies, A. (2009a). “Iffiness”. <http://www.umich.edu/~thony/iffiness-revised-june2008.pdf>. Unpublished manuscript.
- Gillies, A. (2009b). “On Truth-Conditions for *If* (But Not Quite only *If*)”. <http://www.umich.edu/~thony/note-on-if-finalish-preprint.pdf>. Forthcoming in *Philosophical Review*.
- Groenendijk, J. (2008). “Inquisitive Semantics”. In *Proceedings of SALT XVIII*, Amherst, MA. The University of Massachusetts at Amherst.
- Groenendijk, J. and Stokhof, M. (1991). “Dynamic Predicate Logic”. *Linguistics and Philosophy*, 14 (1): 39–100.
- Groenendijk, J., Stokhof, M., and Veltman, F. (1996). “Coreference and Modality”. In Lappin, S., editor, *The Handbook of Contemporary Semantic Theory*, pages 179–214. Blackwell, Cambridge, MA.
- Hacking, I. (1967). “Possibility”. *Philosophical Review*, 76 (2): 143–168.

- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst.
- Kamp, H. (1981). “A Theory of Truth and Representation”. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Formal Methods in the Study of Language, Part I*, pages 277 – 320. Mathematisch Centrum, Amsterdam.
- Kamp, H. and Reyle, U. (1993). *From Discourse to the Logic: Introduction to Modaltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Press, Dordrecht.
- Kamp, H., van Genabith, J., and Reyle, U. (2009). “Discourse Representation Theory”. <http://www.ims.uni-stuttgart.de/~hans/Papers/hpl-drt.pdf>. Unpublished manuscript.
- Kant, I. (1781). *Critik der Reinen Vernunft*. Johann Friedrich Hartknoch, Riga.
- Kratzer, A. (1977). ‘What ‘Must’ and ‘Can’ Must and Can Mean’. *Linguistics and Philosophy*, 1 (3): 337–355.
- Kratzer, A. (1981). “The Notional Category of Modality”. In Eikmeyer, H.-J. and Rieser, H., editors, *Words, Worlds, and Contexts. New Approaches to Word Semantics*, pages 38–74. Walter de Gruyter & Co., Berlin.
- Kratzer, A. (1991). “Modality”. In von Stechow, A. and Wunderlich, D., editors, *Semantics: An International Handbook of Contemporary Research*, pages 639 – 650. Walter de Gruyter & Co., Berlin.
- Levi, I. (1988). “Iteration of Conditionals and the Ramsey Test”. *Synthese*, 76 (1): 49–81.
- MacFarlane, J. (2009a). “Epistemic Modals Are Assessment-Sensitive”. <http://johnmacfarlane.net/epistmod.pdf>. Forthcoming in Egan, A. and Weatherson, B., editors, *Epistemic Modality*, Oxford University Press.
- MacFarlane, J. (2009b). “Non-Indexical Contextualism”. <http://johnmacfarlane.net/nonindexical-contextualism.pdf>. Forthcoming in *Synthese*.
- Price, H. (1983). “Does ‘Probably’ Modify Sense?”. *Australasian Journal of Philosophy*, 61 (4): 396–408.
- Ramsey, F. P. (1990). “General Propositions and Causality”. In Mellor, D. H., editor, *Philosophical Papers*, pages 145–163. Cambridge University Press, Cambridge.
- Roberts, C. (1987). *Modal Subordination, Anaphora, and Distributivity*. PhD thesis, University of Massachusetts at Amherst.
- Roberts, C. (1989). “Modal Subordination and Pronominal Anaphora in Discourse”. *Linguistics and Philosophy*, 12 (6): 683–721.

- Rott, H. (1989). “Conditionals and Theory Change: Revisions, Expansions, and Additions”. *Synthese*, 81 (1): 91–113.
- Sadock, J. M. and Zwicky, A. M. (1975). “Ambiguity Tests and How to Fail Them”. In Kimball, J. P., editor, *Syntax and Semantics 4*. New York Academic Press, New York, NY.
- Stalnaker, R. C. (1978). “Assertion”. In Cole, P., editor, *Syntax and Semantics 9*, pages 315 – 332. New York Academic Press, New York, NY.
- Stephenson, T. (2007). *Towards a Theory of Subjective Meaning*. PhD thesis, Massachusetts Institute of Technology.
- Stephenson, T. (2008). “Judge Dependence, Epistemic Modals, and Predicates of Personal Taste”. *Linguistics and Philosophy*, 30 (4): 487–525.
- Teller, P. (1972). “Epistemic Possibility”. *Philosophia*, 2: 303–320.
- Veltman, F. (1985). *Logics for Conditionals*. PhD thesis, University of Amsterdam.
- Veltman, F. (1996). “Defaults in Update Semantics”. *Journal of Philosophical Logic*, 25 (3): 221–261.
- Yalcin, S. (2007). “Epistemic Modals”. *Mind*, 16 (464): 983–1026.
- Yalcin, S. (2009). “Nonfactualism about Epistemic Modals”. <http://mit.edu/yalcin/www/nfem.pdf>. Forthcoming in Egan, A. and Weatherson, B., editors, *Epistemic Modality*, Oxford University Press.