

Dynamics of Epistemic Modality*

Malte Willer

Abstract

A dynamic semantics for epistemically modalized 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 theory 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 alternative theories. The second part of the story extends the basic framework to cover some complicated data about retraction and the interaction between epistemic modality and tense. A comparison between the suggestion made in this paper and current versions of the orthodoxy is provided.

1 Introduction

Orthodox semantics assigns to epistemically modalized 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 habit of evaluating present tense claims of epistemic modality by testing them against their own perspective. Consider the following example, inspired by [von Fintel and Gillies \(2008a\)](#). 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.

*Published in *Philosophical Review* 122(1): 45–92, 2013

One extraordinary version of the orthodoxy—the relativist version—is that epistemically modalized sentences vary in truth-value across points of assessments (a judge or some other point that varies with the assessor). So Alex’s utterance is literally true when he is the judge and false when Mary is (and vice versa for Mary’s utterance).¹ Another radical position—the most recent contextualist approach to epistemic modals—is that we have to give up the idea of a unique contextually determined proposition expressed by a judgment 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.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) causes any problems. From a naive perspective, the case seems completely innocent. Mary rejects Alex’s judgment 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 dispute about matters of fact there is a difference in what Alex and Mary know. This difference has a *pragmatic* but no *semantic* effect on the discourse: it is relevant for how Alex and Mary *assess* the claim that the keys are in the car, not for what they *say* by asserting that the keys are/are not in the car. This is why Mary can evaluate Alex’s judgment against her own perspective without risk of misinterpreting what he said. The naive perspective

¹See Egan et al. (2005), Egan (2007), MacFarlane (2011), Stephenson (2007, 2008).

²See von Stechow and Gillies (2011). Kratzer (1977, 1981, 1991) is the *locus classicus* of contextualism (Kratzer (2012) offers an up-to-date discussion of her views). See also DeRose (1991), Hacking (1967), and Teller (1972) for some letters of support for this view.

³In the present discussion I have omitted a view that has served as the whipping boy in modern debates about epistemic modality, namely that epistemically modalized sentences are used to describe what is known by a contextually salient group of individuals (in most cases involving the speaker). See Egan et al. (2005), §2 and MacFarlane (2011), §3 for critical discussion and Dowell (2011) for a defense. Part of the exercise here is to demonstrate there is simply no need for non-solipsistic contextualism to make sense of disputes about epistemic modality.

assumes that differences in what is known play exactly the same role when it comes to judgments of epistemic modality: what one knows determines how one assesses a judgment of epistemic modality, but does not matter for what one says by making such a judgment. Everybody, I think, can agree that the naive conception is very attractive: it offers a simple and uniform perspective on disputes about epistemic modality and disputes about matters of fact and, what is more, creates no need for relativism or some new age contextualism.

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 *do* have a semantic effect on Alex's and Mary's dispute about 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 judgment that 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 judgment about Mary's epistemic situation, thus legitimatizing Mary's rejection of Alex's claim on the basis of what she knows. But none of this would be necessary were we to accept the naive view that one's epistemic situation is only relevant for how one assesses a judgment of epistemic modality, and not for what one says by making such a judgment.

What I have done so far is to contrast two possible views about the role of one's epistemic situation for judgments of epistemic modality. On the naive conception the role is only of a pragmatic kind: what one knows is relevant for how one assesses a judgment of epistemic modality, period. On the orthodox view, what one knows is also of relevance for the semantic content of a judgment of epistemic modality. I have also outlined some motivations for adopting the naive conception, and my aim is to show that it indeed articulates a viable alternative to orthodox semantics. This should sound a bit surprising at first, since the naive conception seems to require that epistemically modalized judgments aim at describing facts that are independent of anyone's epistemic situation, on par with a judgment stating that the keys are in the car. I admit that it is difficult to see what such facts could be, but this only means that one should not use the naive conception as the starting point of a truth-conditional semantics for epistemic modals. Accordingly, my semantic story will run under the slogan that epistemically modalized sentences have content, but not truth-conditional content. Such a story may sound even more radical than current versions of the orthodoxy, but as I hope to show in the following sections, it offers a simple and intuitive perspective on the semantics of epistemic modals.

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 of disagreement and to dispel some old prejudices against non-truth-conditional approaches to the semantics of epistemic modals. In §3 I show how the framework developed in this paper takes care of some further tricky data. §4 extends the basic framework with a semantics for tense and explains the various ways in which we may assess

judgments of epistemic modality in the light of new evidence. The final §5 offers a more in-depth comparison between the framework developed in this paper and the orthodox perspective.

2 Basics

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 The section is structured as follows. In §2.1 I give an informal outline of my semantic proposal. The formal details are elaborated in §2.2. §2.3 shows how, given some very simple pragmatic assumptions, the framework can provide the desired uniform perspective on modal and factual disputes. The final §2.4 addresses some additional foundational issues about the dynamic semantic framework developed in this paper, including the possibility of embedding epistemic modals.

2.1 Outline

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 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 relational: they are relations between contexts.⁵ Of course, *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 modalized sentences have context change potential, 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.

⁴Some popular dynamic semantics: Discourse Representation Theory ([Kamp \(1981\)](#); [Kamp and Reyle \(1993\)](#); [Kamp et al. \(2011\)](#)), Dynamic Predicate Logic ([Groenendijk and Stokhof \(1991\)](#)), File Change Semantics ([Heim \(1982\)](#)), Update Semantics ([Veltman \(1985, 1996\)](#)).

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 or minus a bit, the route taken by [Dever \(2006\)](#), §1 or [von Stechow and Gillies \(2008b\)](#), §6.

With so much focus on contexts, it is legitimate to ask what exactly they are supposed to be. The answer depends on the specific goals of one’s dynamic proposal. Since I am primarily interested in the interaction between epistemically modalized sentences and an agent’s state of mind, I will treat contexts as information states.⁶ The resulting approach will offer an appealing perspective on our observation about how utterances of epistemic modality are assessed by different people. The effect of an utterance may vary from information state to information state, and so 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 a position to offer a simple and uniform account of how speakers assess utterances in discourse.

My view then is that epistemically modalized sentences have content in virtue of their context change potential, which in turn is understood as a relation between information states. The more specific suggestion starts with the observation that speakers frequently use *might*-statements to highlight possibilities that they think the hearer should take seriously. The point of Alex’s statement in (1), for example, is to bring Mary to seriously consider the possibility that the keys are in the car. So my suggestion is that *might*-statements are designed to affect an agent’s information state by highlighting the theoretical and practical importance of certain possibilities. To introduce a bit of terminology: *might*-statements are designed to change possibilities that are merely compatible with the agent’s evidence into “live possibilities”—possibilities that are compatible with the agent’s evidence and that the agent takes seriously in inquiry.

Taking a possibility seriously is a dispositional affair: an agent need have no *occurrent* attitude toward some possibility p to take p seriously, and usually will in fact have no such attitude. What is required, instead, is a disposition to take the possibility of p ’s being true into serious consideration whenever it is of practical or theoretical pertinence. To make this more precise, think of inquiry as an attempt to arrive at an answer to an open question. Each open question allows for more than one answer that is compatible with the available evidence, but often the inquiring agent only considers a subset of those as the set of *relevant* alternatives. If p belongs to the set of relevant alternatives, then the inquiring agent will not adopt an answer that is incompatible with p before the agent has eliminated p as a possibility. If p is not a relevant alternative, then the inquiring agent effectively ignores p and may adopt an answer that is incompatible with p without having eliminated p as a possibility. So to give an

⁶It is no coincidence that Veltman (1985, 1996) treats contexts in a way that is similar to mine, since my proposal bears some resemblance to his Update Semantics. However, the differences between the upcoming proposal and Update Semantics are substantial, primarily because the latter cannot explain some crucial observations about the role of epistemic modals in information exchange. So the reader may wish to take a look at Veltman’s original work, and also take a look at Beaver (2001), Does et al. (1997), Gillies (2001), and Groenendijk et al. (1996) for some applications and extensions of Veltman’s proposal. This proposal is also the one favored by von Fintel and Gillies (2008b), but does not play a role in their later work on epistemic modals. The domain semantics for epistemic modals developed by Yalcin (2007, 2011) is, from an abstract point of view, a static version of Update Semantics and as such alike to Veltman’s proposal in its virtues and limitations, some of which we will see momentarily.

example with a familiar ring, consider an agent who sees an equid with black and white stripes and who wonders what kind of animal he or she sees. In most cases, the agent will ignore the possibility that the animal is a cleverly disguised mule and thus will see no reason not to accept the answer that the animal is a zebra on the basis of its outer appearance. But if the agent treats the possibility that the animal is a cleverly disguised mule as a relevant alternative, then the agent will not accept the answer that the animal is a zebra until he or she has eliminated the possibility that it is a cleverly disguised mule, which is just to say that the agent will not arrive at that answer just on the basis of the animal's outer appearance.⁷

My simple suggestion, then, is that p is a live possibility for an agent S — S takes p seriously in inquiry—just in case S is disposed to treat p as a relevant alternative whenever S is concerned with a question to which p is an answer. Combined with the dynamic semantic proposal that was articulated earlier, this view predicts—correctly, I think—that *might*-statements are frequently used with the intention of *expanding* the set of relevant alternatives from which an audience selects an answer to a question that is currently under consideration. It should then be unsurprising that such statements can be used to undermine the audience's confidence in a previously accepted answer to the question: that answer may very well conflict with a newly introduced alternative that is not conclusively ruled out by the audience's evidence.

The working assumption of my semantic proposal, then, is that inquiring agents distinguish between live possibilities and possibilities that are merely compatible with their evidence (which I will sometimes call “plain” or “mere” possibilities); *might*-statements are designed to change possibilities of the latter kind into live possibilities. The proposal does not require a detailed story about *how* agents draw this distinction in inquiry, and so I restrict myself to stating the obvious: there are various reasons an agent may have for taking a possibility seriously, including practical concerns, but it is safe to say that in at least many cases agents decide which possibilities to take seriously on the basis of general considerations of plausibility. Some possibilities are really far-fetched while others seem to be more reasonable, and in general only possibilities that meet a certain standard of plausibility play a significant role in one's practical and theoretical deliberation. One immediate consequence is that the possibilities an agent is aware of need not coincide with his or her live possibilities, as one may be aware of a possibility that, at least in the absence of additional information, does not appear sufficiently plausible to be taken seriously.

The preceding considerations show that we can make good intuitive sense of the distinction between an agent's live possibilities and those possibilities that are merely compatible with the agent's evidence, but we still need to decide how

⁷The terminology I use here is, of course, very familiar from relevant alternative theories à la Dretske (1970, 1981) and Lewis (1996), but I am not advancing any claims about knowledge. The point is: that inquiring agents distinguish between relevant and irrelevant alternatives is important for our best theory of rational inquiry and a dynamic semantics for epistemic modals. Relevant alternative theories about knowledge are consistent with but certainly not entailed by this observation.

to capture what we have said so far in a formal representation of information. [Stalnaker \(1984\)](#) models epistemic and doxastic attitudes as attitudes toward possible states of the world, and it is straightforward to adopt this strategy for our purposes. Start with an idea that was already implicit in what I said earlier: that an agent’s inquiry is guided by a conception of which possible worlds matter for inquiry. We may then say that to take p seriously is to be in a state that *includes* a p -world within the set of possible worlds that matter for inquiry. To rule out p , in turn, is to be in a state that *excludes* all p -worlds from the set of possible worlds that matter for inquiry.

The approach outlined so far preserves the familiar conception of attitudes in inquiry as attitudes toward possible states of the world. The only wrinkle we have to add to the story is that a state of information will often fail to *completely* specify which possible worlds matter for inquiry. An agent who ignores a possibility p even though p is compatible with his or her evidence fails to take p seriously but does not rule out p either. The model I have suggested predicts that the agent’s state of information leaves it *unspecified* whether some p -world matters for inquiry: it does not include a p -world within the set of possible worlds that matter for inquiry since p is not taken seriously; it does not exclude all p -worlds from the set of possible worlds that matter for inquiry since p is not ruled out. This result is in principle unproblematic, but it follows that we cannot simply identify an agent’s information state with a distinct set of possible worlds that the agent recognizes in inquiry. The good news is that we already know from the literature on supervaluationism how to model a merely incomplete specification of a set of possible worlds.⁸ Let me explain.

The suggestion is to represent an agent’s state of information as the set of sets of possible worlds that are in accordance with the agent’s conception of what matters in inquiry, and then to capture the features of an information state that interest us in a supervaluationist fashion.⁹ The idea is simple: whenever an agent thinks that a p -world matters for inquiry, every set of possible worlds that is in accordance with the agent’s conception of what matters in inquiry will include a p -world. So we say that a state of information Σ *includes* a p -

⁸Applications of supervaluationist techniques to various philosophical topics can be found in discussions by [Fine \(1975\)](#), [van Fraassen \(1966\)](#), [Kamp \(1975\)](#), [Stalnaker \(1981\)](#), and [Thomason \(1970\)](#), among many others. [Beaver \(2001\)](#), §9, recognizes the importance of supervaluationist techniques in modeling an agent’s state of mind, albeit for purposes that are very different from mine: his interest is to account for the possibility that participants in a discourse may sometimes not know what the common ground is. See also [Rothschild \(2012\)](#), who models mental states as sets of probability functions, and [Yalcin \(2012\)](#), who models information states as sets of probability spaces.

⁹To see that this suggestion also makes perfect formal sense, notice that we can think of an information state as a potentially gappy function f specifying a set of possible worlds by mapping possibilities to the values 0 and 1, where p is mapped to 1 in case the set contains at least one possible world at which p is true, and p is mapped to 0 in case the set contains no such world. Every complete extension of f that results from filling in its gaps in one way or another then identifies a distinct set of possible worlds, and, furthermore, this set meets all the constraints carried by f in the sense that it is guaranteed to include a p -world if f maps p to 1 and guaranteed to include no p -world in case f maps p to 0. Questions about the properties of an information state can then be meaningfully reduced to questions about properties of its complete extensions in the familiar supervaluationist fashion.

world within the set of possible worlds that matter for inquiry just in case *every* element of Σ contains a p -world. Similarly, we say that Σ *excludes* all p -worlds from the set of possible worlds that matter for inquiry just in case *no* element of Σ contains a p -world. In combination with the possible worlds analysis of an agent’s attitudes that was suggested earlier, this leads to the following formal setup: if an agent S is in a state Σ , then p is a live possibility for S just in case every element of Σ contains a p -world. p is eliminated by S in inquiry just in case no set of possible worlds in Σ contains a p -world. And whenever S fails to take p seriously but does not rule out p either, then some but not all elements of Σ contain a p -world, i.e. Σ fails to specify whether or not the set of possible worlds that matter for inquiry includes a p -world.

Our supervaluationist model draws the distinctions we need while preserving the familiar approach of modeling an agent’s attitudes in inquiry as attitudes toward possible states of the world.¹⁰ It also allows us to formalize the process of highlighting possibilities as live possibilities in a simple and intuitive fashion. To raise p from a mere possibility to a live possibility is to transform a state that fails to specify whether a p -world matters for inquiry to a state that includes a p -world within the set of possible worlds that matter for inquiry. Accordingly, the process of highlighting a possibility p as a live possibility is to be modelled as a process of eliminating, from the state under consideration, all elements that fail to contain at least one p -world. The fact that such operations are very easy to handle from a technical perspective will prove to be advantageous once we look at the formal details of the proposal, to which I will turn now.¹¹

¹⁰There are, of course, alternative ways of modeling the distinction between plain possibilities and live possibilities. Yalcin (2011), for instance, models information states as sets of possible worlds but adds a partition as a bit of additional structure. This approach can in principle be combined with a dynamic approach to the semantics of epistemic modals: the role of *might*-statements is then modelled as similar to the role of questions in that they raise issues, with elements of a partition representing possibilities that are live for an inquiring agent. One major problem with this alternative is that even if both ϕ and $\neg\phi$ are compatible with an agent’s evidence, the agent may take only *one* of these possibilities seriously—this is hard to explain if we use partitions in modeling live possibilities. Another major problem is that issue raising operations do not embed at all or at least not in a way that suits our purposes, despite the progress that has been recently made toward a semantics that allows for embeddings of questions by, e.g., Asher (2007), Asher and Lascarides (2009), and Groenendijk (2009). On the other hand, and as an anonymous reviewer correctly observes, using partitions in identifying an agent’s live possibilities may have a potential advantage over the method endorsed here since the latter, but not the former, predicts that live possibilities are closed under logical entailment. Notice, however, that it is open for us to appeal to partitions in modeling the possibilities that an agent treats as *relevant* in inquiry, i.e. those possibilities that are live and pertain to issues currently addressed by the agent. The possibilities that an agent actively considers in inquiry are then correctly predicted not to be closed under logical entailment, and so an agent’s live possibilities being closed under logical entailment does not lead to implausible predictions about the way agents actually proceed in inquiry.

¹¹I am grateful to an anonymous reviewer for pressing me on the informal motivations behind the technical apparatus to be presented in the next section, and in particular for the pointer to supervaluationism as a potentially helpful tool for illuminating some important aspects of my approach. The usual disclaimers apply.

2.2 Details

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

Definition 1 (Language) \mathcal{L} is the smallest set that contains a set of 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 suggestion is to understand the meaning of a sentence in terms of its context change potential, with contexts being treated as information states. Such states need to keep track of the possibilities left uneliminated by an inquiring agent, but to suit our purposes they also must identify which of those possibilities the agent takes seriously. The framework developed here stays faithful to the familiar strategy of modeling an agent's attitudes in inquiry as attitudes toward possible states of the world. A state of information represents the agent's conception of which possible states of the world matter for inquiry and is modeled as the set of sets of possible worlds in accordance with that conception:

Definition 2 (Possible Worlds, Information States) w is a *possible world* iff $w: \mathcal{A} \rightarrow \{0, 1\}$. W is the set of such w 's, $\mathcal{P}(W)$ is the powerset of W . Σ is an *information state* iff $\Sigma \subseteq (\mathcal{P}(W) \setminus \{\emptyset\})$, 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 *absurd* information state Σ_{\emptyset} is identical with \emptyset .

For each element of \mathcal{L}_0 we can define a set of indices at which it is true. Such sets of indices will not play their usual role as carriers of meaning but figure in the formalization of our notion of a (live) possibility:

Definition 3 (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) $\llbracket \phi \wedge \psi \rrbracket = \llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket$

In drawing the distinction between plain and live possibilities, the crucial question is which possible worlds an agent's state of information includes within the region in logical space that matters for inquiry, and which it excludes from that region. The analysis is provided in a supervaluationist fashion:

Definition 4 (Possibilities) Consider any $\Sigma \in I$ and $\phi \in \mathcal{L}_0$:

1. ϕ is a *possibility* in Σ iff $\exists \sigma \in \Sigma \exists w \in \sigma: w \in \llbracket \phi \rrbracket$
2. ϕ is a *live possibility* in Σ iff $\forall \sigma \in \Sigma \exists w \in \sigma: w \in \llbracket \phi \rrbracket$

For ϕ to be a possibility for an agent, all it takes is that the agent’s conception of what matters in inquiry is compatible with the existence of a ϕ -world. For ϕ to be a live possibility for an agent, it is required that the agent’s conception of what matters in inquiry entails the existence of a ϕ -world.¹² Notice that while the proposed model of an information state is a bit non-classical, it does not stray too far away from the familiar approach to modeling states of information. Specifically, we preserve the conception of an information state as a representation of the way things are. Given an information state Σ , we may legitimately ask whether the information it carries excludes the actual world w_a , 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 it carries 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.

The general idea is to model the meaning of formulas in \mathcal{L} in terms of their update effects on states of information. The more specific suggestion is that such formulas are designed to refine an agent’s conception of what matters in inquiry either by eliminating certain possibilities from inquiry in a way familiar from Stalnaker’s discussion of assertion or—the new wrinkle—by introducing hitherto ignored possibilities as live possibilities. The formal implementation of this suggestion is straightforward. We first 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 5 (Updates on Elements of Information States) Consider any $\sigma \subseteq W$, $p \in \mathcal{A}$ and $\phi, \psi \in \mathcal{L}$. An update on an element of an information state is a function $\uparrow: \mathcal{P}(W) \rightarrow \mathcal{P}(W)$ defined by the following recursion:

- (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 σ with an atom p eliminates all possible worlds from σ in which p is false. According to clause (2), updating σ with $\lceil \neg\phi \rceil$ comes down to leaving everything in σ that gets eliminated by an update with ϕ . To update with a conjunction, update with the first conjunct and then update the result with the second conjunct (cf. (3)). Clause (4) captures a

¹²Notice here that S ’s conception of what matters in inquiry need not entail that the existence of a *particular* ϕ -world.

test-conception of claims of epistemic modality. Updating σ with a formula of the form $\lceil \diamond \phi \rceil$ is to run a test: if updating σ with ϕ does not return the empty set, then σ passes the test. Otherwise, we get back the empty set, and we do not need to consider σ 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.¹³

We can now describe what impact 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 relational: they are relations between information states.

Definition 6 (Updates on Information States) Consider any $\Sigma \in I$ and $\phi \in \mathcal{L}$. An update on an information state is a function $[\cdot] : I \rightarrow I$ 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 sets of possible worlds together, leaving out the empty set. This yields the output information state. The following notions will turn out useful in later discussion:

Definition 7 (Settledness, Admission, Entailment, Equivalence) 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 $\forall \Sigma : \Sigma[\phi] \models \psi$
4. ϕ and ψ are *equivalent*, $\phi \cong \psi$, iff, $\forall \Sigma : \Sigma[\phi] = \Sigma[\psi]$

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

We are now in a position to lay out a few important features of the current framework. First, we find that for all Σ and $\phi \in \mathcal{L}_0$, Σ supports $\lceil \diamond \phi \rceil$ just in case ϕ is a live possibility in Σ :

Fact 1 For all Σ and $\phi \in \mathcal{L}_0$: $\Sigma \models \diamond \phi$ iff ϕ is a live possibility in Σ

¹³These rules and in particular the test-conception of epistemically modalized sentences, lie at the heart of Veltman's Update Semantic. As an anonymous reviewer remarks, the semantic proposal for *might* can be traced back to an informal proposal articulated by Stalnaker (1970). We will see momentarily that what has been said so far is only one component among many others in a complete dynamic story about *might* and *must*.

Since $\Sigma[\diamond\phi] \models \diamond\phi$, we get the desired result that an update with an epistemically modalized sentence raises the prejacent to a live possibility. Specifically, any admissible update with $\ulcorner \diamond\phi \urcorner$ raises ϕ from a plain possibility to a live possibility.

Judgments of epistemic modality and factual judgments stand in the right logical relations to one another. 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 judgment of epistemic necessity:

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

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

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

If ϕ is a live possibility in Σ , Σ does not admit the judgment 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 one's information state supports ϕ it also supports $\ulcorner \Box\phi \urcorner$; but it does not follow that $\ulcorner \neg\phi \urcorner$ is settled whenever the falsity of ϕ is a live possibility.

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

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

As in S5, any combination of boxes and diamonds reduces to the innermost operator. This is an acceptable result, I think, as embedding an epistemic modal under another epistemic modal does not in general have any interesting semantic effects.¹⁴ It is also of relevance for the explanatory adequacy of the theory developed so far. As we have seen, the framework makes good sense of the semantics of epistemically modalized sentences whenever the modal takes scope over an element of \mathcal{L}_0 : *might* highlights the prejacent as a live possibility, and *must* is just the dual of *might*. Since the result of embedding such sentences under another epistemic modal is just equivalent to the embedded sentence, we can conclude that the semantic proposal generalizes to cover epistemically modalized sentences of arbitrary complexity.

This is all I have to say about the semantics of \mathcal{L} .¹⁵ And with a bit of

¹⁴Sorensen (2009) presents examples that purport to demonstrate that no adequate logic for epistemic modals can validate the characteristic axiom of S5. I am not convinced, but the cases Sorensen presents raise issues that are too complex to be discussed here. See Yalcin (2009) for a response to his critique.

¹⁵At 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

pragmatic underpinning, it is all we need to deliver the promised perspective on disputes about epistemic modality.

2.3 Disputes

First, here is an appealingly simple picture about pragmatics. 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 context change potential, 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 S with information state Σ_S . Then S will by default assess an utterance of ϕ as follows:

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

In words, S will agree with ϕ if S 's information state already encodes the information encoded in ϕ . If S is agnostic about ϕ , then we expect S to admit that ϕ is the case. And finally, if S 's information is incompatible with ϕ , then we expect S to reject an assertion of ϕ . The last bit of the suggestion—that agents by default reject an assertion of ϕ in case they cannot add ϕ consistently to what they believe—makes good sense since rational belief change must aim at preserving consistency. Of course, it is anything but unheard of that agents revise their beliefs to accommodate a conflicting bit of information ϕ but this does, at least in general, require *weakening* of one's information state to avoid inconsistency: information needs to be retracted to arrive at an information state that can be consistently updated with ϕ . Weakening, in turn, is very complex—deciding which beliefs to give up and which to retain in accommodating a conflicting bit of information is anything but a trivial affair—and in any case it occurs at the cost of information loss,¹⁶ and so a rational agent will avoid this procedure—will reject the conflicting bit of information—unless there are good reasons to do otherwise (such as: the conflicting bit of information comes

that *might* and *must* are duals. The fact that *might* and *must* behave differently when it comes to modal subordination (Roberts (1987, 1989)) may count as evidence 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).

¹⁶For a recent survey of the literature on belief revision (including the one addressing issues of computational complexity) see Rodrigues et al. (2011). The idea that information loss is a cost to the inquirer who changes a belief state via weakening is articulated by Levi (1996). As we will see momentarily, an agent who is in an information state Σ such that $\Sigma[\phi] = \emptyset$ will not always be required to weaken Σ to accommodate ϕ , and in such cases revision does not have its usual costs. I will come back to this issue in §3.2.3.

from a source that the agent believes to know more about the issue than he or she does).

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 since 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 is not relevant for *what* they say when they make their utterances, but for *why* they say it. Let p stand 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 has a *pragmatic* but no *semantic* effect on epistemic modals: 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 modalized sentence. And we can hold this view because we have a semantics that does not ascribe truth-conditions to such sentences. Instead, the content of epistemically modalized sentences is described in terms of how they affect contexts, here understood as information states.

2.4 Embeddings and Related Issues

In §2.1 I outlined Stalnaker’s picture about context-content interaction. This picture is truth-conditional at heart since all context change is mediated by truth-conditional content. It is important to notice that Stalnaker’s view remains intact as long as we only consider the non-modal fragment \mathcal{L}_0 of our toy language. Updating an element of an information state with some $\phi \in \mathcal{L}_0$ comes down to adding the proposition associated with ϕ to that element:

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

But all of this is quite consistent with the claim that we cannot understand the context change that is initiated by an epistemically modalized sentence in a Stalnakerian fashion. Updating with such a sentence eliminates either all or no possibilities, depending on global features of the state.¹⁷ Hence we will not find some proposition that mediates the kind of context change that is induced by an epistemically modalized sentence. If it were not for these constructions,

¹⁷Another way of putting the point is that unlike updating σ with an element of \mathcal{L}_0 , updating σ with ‘ $\diamond\phi$ ’ fails to be *distributive*: $\sigma \uparrow \phi \neq \bigcup_{w \in \sigma} \{w\} \uparrow \phi$. The same applies to updates with *must*.

then, there would be no departure from the foundations of truth-conditional semantics.

MacFarlane (2011) articulates the common worry that a non-truth-conditional semantics for epistemically modalized sentences is not viable since such sentences embed (though not as freely as factual sentences). 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.¹⁸ Such theories are clearly in bad shape when it comes to handling embeddings but 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.¹⁹ So there is no reason at all to believe that there is any embedding problem for the framework developed so far. This constitutes a major advantage of the proposal developed here over other non-truth-conditional semantic frameworks for epistemic modals.

One may object that the embedding problem is not really solved since we also know that epistemically modalized sentences may occur under the scope of attitude ascriptions. Consider:

- (3) a. Alex believes that the keys might be in the car.
- b. Mary does not believe that John must be in Chicago.

The fact that epistemically modalized 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. The problem, however, is unreal. What is true is that if such verbs denote relations between individuals and truth-condition determining propositions, then every semantics will be in trouble that does not assign to epistemically modalized sentences such propositions. But we may instead understand attitude verbs as denoting relations between individuals and CCPs.²⁰ 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. Elements of information 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

¹⁸This approach to epistemic modality operators as comments on propositions is certainly the standard non-truth-conditional view. See, among many others, Kant (1781), Frege (1879), Price (1983). Yablo (2011) is a notable exception—in fact, his semantics for *might* and *must* has a distinct dynamic flavor—but the theory faces unresolved embedding problems as well.

¹⁹And of course, it is straightforward to combine our dynamic story for *might* and *must* with the no less dynamic story for *if* that Gillies (2004, 2009, 2010) puts to good use for various purposes. The outcome of this exercise is not immediately relevant for current purposes, and so I omit the details here.

²⁰If 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. I am primarily interested in how the conception of attitude verbs as denoting relations between individuals and CCPs may help us resolve a quite specific embedding problem for non-truth-conditional semantics. See Heim (1992) and Zeevat (1992) for a discussion of how this conception may lead to an adequate theory of presuppositions in attitude contexts.

attitude verbs denote relations between individuals and CCPs, they are of type $\langle\langle\langle s, t \rangle, t \rangle, \langle\langle\langle s, t \rangle, t \rangle, t \rangle\rangle, \langle e, t \rangle\rangle$. The only notable 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.

So far I have treated the belief-relation as basic, but we might instead analyze belief as a matter of what an agent’s information state as whole supports. On this view, an agent S who is in an information state Σ believes ϕ just in case adding ϕ to Σ does not induce a change. Thus our lexical semantics would associate the following denotation with *to believe*:

$$(4) \quad [\text{believe}]^i = \lambda R. [\lambda x. \exists \Sigma: x \text{ is in } \Sigma \text{ at } i \text{ and } \langle \Sigma, \Sigma \rangle \in R]$$

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 . This fits nicely with our treatment of epistemic modals since it allows us to predict that, for instance, Alex believes that the keys might be in the car just in the case the keys being in the car is a live possibility for him. The semantics of belief ascriptions 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. The moral: while it is compatible with the current proposal to take the belief-relation as basic, it is also straightforward to provide a non-trivial analysis in terms of information states and CCPs.

The strategy outlined here for handling the notion of belief in dynamic semantics can be extended to other important attitudes in inquiry, but I will not do this here. Instead, let me do a bit more to dispel the worry that something important gets lost in moving from a classical truth-conditional semantics to a dynamic semantics for epistemic modals. A remaining concern is that even if dynamic semantics can in principle make sense of the notion of belief, it may be unable to distinguish between correct and incorrect beliefs (or speech acts) about what might or must be the case. The source of concern is that our dynamic setup does not ascribe to such beliefs some propositional content that can be categorized as true or false, and it is not obvious what other criterion one should appeal to in distinguishing between correct and incorrect beliefs.

In response, it makes sense to categorize epistemic modal beliefs as correct or incorrect even if they do not have propositional content. The key observation is that while epistemic modal beliefs are not themselves true or false, updating an information state with their content frequently induces factual commitments that can be evaluated in a classical fashion. So instead of characterizing the correctness of an agent’s epistemic modal belief directly in terms of truth, we will take a slightly indirect route and say that such a belief is correct just in case adopting the belief under consideration would not commit someone equipped with the agent’s evidence to factual error. And this correctness criterion has

some real bite. It predicts, for instance, that an agent’s belief that ϕ must be the case is incorrect in case ϕ is false: *must* is strong— $\lceil \Box \phi \rceil$ entails ϕ —and thus a belief that ϕ must be the case commits the agent to the belief that ϕ is the case.²¹ Another result is that an agent’s belief that ϕ might be the case is incorrect in case ϕ is incompatible with the agent’s evidence: someone who is equipped with that evidence and adopts the belief that ϕ might be the case is guaranteed to end up in the absurd information state and thus committed to a contradiction. Some further predictions—in particular, the result that a belief that ϕ might be the case is not guaranteed to be correct just because ϕ is compatible with the believer’s evidence—will turn out to be of relevance in the next section.

Correctness as characterized above turns on an individual’s evidential situation and thus we leave room for the possibility that, given adequate variation in what is known, Alex correctly believes that the keys might be in the car, while Mary correctly believes that the keys cannot be in the car. This, I think, is a desirable result, and it is also obvious how to generalize the proposal so that it covers speech acts and judgments more generally. But we are well advised to bear in mind that correctness thus understood cannot serve as a reliable guide to the semantics and pragmatics of epistemic modals. Such caveats aside, the proposal usefully complements the dynamic story told so far in that it delivers a reasonable correctness criterion for epistemic modal beliefs. The fact that this is possible without directly ascribing to epistemic modal beliefs propositional content should make us confident that moving away from the orthodoxy does not have any negative side effects.²²

3 Bonus

The dynamic framework developed so far offers an intuitively appealing semantics for epistemic modals that easily accounts for our intuitions about cases of modal disagreement, does not suffer from Frege-Geach style problems, and can account for the possibility of (correct or incorrect) epistemic modal beliefs. In addition to that, it can explain a range of data about epistemic modals that go

²¹For a recent defense of the view that *must* is strong, I recommend the discussion by [von Fintel and Gillies \(2010\)](#), who also look at an issue that is important but not immediately relevant for our purposes: the evidential dimension of epistemic modals.

²²[Yalcin \(2011\)](#) offers correctness criteria for epistemic modal speech acts that ask whether someone equipped with a certain body of evidence would be reacting appropriately to that evidence by adopting the content of the speech act under consideration, where the body of evidence is either the speaker’s or the one of someone with full information about the relevant situation. Obviously, these criteria are not too different from what I have said about the correctness of epistemic modal beliefs, though Yalcin does not specify what exactly it takes to react appropriately to one’s evidence in adopting the content of a speech act. And clearly, it would be unproblematic to take a bit of inspiration from Yalcin and combine the correctness criteria developed here with one that evaluates epistemic modal beliefs in light of the evidence of someone equipped with full information about the relevant situation. The pros and cons connected with this move are not immediately relevant to the issues addressed here, and so I will set this option aside.

beyond issues of disagreement. I will first describe the observations (§3.1) and then explain how the framework accounts for them (§3.2).

3.1 Data

3.1.1 Epistemic Contradiction

One observation that has received a lot of attention in recent discussion is that there is an interesting logical connection between factual judgments and present tense claims of epistemic modality. Thus Yalcin (2007, 2011) observes that the following sentence (5a) sounds terrible, and cannot even be supposed to be true:

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

It may be tempting to treat the infelicity of (5a) as a matter of pragmatics, analogous to Moorean paradoxical constructions of the form “ ϕ but I do not believe that ϕ .” However, (5a) differs from Moore’s paradoxical constructions in that the latter *can* be supposed to be true: it is unproblematic for me to suppose that it is raining and I do not believe that it is raining. Hence the defect of (5a) must go deeper. This is a puzzle for the standard analysis of *might* as an existential quantifier ranging over a set of possible worlds compatible with what is known since there is nothing problematic about supposing the existence of unknown truths.

3.1.2 Informativity and Ignorance

A fact that is less commonly noticed is that *might*-statements are sometimes informative. Consider the following example:

- (6) 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 is 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. This suggests 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.

A related observation is that agents are sometimes agnostic about a certain epistemic possibility even though the prejacent is a question under discussion. Consider the following case. A cancer test is run for John. If the result is negative, cancer is ruled out; if the result is positive, John might or might not have cancer: further tests are necessary. As DeRose (1991) points out, the following exchange sounds perfectly fine if neither Mary nor Alex are informed about the test result:

- (7) 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 John's having cancer is compatible with what they respectively know.

3.1.3 Must

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 would 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:

- (8) 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?

Even though it is compatible with what Mary knows that Colin is not in Chicago, it is natural for her not to reject Alex's utterance and instead to integrate the encoded information into her epistemic state.

In contrast, consider a case in which an assessor believes 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:

- (9) 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 that Mary's reaction to Alex's judgment 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 (9) may continue as follows:

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

The observation then is that agents do not always reject a judgment of epistemic necessity in case they do not rule out the falsity of the prejacent. But there are also cases in which it is natural to respond to a *must*-statement by pointing to an open possibility that falsifies the prejacent. A good semantics and pragmatics for epistemic modals should be able to account for these data.

3.2 Explanations

3.2.1 Epistemic Contradiction

Yalcin’s observation 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.²³ So an attempt at updating one’s information state with (5a) 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 (5a) 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 (5a) is pragmatically infelicitous.²⁴

Yalcin acknowledges the explanatory potential of the dynamic perspective on epistemic modality but bases his theory on a static conception of logical consequence. Like my proposal, his theory treats “It is raining and it might not be raining” as a contradiction, but unlike mine it treats the result of reversing the order of conjuncts in the same vein:

(5) c. # It might not be raining and it is raining.

Thus on Yalcin’s account one would expect that epistemic contradictions are insensitive to the order of conjuncts. This is anything but uncontroversial, however, and the fact that (5c) is dynamically consistent readily explains why some maintain that reversing the order of conjuncts improves the odd-sounding (5a).²⁵ Yet (5c) still strikes many speakers as odd and it is legitimate to ask—as an anonymous reviewer does—how the dynamic framework can account for this fact. The answer is that while (5c) is consistent from a dynamic perspective, it is *incoherent* in the sense that no single information state other than the absurd one can support it: there are no ϕ and $\Sigma \neq \Sigma_{\emptyset}$ such that $\Sigma \models \diamond\neg\phi$ and $\Sigma \models \phi$.²⁶ Insofar as the point of a supposition is to adopt a state of information that supports a certain hypothesis, we then expect that supposing “It might not be raining and it is raining” is just as odd as supposing “It is raining and it might not be raining.” Plain assertions of (5c), thus, are predicted to be infelicitous, and we can say all this without neglecting the subtle differences between (5a) and (5c) that follow from the noncommutativity of dynamic conjunction.

²³But 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.”

²⁴Yalcin also observes that (5a) cannot occur felicitously in the antecedent of a conditional. I have not presented a full theory of conditionals here but I will say that much: evaluating a conditional proceeds by evaluating its consequent under the supposition of its antecedent, and so what I have said about supposing (5a) also makes sense of Yalcin’s observation about conditional antecedents.

²⁵See, for instance, [Sorensen \(2009\)](#) and [Dorr and Hawthorne \(2012\)](#) for discussion.

²⁶On the notion of coherence, see [Groenendijk et al. \(1996\)](#) and, more recently, [Gillies \(2004\)](#), who prefers the label “cohesiveness.”

3.2.2 Informativity and Ignorance

The previous subsection illustrated how the semantics for *might* interacts with the dynamic conception of logical consequence to account for Yalcin’s observation about epistemic contradictions. The distinction between plain and live possibilities—another important feature of the framework developed here—is not crucial for the logical connection between factual judgments and present tense claims of epistemic modality, but it plays a key role in explaining the observations in §3.1.2. Whenever a possibility is neither ruled out nor taken seriously by an inquiring agent S , we predict that S is agnostic about the corresponding judgment of epistemic possibility: S ’s information state neither supports the possibility judgment nor its negation. The existence of information states that admit epistemically modalized sentences guarantees that updates with such sentences can be informative, as observed in (6). Let’s first get clear on what it means for an update to be informative.

Definition 8 (Informativity) Consider any $\phi \in \mathcal{L}$ and $\Sigma \in I$:

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

This conception of informativity translates the insights from [Stalnaker \(1978\)](#) into the setup developed here: adding the information encoded in ϕ to Σ should eliminate some but not all elements of Σ . The following fact is obvious:

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 ‘ $\diamond\phi$ ’-admitting information state, ‘ $\diamond\phi$ ’ will be informative with respect to that information state.

The formal result that updates with *might* sometimes have non-trivial effects on an information state may not be entirely satisfactory: if updating with *might* is not guaranteed to eliminate certain ways the world could be from an information state, in what sense do the observed non-trivial effects count as information change? Both factual statements and statements of epistemic modality refine an agent’s conception of the set of possible worlds that matter for inquiry; they only do so in distinct ways: the former eliminate certain possibilities in a way familiar from Stalnaker’s discussion of assertion, the latter introduce hitherto ignored possibilities as live possibilities. Since both update effects matter for how an inquiring agent proceeds in theoretical and practical inquiry, there is no reason to think that one of them is more fundamental than the other. Quite to the contrary, we may observe that while an agent’s live possibilities do not represent reality as being one way or another, they render some ways the world could be more plausible than others. More precisely, observe that an agent’s state of information Σ does not only select a subset of W as the set of possible worlds compatible with the information that Σ carries, but also imposes a ranking on W that reflects the agent’s conception of which possibilities to take seriously in inquiry.

Definition 9 (Rankings) \leqslant is a function mapping an information state Σ to a relation on W so that $w \leqslant_{\Sigma} w'$ iff for all $\phi \in \mathcal{L}_0$ such that $\Sigma \models \diamond\phi$ and $\Sigma \not\models \diamond\neg\phi$: if ϕ is true at w' , then ϕ is true at w . \leqslant_{Σ} is the *ranking imposed by Σ on W* .

Ideally, the actual world w_a is not only compatible with the information Σ carries but also minimal in \leqslant_{Σ} (most plausible in light of the possibilities that Σ recognizes) in the sense that for all $w' \in W$, if $w' \leqslant_{\Sigma} w_a$ then $w_a \leqslant_{\Sigma} w'$. Regardless of whether this is the case, it is obvious that *might*-statements are alike to statements of fact in that they potentially affect the ranking an information state imposes on W . I submit that insofar as we have any intuition that *might* is informative, it is adequately captured by the dynamic framework developed here.²⁷

The dynamic framework developed here also leaves room for Alex's response in (7) that he does not know whether John might have cancer: he is not in a position to rule out the possibility that John has cancer but refrains from committing to this being a live possibility until he knows the test results. There is, however, a non-trivial question about Alex's state of mind here: the possibility of John's having cancer is contextually salient and not implausible in light of the available evidence to Alex, and so his agnostic response is a bit of a mystery. The common orthodox reaction is that Alex's use of *might* refers to some third party knowledge. When Alex says that he does not know whether John might have cancer, he claims to be ignorant as to whether John's having cancer is ruled out by what is known to a group of people including, for instance, John and his doctors. The orthodox stance, however, is not without alternatives, and in any case it fails to tell us everything we want to know about the case under consideration.

That third party information plays *some* crucial role in Alex's agnostic response is undeniable, especially since agnostic responses to questions about *might* are marginal at best in case no such third party information is salient, as an anonymous reviewer observes. However, merely pointing to the possibility of non-solipsistic uses of epistemic *might* does not answer the basic question *why* Alex reacts the way he does instead of giving an affirmative response like the following:

- (11) Mary: I've heard that John is sick. Might it be cancer?
 Alex: It might be cancer but we'll know more tomorrow once the test results are in.

How can we account for Alex's reaction in DeRose's example? As follows: earlier I said that which possibilities an agent takes seriously in discourse and reasoning

²⁷Notice that in accounting for the possibility of informative updates with *might*, the dynamic framework developed here substantially improves on Veltman's proposal. The model of an information state that figures in Veltman's update semantics predicts that an update with *might* idles whenever the prejacent is compatible with the information, leaving no room for informativity.

may depend on practical concerns, and Yalcin (2007) observes that we sometimes aim at bringing our presuppositions in line with what is known by the experts or some other carrier of information. In such cases it makes perfect sense to take an agnostic stance toward a question about *might*. A plausible suggestion then is that when Alex says that he does not know whether John might have cancer, he does so because he aims at keeping his live presuppositions in line with what is compatible with the test results: Alex ‘defers’ the question whether John having cancer is a possibility to be taken seriously to a source of information that he is, as a matter of fact, in suspense about. It does not follow from this that he takes an agnostic stance on a judgment of epistemic modality with non-solipsistic truth-conditions. Alex, it is plausible to say, evaluates a *might*-statement in the familiar fashion against his own perspective, the only wrinkle being that his inquisitive intentions put him into a conditional state of mind: cancer is a live possibility if, and only if, the test comes back positive. So adopting the belief that John might have cancer would commit Alex to the belief that the test is positive, which—given the correctness criteria outlined in §2.4—is just to say that whether it is correct that John might have cancer depends on whether it is true that the test is positive. Since Alex does not know the test results, he does not know whether John having cancer is a possibility to be taken seriously, and thus chooses an agnostic response to Mary’s question.

3.2.3 Must

The contrast between plain and live possibilities also accounts for our earlier observations about epistemic *must*. Notice that whenever Σ admits $\ulcorner \diamond \phi \urcorner$ it also admits $\ulcorner \Box \neg \phi \urcorner$:

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

But when Σ supports $\ulcorner \diamond \phi \urcorner$, $\Sigma[\Box \neg \phi] = \emptyset$ since *might* and *must* are duals.

Thus we expect that assessments of *must* vary depending on whether the falsity of the prejacent is a plain or a live possibility. Going back to our observations in §3.1.3, the explanation for Mary’s reactions seems to be that in (8) it is merely compatible with what she knows that Colin is not in Chicago, i.e. Mary’s information state admits Alex’s claim that Colin must be in Chicago. The situation is different in (9) where it is a live possibility for Mary that Colin forgot to turn off the lights when he left the house: here Mary’s information state rejects the claim that Colin must be at home. Accordingly, our theory predicts uptake in (8) but rejection in (9), as desired.

On the present proposal, the pragmatics of assessment for $\ulcorner \Box \phi \urcorner$ differs from the one for ϕ in that the former, but not the latter, is subject to rejection in case $\ulcorner \neg \phi \urcorner$ is a live possibility. To see why this is right, consider the following variant of the discourse in front of Colin’s house:

(12) Alex: Colin is at home.

Mary: # No, he might be out—maybe he just forgot to turn off the lights.

Mary’s response in (12) is marginal, and this is just what we expect: that Colin might have forgotten to turn off the lights means that he does not *have* to be at home—Mary is in a position to reject Alex’s judgment of epistemic necessity in (9)—but does not license Mary’s rejection of Alex’s statement of fact.²⁸

I am about to conclude that the distinction between plain and live possibilities plays a central role in every successful pragmatic story about epistemic *must*. Before that, let me clarify that nothing said so far excludes the possibility that agents sometimes go along with a judgment of epistemic necessity even though the falsity of its prejacent is a live possibility—revision to accommodate a conflicting bit of information is always an option. But there is more to be said here: earlier in §2.3 I justified the suggestion that, all things being equal, rejection is preferred to revision since the latter in general requires weakening of one’s information state. This is right and also gets the data about disagreement straight: if Mary believes that it is not raining, she is likely to reject an utterance that it is (or might be) raining, and this is just what we expect if revision requires weakening and is thus avoided unless there is compelling reason to proceed otherwise. Revision with the aim of accommodating epistemic *must*, however, is special in a way that leaves room for this phenomenon to be a quite common phenomenon. Let me explain.

The crucial observation is that whenever $\Sigma_S \models \diamond\neg\phi$, Σ_S must be revised to allow for consistent update with $\lceil \Box\phi \rceil$, but the revision procedure differs from the more familiar ones in that it does not require weakening: all the agent needs to do to arrive at an information state consistent with $\lceil \Box\phi \rceil$ is to update Σ_S with ϕ . When it comes to judgments of epistemic necessity, then, revision does not have its usual costs, and so it is compatible with both the letter as well as the spirit of the pragmatic story told here that accommodation of epistemic *must* is a quite common phenomenon. In fact, the previous observation suggests that agents are willing to accommodate a judgment of epistemic necessity to the extent that they are willing to update their information state with its prejacent. The upshot: while the pragmatics of assessment for *must*-statements differs in non-trivial ways from the one for utterances of their bare prejacents, we predict—correctly, I think—that in many cases an agent’s reaction to the former is a function of how he or she would react to the latter.²⁹

We have seen that the dynamic framework developed so far accounts for a range of problems that go beyond issues of modal disagreement. We have also seen that the heavy duty components of the technical apparatus—in particular, the distinction between plain and live possibilities—possess an intuitive

²⁸The fact that Mary is not in a position to *reject* Alex’s judgment in the scenario under consideration is, of course, compatible with the observation that she might point to the live possibility to achieve a different discourse goal. For instance, Mary is free to challenge Alex to explain his grounds for excluding the possibility that Colin is out and simply forgot to turn off the lights.

²⁹Thanks to an anonymous reviewer for pressing me on the pragmatic issues about *must* discussed in this and the preceding paragraph.

foundation in everyday discourse and reasoning. Since the resulting semantics avoids the most common difficulties faced by non-truth-conditional theories of meaning—it does not have a Frege-Geach problem and makes sense of correct or incorrect epistemic modal beliefs—there is good reason to think that we have found the right semantics for *might* and *must*. But there are some more complex data that require an extension of the existing framework with a basic semantics for tense. The details are elaborated in the next section.

4 Tense

This section extends the basic framework with a simple semantics for tense. The resulting framework explains complicated data about the ways speakers may assess previously made claims of epistemic modality in the light of new information.

4.1 The Observations

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

- (13) 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 Stechow and Gillies \(2008a\)](#) 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:

- (14) 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.³⁰ And they might have been.

Here Alex sticks to his guns, and the question of why this is a legitimate move is as important as the question of why it is also sometimes natural to withdraw a commitment to *might* 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

³⁰Speakers' 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 Stechow and Gillies, avoid the backshift. Such minor disputes need not be resolved here.

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. Thus consider:

- (15) 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 (15a) says is that the keys being in the car was once a possibility, even though it is not any longer. What the first disjunct in (15b) says is that it is now a possibility that the keys were in the car. An alternative way of putting (15b) is the following:

- (15) 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 (15a) with the following strange example:

- (15) 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 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 a possibility. This is just the moral drawn from the examples in (13) and (14). Second, saying that the truth of the prejacent was a possibility is importantly different from saying that it is a 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 (15a) and (15d).

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

4.2 Tense

The exercise is to provide a semantics for a \mathcal{L}^+ , which is just the result of extending our base language \mathcal{L} with the temporal operators for *past* (P) and *future* (F). Everything that has been said about formulas of \mathcal{L} remains valid, so all we have to deal with are the newly introduced operators.

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. 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. Possible worlds now assign truth-values to elements of \mathcal{A} relative to points of time. An agent's conception of what matters in inquiry extends over time and is thus modeled as a function from points of time to sets of such possible worlds. Once again, we rely on supervaluationism to capture the fact that an agent's state of information does not completely specify, for each point in time, which possible worlds matter for inquiry.

Definition 10 (Tense) Let $Temp$ be any nonempty set, $<$ be a transitive relation on $Temp$ that is also irreflexive and linear. v is a *temporally extended possible world* iff $v: Temp \rightarrow (\mathcal{A} \rightarrow \{0, 1\})$. V is the set of such v 's. T is a *temporal information state* iff T is a set of functions $\tau: Temp \rightarrow (V \rightarrow \{0, 1\})$ such that for all $\tau \in T, t \in Temp: \tau(t) \neq \emptyset$, i.e. T is a set of functions from temporal indices to non-empty sets of temporally extended possible worlds. I^+ is the set of such T 's. The *absurd* temporal information state T_\emptyset is identical 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. Elements of temporal information states change over time and are modeled as functions from temporal indices to sets of (temporally extended) possible worlds. So the natural thing to do is to understand updates on such elements as resetting the values assigned to their temporal indices. For this proposal to make sense, one needs to say how the values assigned to temporal indices are changed through updates and what it means to reset the values assigned to temporal indices. Let me address these issues in turn.

Definition 11 (Updates on Elements of Temporal Information States) Consider arbitrary $\tau: Temp \rightarrow (V \rightarrow \{0, 1\})$, $t \in Temp$, $p \in \mathcal{A}$ and $\phi, \psi \in \mathcal{L}^+$. An update on an element of a temporal information state is a function $\uparrow: \wp(V) \rightarrow \wp(V)$ 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) $\tau(t) \uparrow P\phi = \{v \in \tau(t): \exists t' < t: v \in \tau(t') \uparrow \phi\}$
- (6) $\tau(t) \uparrow F\phi = \{v \in \tau(t): \exists t' > t: v \in \tau(t') \uparrow \phi\}$

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 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 τ 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)-(4) should be obvious from what has been said earlier. Clauses (5) and (6) 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, an update of $\tau(t)$ with $\lceil P\phi \rceil$ leaves those worlds in $\tau(t)$ that 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 semantics of the future tense operator should be obvious.

Elements of an information state change over time, and we agreed that since they are functions from points of time to sets of possible worlds, this should be understood as a resetting of such functions.

Definition 12 (Reset) Consider any $\tau: Temp \rightarrow (V \rightarrow \{0,1\})$, $t \in Temp$, $\phi \in \mathcal{L}^+$:

1. $\tau^{\tau(t)\uparrow\phi}$ is the function from temporal indices to sets of possible worlds that is just like τ except that for all $t' \prec t$, $\tau^{\tau(t)\uparrow\phi}(t') = \tau(t) \uparrow \phi$
2. $\tau \uparrow^t \phi = \tau^{\tau(t)\uparrow\phi}$

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 13 (Updates on Temporal Information States) Consider any $T \in I^+$, $t \in Temp$, $\phi \in \mathcal{L}^+$. An update on a temporal information state at t is a function $[\cdot]^t: I^+ \rightarrow I^+$ 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 elements together, leaving out those assigning to t the empty set. This yields the output temporal information state.

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

Definition 14 (Temporal Settledness and 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. 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 toward 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 that one can talk about a possible past or a past possibility. Remember our earlier examples:

- (15)
- 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.

(15a) is about a past possibility, and thus we expect the epistemically modalized sentence to be of the form ' $P\Diamond\phi$ '. On the other hand, the first disjuncts in (15b) and (15c) are concerned with a possible past, and that means that they are of the logical form ' $\Diamond P\phi$ '. 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 (15d) sounds terrible (repeated):

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

And the infelicity of (15d) 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. Accordingly, updating with the second conjunct results in the absurd state. So updating with (15d) 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.

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 focused on the commitments that *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 judgments of

epistemic modality after learning something new. The question then 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 $T \in I^+$, $t \in Temp$: 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 that 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 .³¹

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 an agent S with temporal information state T_S such that $T_S \models_t \phi$. Then S will by default assess his or her commitment to ϕ at t as follows after updating with ψ :

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

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 he or she was wrong in believing that ϕ .

We now have everything we need to explain the complicated data about how agents assess their judgments of epistemic possibility after learning that the prejaacent in fact is false, the crucial point being that updating temporal information states fails to preserve previous commitments in just the right way. The puzzle, remember, 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

³¹Preservation is a major issue in the literature on belief change. Alchourrón et al. (1985) and Gärdenfors (1988) predict that unless one already believes ‘ $\neg\phi$ ’ in a prior state, revising that state with ϕ results in a posterior state that carries at least as many commitments as the prior state. It is a well-worn story that so much doxastic conservatism leads to trouble when one also has beliefs about what might or might not be the case: one may believe that ‘ $\diamond\neg\phi$ ’ without believing that ‘ $\neg\phi$ ’, and still a commitment to ‘ $\diamond\neg\phi$ ’ needs to go once one learns that ϕ is the case. See Levi (1988), Fuhrmann (1989), and Rott (1989) for seminal discussion. As Gillies (2006) has shown, dynamic semantics gives us independent reason to reject the idea that updating is always preservative updating.

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:

- (16) 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 $\neg p$ preserves $P\Diamond p$ but not $\Diamond p$, and this is not a coincidence: for all $\phi \in \mathcal{L}$, updating preserves $\ulcorner P\phi \urcorner$ —plain past possibilities never get eliminated—but updating with $\ulcorner \neg\phi \urcorner$ does *not* preserve $\ulcorner \Diamond\phi \urcorner$ unless ϕ is a tautology. Specifically, 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 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 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, (16a) as well as (16b) 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 possibility, and this is how it should be. Current commitments to such 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 *might* after learning something new.

5 Comparisons

The goal of this paper has been to look for an alternative to the orthodox view that epistemically modalized 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:

Definition 16 (Contextualism) Consider a 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) $\llbracket B \rrbracket^{c,i,a} = \{v : v \text{ is compatible with what } j_a \text{ knows at } t_a \text{ in } w_a\}$

Contextualists and relativists agree that modals are quantifiers over a suitably restricted set of possible worlds, which is often called the *modal base* and denoted by B in the above definitions. Epistemic modals quantify over the set of possible worlds that are compatible with a relevant body of knowledge. A judgment of the form ‘ $\Diamond\phi$ ’ is true just in case ϕ is true at some world in the modal base. A judgment of the form ‘ $\Box\phi$ ’ is true just in case ϕ is true at every world in the modal base. The dispute between contextualists and relativists concerns the way in which the modal base is determined. As one can see in the above definitions, contextualists maintain that the modal base is completely determined by the index of evaluation and context of utterance. In contrast, relativists hold that the modal base may *also* depend on features of the situation in which a judgment of epistemic modality is assessed. It follows that judgments of epistemic modality may vary in truth-value across points of assessment, even if the context of utterance and index of evaluation remain fixed. Relativists commonly insist that so much flexibility is needed to account for the complicated data about modal disagreement. My claim is that both contextualism and relativism are equally problematic views.³²

The key reason to be skeptical of the orthodox view is the intuition that variation in what is known should have a pragmatic but no semantic effect on discourses involving *might* and *must*. According to my theory, variation in what is known only matters for how *might*- and *must*-statements are assessed. As a result, it 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.

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 seems right since speakers frequently make claims of epistemic possibility in case the prejacent is compatible with what they know. But the assumption that speakers sometimes use epistemic modals with the intention of reporting on what is or is not

³²See fn. 1 and fn. 2 for proponents of relativism and contextualism, respectively. 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. We may also distinguish contextualism and relativism as defined in the main text from so-called “non-indexical” versions of contextualism and relativism (see MacFarlane (2009) for discussion). For current purposes, it is harmless to set these complications aside.

compatible with what someone else knows is less obvious. Without this assumption, however, orthodoxy has a hard time explaining how such utterances of epistemically modalized sentences behave in discourse.

The point becomes immediately clear when we consider what contextualists tend to say about the original dispute between Alex and Mary (repeated):

- (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.

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. One response on behalf of contextualism starts with the idea that epistemic modals are by design 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. This discourse effect legitimates Mary's denial of Alex's utterance.³³

The contextualist's response that multiple propositions are floated into the discourse that may be rejected by the audience is dubious for at least two reasons. First, there is no compelling evidence for the underlying hypothesis that epistemically modalized sentences are ambiguous. Common ambiguity tests fail to deliver the desired support.³⁴ 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 (17):

- (17) # Colin might be in Chicago and he can't be in Chicago.

(17) strikes the ear as a glaring contradiction.³⁵ This would be rather surprising if the epistemically modalized sentence had the multiple readings that some contextualists suggest. Of course, ambiguity tests have to be handled with care and in any case the contextualist is free to insist that epistemic modals are ambiguous in a very special way that is not detectable in the usual manner. None of these reservations, however, affects the basic observation that the ambiguity thesis lacks independent support, putting the contextualist proposal on rather shaky grounds. Second, even if there were an ambiguity, it would be hard to find a reasonable semantic or pragmatic principle predicting 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 in discourse would put into play all available readings, which then become subject to legitimate rejection by the hearer. In fact, when confronted with an

³³See von Stechow and Gillies (2011).

³⁴See Zwicky and Sadock (1975) for a discussion of a whole range of such tests.

³⁵If you think that this is just a pragmatic issue, notice that (17) remains more than bad even if it is just *supposed*.

ambiguous expression, cooperative hearers assume the most reasonable reading or find out what the speaker really intended to say. And even if we abstract away from all of these difficulties, it remains a strong point in favor of the view developed here that it does not treat epistemically modalized sentences as ambiguous.

A similar worry applies to relativism, though the issue is a bit more subtle.³⁶ To make sense of simple disputes like (1), relativists need not depart from the fairly straightforward view that the assessor's knowledge always 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 evaluated when he is the assessor, and it is Mary's when she is the judge. Thus relativists 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 becomes more complicated once we move on to claims of epistemic necessity. We know of 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 (8) is false from Mary's point of assessment, leaving it unexplained why she does not reject Alex's utterance (repeated):³⁷

- (8) 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?

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 it is now unclear whether relativism can claim any advantage over contextualism: very much as before, we lack any plausible principle to explain 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 worrisome because a claim of epistemic necessity is not always assessed against the knowledge distributed among the assessor and the speaker, as the dialogue in (9) shows (repeated):

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

³⁶The challenge I am about to articulate is only one among many that relativism faces. See, e.g., von Stechow (2008a), Wright (2007), and Yalcin and Knobe (2010) for critical discussion.

³⁷At 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 solves some puzzling cases involving epistemic modality without diverging too much from classical truth-conditional semantics.

Here it seems that Mary rejects Alex’s claim because she thinks that Alex has overlooked a possibility (regardless of whether he did so in fact). So Mary’s point of assessment selects her knowledge as the standard of evaluation, not the one distributed among her and Alex. When it comes to the more complex cases, relativists thus cannot offer a non-arbitrary way of selecting the modal base over which *might* and *must* ranges, at least as long as this way of selecting the modal base is supposed to yield truth-values that do justice to how the relevant claim of epistemic modality is assessed.³⁸

The dynamic framework developed here does not deny that we can find reasonable correctness criteria for judgments of epistemic modality—as we have seen, such criteria are identifiable even in case these judgments lack propositional content. The point is that such correctness criteria do not serve as a reliable guide to the semantics and pragmatics of epistemic modals. A speaker’s judgment that so-and-so might be the case may be correct—the speaker may be responding correctly to his or her evidence—and still be rejected by a better informed assessor. And this is unsurprising if we think of meaning as context change potential and of assertions as made with the intention that hearers update their information state with what is said. A hearer’s reaction to a judgment involving *might* or *must* is then expected not to be guided by considerations of correctness but rather by the judgment’s update effect on the hearer’s information state. We have seen that this approach easily accounts for the observation that ordinary speakers assess judgments involving *might* by testing them against their own perspective and immediately extends to cover related observations about *must*.

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, since 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 plain past possibilities untouched, but may eliminate current 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 point to a commitment (that the keys might have been in the car) that remains untouched by the new information. To my knowledge, no other account is in a position to offer an equally straightforward explanation of the complicated data about post-update assessment.³⁹ Since my story eas-

³⁸Could the relativist appeal to accommodation to explain Mary’s reaction in (8), as an anonymous reviewer suggests? Perhaps, but now relativist truth-conditions no longer do the heavy lifting in explaining how speaker assess judgments of epistemic necessity in discourse—instead, we look at the effects that such judgments have on a recipient’s state of information, just as proposed by the dynamic story that has been told here.

³⁹Stephenson (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 that is so central to my theory.

ily handles other embeddings as well—contrary to what is so often said about non-truth-conditional approaches—I conclude that it is superior to the orthodox alternatives currently on the market.

Acknowledgments For comments and discussion, I would like to thank Nicholas Asher, David Beaver, Fabrizio Cariani, Josh Dever, Hans Kamp, Mark Sainsbury, Seth Yalcin, two anonymous *Philosophical Review* referees, and audiences at New York University, Northwestern University, the University of Texas at Austin, and the 2010 Eastern APA.

References

- Alchourrón, Carlos E., Peter Gärdenfors, and David Makinson. 1985. “On the Logic of Theory Change: Partial Meet Contraction and Revision Functions.” *Journal of Symbolic Logic* 50(2): 510–530.
- Asher, Nicholas. 2007. “Dynamic Discourse Semantics for Embedded Speech Acts.” In *John Searle’s Philosophy of Language*, ed. Savas L. Tsohatzidis, 211–244. New York: Cambridge University Press.
- Asher, Nicholas, and Alex Lascarides. 2009. “The Interpretation of Questions in Dialogue.” In *Proceedings of Sinn und Bedeutung 13*, ed. Arndt Riester and Torgrim Solstad, 17–30. Stuttgart: OPUS.
- Asher, Nicholas, and Eric McCready. 2007. “Were, Would, Might, and a Compositional Account of Counterfactuals.” *Journal of Semantics* 24(2): 93–129.
- Beaver, David I. 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford: CSLI Publications.
- DeRose, Keith. 1991. “Epistemic Possibilities.” *Philosophical Review* 100(4): 581–605.
- Dever, Josh. 2006. “Living the Life Aquatic: Does Presupposition Accommodation Mandate Dynamic Semantics?” <https://webpace.utexas.edu/deverj/personal/papers/aquatic.pdf>. Manuscript, University of Texas at Austin.
- Does, Jaap van der, Willem Groeneveld, and Frank Veltman. 1997. “An Update on Might.” *Journal of Logic, Language and Information* 6(4): 361–380.
- Dorr, Cian, and John Hawthorne. 2012. “Embedding Epistemic Modals.” <http://users.ox.ac.uk/~sfop0257/papers/Yalcin.pdf>. Manuscript, Oxford.
- Dowell, Janice L. 2011. “A Flexible Contextualist Account of Epistemic Modals.” *Philosophers’ Imprint* 11(14).
- Dretske, Fred. 1970. “Epistemic Operators.” *Journal of Philosophy* 67(24): 1007–1023.

- . 1981. “The Pragmatic Dimension of Knowledge.” *Philosophical Studies* 40(3): 363–378.
- Egan, Andy. 2007. “Epistemic Modals, Relativism, and Assertion.” *Philosophical Studies* 133(1): 1–22.
- Egan, Andy, John Hawthorne, and Brian Weatherson. 2005. “Epistemic Modals in Context.” In *Contextualism in Philosophy*, ed. Gerhard Preyer and Georg Peter, 131–169. New York: Oxford University Press.
- Fine, Kit. 1975. “Vagueness, Truth and Logic.” *Synthese* 30(3-4): 265–300.
- von Fintel, Kai, and Anthony S. Gillies. 2008a. “CIA Leaks.” *Philosophical Review* 117(1): 77–98.
- . 2008b. “An Opinionated Guide to Epistemic Modality.” In *Oxford Studies in Epistemology 2*, ed. Tamar Gendler and John Hawthorne, 32–62. New York: Oxford University Press.
- . 2010. “Must ...Stay ...Strong.” *Natural Language Semantics* 18(4): 351–383.
- . 2011. “‘Might’ Made Right.” In *Epistemic Modality*, ed. Andy Egan and Brian Weatherson, 108–130. New York: Oxford University Press.
- Frege, Gottlob. 1879. *Begriffsschrift: Eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Halle: Louis Nebert.
- Fuhrmann, André. 1989. “Reflective Modalities and Theory Change.” *Synthese* 81(1): 115–134.
- Gärdenfors, Peter. 1988. *Knowledge in Flux*. Cambridge, MA: MIT Press.
- Gillies, Anthony S. 2001. “A New Solution to Moore’s Paradox.” *Philosophical Studies* 105(3): 237–250.
- . 2004. “Epistemic Conditionals and Conditional Epistemics.” *Noûs* 38(4): 585–616.
- . 2006. “What Might Be the Case after a Change in View.” *Journal of Philosophical Logic* 35(2): 117–145.
- . 2009. “On Truth-Conditions for *If* (but Not Quite Only *If*).” *Philosophical Review* 118(3): 325–349.
- . 2010. “Iffiness.” *Semantics and Pragmatics* 3(4): 1–42.
- Groenendijk, Jeroen. 2009. “Inquisitive Semantics: Two Possibilities for Disjunction.” In *Lecture Notes in Artificial Intelligence: Logic, Language, and Computation, TbiLLC 2007*, ed. Peter Bosch, David Gabelaia, and Jérôme Lang, 80–94. Berlin: Springer-Verlag.

- Groenendijk, Jeroen, and Martin Stokhof. 1991. "Dynamic Predicate Logic." *Linguistics and Philosophy* 14(1): 39–100.
- Groenendijk, Jeroen, Martin Stokhof, and Frank Veltman. 1996. "Coreference and Modality." In *The Handbook of Contemporary Semantic Theory*, ed. Shalom Lappin, 179–214. Cambridge, MA: Blackwell.
- Hacking, Ian. 1967. "Possibility." *Philosophical Review* 76(2): 143–168.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- . 1992. "Presupposition Projection and the Semantics of Attitude Verbs." *Journal of Semantics* 9(3): 183–221.
- Kamp, Hans. 1975. "Two Theories about Adjectives." In *Formal Semantics of Natural Language*, ed. Edward L. Keenan, 123–155. Cambridge: Cambridge University Press.
- . 1981. "A Theory of Truth and Representation." In *Formal Methods in the Study of Language, Part I*, ed. Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, 277–320. Amsterdam: Mathematisch Centrum.
- Kamp, Hans, Josef van Genabith, and Uwe Reyle. 2011. "Discourse Representation Theory." In *Handbook of Philosophical Logic*, 2nd ed., vol. 15, ed. Dov M. Gabbay and Franz Guenther, 125–394. Dordrecht: Springer.
- Kamp, Hans, and Uwe Reyle. 1993. *From Discourse to the Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer Academic Press.
- Kant, Immanuel. 1781. *Critik der Reinen Vernunft*. Riga: Johann Friedrich Hartknoch.
- Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1(3): 337–355.
- . 1981. "The Notional Category of Modality." In *Words, Worlds, and Contexts. New Approaches to Word Semantics*, ed. Hans-Jürgen Eikmeyer and Hannes Rieser, 38–74. Berlin: Walter de Gruyter & Co.
- . 1991. "Modality." In *Semantics: An International Handbook of Contemporary Research*, ed. Armin von Stechow and Dieter Wunderlich, 639–650. Berlin: Walter de Gruyter & Co.
- . 2012. *Modals and Conditionals*. New York: Oxford University Press.
- Levi, Isaac. 1988. "Iteration of Conditionals and the Ramsey Test." *Synthese* 76(1): 49–81.

- . 1996. *For the Sake of the Argument*. New York: Cambridge University Press.
- Lewis, David K. 1996. “Elusive Knowledge.” *Australasian Journal of Philosophy* 74(4): 549–567.
- MacFarlane, John. 2009. “Nonindexical Contextualism.” *Synthese* 166(2): 231–250.
- . 2011. “Epistemic Modals Are Assessment-Sensitive.” In *Epistemic Modality*, ed. Andy Egan and Brian Weatherson, 144–178. New York: Oxford University Press.
- Price, Huw. 1983. “Does ‘Probably’ Modify Sense?” *Australasian Journal of Philosophy* 61(4): 396–408.
- Roberts, Craige. 1987. *Modal Subordination, Anaphora, and Distributivity*. Ph.D. thesis, University of Massachusetts, Amherst.
- . 1989. “Modal Subordination and Pronominal Anaphora in Discourse.” *Linguistics and Philosophy* 12(6): 683–721.
- Rodrigues, Odinaldo, Dov M. Gabbay, and Alessandra Russo. 2011. “Belief Revision.” In *Handbook of Philosophical Logic*, 2nd ed., vol. 16, ed. Dov M. Gabbay and Franz Guenther, 1–114. Dordrecht: Springer.
- Rothschild, Daniel. 2012. “Expressing Credences.” *Proceedings of the Aristotelian Society* 112(1pt1): 99–114.
- Rott, Hans. 1989. “Conditionals and Theory Change: Revisions, Expansions, and Additions.” *Synthese* 81(1): 91–113.
- Sorensen, Roy. 2009. “Meta-agnosticism: Higher Order Epistemic Possibility.” *Mind* 118(471): 777–784.
- Stalnaker, Robert C. 1970. “Pragmatics.” *Synthese* 22(1-2): 272–289.
- . 1978. “Assertion.” In *Syntax and Semantics 9*, ed. Peter Cole, 315–332. New York: New York Academic Press.
- . 1981. “A Defense of Conditional Excluded Middle.” In *Ifs: Conditionals, Belief, Decision, Chance, and Time*, ed. William L. Harper, Robert C. Stalnaker, and Glenn Pearce, 87–104. Dordrecht: Reidel.
- . 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stephenson, Tamina. 2007. *Towards a Theory of Subjective Meaning*. Ph.D. thesis, MIT.
- . 2008. “Judge Dependence, Epistemic Modals, and Predicates of Personal Taste.” *Linguistics and Philosophy* 30(4): 487–525.

- Teller, Paul. 1972. “Epistemic Possibility.” *Philosophia* 2: 303–320.
- Thomason, Richmond H. 1970. “Indeterminist Time and Truth-Value Gaps.” *Theoria* 36(3): 264–281.
- van Fraassen, Bas. 1966. “Singular Terms, Truth-value Gaps, and Free Logic.” *Journal of Philosophy* 63(17): 481–495.
- Veltman, Frank. 1985. *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam.
- . 1996. “Defaults in Update Semantics.” *Journal of Philosophical Logic* 25(3): 221–261.
- Wright, Crispin. 2007. “New Age Relativism and Epistemic Possibility: The Question of Evidence.” *Philosophical Issues* 17, The Metaphysics of Epistemology: 262–283.
- Yablo, Stephen. 2011. “A Problem about Permission and Possibility.” In *Epistemic Modality*, ed. Andy Egan and Brian Weatherson, 270–294. New York: Oxford University Press.
- Yalcin, Seth. 2007. “Epistemic Modals.” *Mind* 116(464): 983–1026.
- . 2009. “More on Epistemic Modals.” *Mind* 118(471): 785–793.
- . 2011. “Nonfactualism about Epistemic Modals.” In *Epistemic Modality*, ed. Andy Egan and Brian Weatherson, 295–332. New York: Oxford University Press.
- . 2012. “Context Probabilism.” *Lecture Notes in Computer Science* 7218: 12–21.
- Yalcin, Seth, and Joshua Knobe. 2010. “*Fat Tony Might Be Dead*: An Experimental Note on Epistemic Modals.” <http://www.semanticsarchive.net/Archive/TdiZjA3N/yalcin.knobe.2010.eavesdropping.pdf>. Manuscript, Berkeley/Yale.
- Zeevat, Henk. 1992. “Presupposition and Accommodation in Update Semantics.” *Journal of Semantics* 9(4): 379–412.
- Zwicky, Arnold M., and Jerrold M. Sadock. 1975. “Ambiguity Tests and How to Fail Them.” In *Syntax and Semantics 4*, ed. John P. Kimball, 1–36. New York: New York Academic Press.