

Dynamics of Epistemic Modality

Malte Willer

This writing sample is taken from a longer version of my ‘Dynamics of Epistemic Modality’, which is available at http://www.maltewiller.net/Epistemic_Modality.pdf. The longer paper contains two sections, sections 4 and 5, that are not included in this writing sample. Section 4 mainly explains the various ways in which we may assess judgments of epistemic modality in the light of new evidence. Section 5 offers an in-depth comparison between the framework developed here and its major truth-conditional competitors.

1 Introduction

The goal of this paper is to develop an alternative to the orthodox view that our best semantic theory ascribes to epistemically modalized sentences truth-conditions relative to what is known. According to the alternative I am proposing, such sentences have content, but not truth-conditional content. Instead, the meaning of an epistemically modalized sentence is best understood in terms of its *context change potential* (CCP). Consider the following example. 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: Oh, OK, I’ll go and check.

In this simple discourse, Alex tells Mary that the keys might be in the car. Mary *learns* something by being told that the keys might be in the car: there is a *change* in her epistemic situation. Intuitively, what has happened is that she is now *aware* of the possibility of the keys being in the car. It is this *dynamic* effect of Alex’s utterance—the type of change it is designed to induce in the hearer’s information state—that I take as the starting point of my proposal. Adequately elaborated, this framework delivers a non-truth-conditional yet strictly compositional semantics.

I offer two arguments for the semantic framework developed in this paper. First, I demonstrate that, unlike the orthodox view, the framework provides a simple and unifying perspective on disagreement about epistemic modality and disagreement about matters of fact. Second, I show that the framework also explains a range of puzzling observations about *might* and *must* that go beyond the phenomenon of modal disagreement.

My plan is as follows. In the remainder of this section I provide some initial motivation for departing from orthodox semantics: such a semantics fails to deliver a simple and uniform perspective on disagreement about epistemic modality and disagreement about matters of fact. In §2 I will tell just enough of my dynamic story to offer the desired perspective and to show that it does not have a Frege-Geach problem. In §3 I show how the framework developed in this paper takes care of some further tricky data.

Virtually everyone thinks that epistemically modalized sentences have truth-conditions relative to what is known, so I should say a bit about why I reject the orthodox view. 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 people have the habit of evaluating present tense claims of epistemic modality by testing them against their own perspective. To motivate the problem, consider the following variation of the first example (from [von Fintel and Gillies \(2008b\)](#)):

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

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

One extraordinary version of the orthodoxy—the relativist version—is that epistemically modalized sentences vary in truth-value across points of assessments (a judge or some other sequence which varies with the assessor). So Alex's utterance is literally true when he is the judge and false when 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.

¹See [Egan et al. \(2005\)](#), [Egan \(2007\)](#), [Yalcin \(2007\)](#), [Stephenson \(2007, 2008\)](#), [MacFarlane \(2009\)](#).

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 (2) is problematic. 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:

- (3) 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* when they state 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 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

²See von Fintel and Gillies (2009). Kratzer (1977, 1981, 1991) is the *locus classicus* of contextualism. See also Hacking (1967), Teller (1972), and DeRose (1991) for some letters of support for this view.

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 always used to describe what is known by a contextually salient group of individuals (in most cases involving the speaker). For critical discussion see Egan et al. (2005), §2 and MacFarlane (2009), §3.

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 which is true given Alex's epistemic situation and false from Mary's perspective. New age contextualists enrich what Alex has said so that he also makes a wrong 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* when one makes 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, the role is *also* of a semantic kind: what one knows is of relevance for the truth-conditions, and hence the semantic content, of a judgment of epistemic modality. I have also outlined some motivations for adopting the naive conception. The conclusion is that we have reason to look for a theory that stays faithful to the naive conception.

In the remainder of this article I demonstrate that the naive conception is indeed a viable position. My suggestion then is to throw orthodox semantics for epistemically modalized sentences overboard. True, no one knows how to give sane truth-conditions for epistemically modalized sentences without signing up for the orthodoxy, but this just means that we should not give them truth-conditions in the first place. Our semantic story should run under the slogan that epistemically modalized sentences have content, but not truth-conditional content. Let me now show that such a story can be told.

2 Dynamics

In §1 I said that our semantics for epistemically modalized sentences should look at the context change potential (CCP) of these sentences rather than their truth-conditions. Hence the semantics I intend to develop is a dynamic semantics.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 shows that my semantics avoids the notorious Frege-Geach problem even though it is non-truth-conditional.

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

With so much focus on contexts, it is legitimate to ask what exactly they are supposed to be.

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

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

The answer depends on the specific goals of one’s dynamic proposal. I am primarily interested in the interaction between epistemically modalized sentences and an agent’s epistemic situation, and so I will treat contexts as information states.⁶ Such a dynamic approach is not only interesting due to its lack of commitment to ascribing truth-conditions to every epistemically modalized sentence. It also offers an appealing perspective on our observation about how utterances of epistemic modality are assessed by different people: since the effect of an utterance may vary from information state to information state, we should expect that speakers may have different reactions to one and the same utterance, let it be factual or modal. Thus the story I wish to tell will be in the position to offer a simple and uniform account of how speakers assess utterances in discourse.

It is common to model information states as sets of possible worlds compatible with what is known by the agent. This view has a problem, but it is not the problem that one often finds in the literature. The problem of logical omniscience—that if an agent knows ϕ , then the agent also knows all the logical consequences of ϕ —is deep and deserves further analysis, but not here. The problem I have in mind is that if we just consider the set of possible worlds compatible with what is known and call those the agent’s epistemic possibilities, then all epistemic possibilities are created equal. This may often be unproblematic, but it is inappropriate for current purposes. Let me explain.

Intuitively, an agent who learns that so-and-so might be the case becomes aware of a certain epistemic possibility—something which has been previously merely compatible with what is known becomes a ‘live’ epistemic possibility. p is a live epistemic possibility for an agent S if and only if, whenever the agent considers how things might completely be, S will include p . When this is not so yet S has no information ruling out p , then p is just a plain epistemic possibility. This intuition underlies my proposal for the semantics of sentences involving *might*: such sentences change an agent’s information state by making that agent aware of certain epistemic possibilities. In other words, such sentences change plain epistemic possibilities into live epistemic possibilities. To capture this intuition, our information states need to be fine-grained enough so that we can distinguish between what is merely compatible with the agent’s information and what is a live epistemic possibility—a possibility that the agent is aware of. This is

⁶It is no coincidence that [Veltman \(1985, 1996\)](#) treats contexts in the same way as I do, 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 why utterances of epistemically modalized sentences can be informative for the audience. So you may wish to take a look at Veltman’s work in case you are interested in the original. See also [Groenendijk et al. \(1996\)](#), [van der Does et al. \(1997\)](#), [Beaver \(2001\)](#), and [Gillies \(2001\)](#) for some applications and extensions of Veltman’s proposal. This proposal is also the one favored by [von Fintel and Gillies \(2008a\)](#), but does not play a role in their later work on epistemic modals.

why we cannot simply model information states as sets of possible worlds compatible with what is known.

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

The general idea, then, is to understand the meaning of a sentence as an instruction to move from one information state to another. *Might* is designed to transform an information state by highlighting possibilities that have so far been merely compatible with what is known. To elaborate this idea, information states must be fine-grained enough so that we can distinguish between live epistemic possibilities and what is merely compatible with what is known. Let me now move on to the details of the proposal.

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):

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

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

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. Modeling information states as mere sets of possible worlds is not good enough: this strategy would fail to provide the required distinction between live epistemic possibilities and what is merely compatible with what is known. My preferred alternative is to treat information states as sets of sets of possible worlds. The relevant definition is as follows:

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

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

To get a feel for the notion of an information state, it is helpful to make more precise the two ways in which p may be an epistemic possibility for an agent:

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

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

$$\forall \sigma \in \Sigma \exists w \in \sigma: w(p) = 1$$

There are information states in which p is an epistemic possibility— p is not excluded by the agent's information—but p is not a live epistemic possibility. In such an information state there is a state without a p -world.

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

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

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

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

⁹If necessary, the update effects of elements of \mathcal{L} can be relativized to contexts of utterance. But it is not necessary now, as epistemic modals are not context-sensitive.

is at least one p -world in σ , and returns the empty set otherwise.¹⁰ At the level of states, then, epistemic possibility is an all-or-nothing affair. That just follows from the idea that states are complete answers to the question what might be the case. No accessibility relation is appealed to or will be needed to give us the result we want.

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 fully relational: they are relations between information states.

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

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

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

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

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

The following notions will turn out useful in later discussion.

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

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

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

3. ϕ entails ψ , $\phi \models \psi$, iff $\forall \Sigma: \Sigma[\phi] \models \psi$

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

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

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

Since $\Sigma[\diamond\phi] \models \diamond\phi$, we get the desired result that an admissible update with an epistemically modalized sentence raises the prejacent to a live epistemic possibility. Specifically, any admissible update with ' $\diamond\phi$ ' raises ϕ from something merely compatible with what is known to a live epistemic possibility. And that is just what seems to happen in our very first example (1). There is a change in Mary's information state when she is told that the keys might be in the car. And the change is just as the framework predicts it to be: she is now aware of the possibility of the keys being in the car.

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 $\phi \in \mathcal{L}_0$: $\phi \models \Box\phi$

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

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

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

If ϕ is a live epistemic possibility in Σ , Σ won't admit the 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 your information state supports ϕ it also supports $\lceil \Box\phi \rceil$; but it does not follow that $\lceil \neg\phi \rceil$ is settled whenever the falsity of ϕ is a live epistemic possibility.

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

Fact 4 $\Diamond\phi \models \Box\Diamond\phi$

Thus any combination of boxes and diamonds reduces to the innermost operator. This is an acceptable result, I think, as embedding an epistemic modal under another epistemic modal does not in general have any interesting semantic effects.¹¹

This is all I have to say about the semantics of \mathcal{L} .¹² And with a bit of 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 CCP, of ϕ . How hearers assess an utterance of ϕ will then depend on the update effects of ϕ on their information states. Thus we have a two-stage picture: first determine how well ϕ plays with the current information state and then assess the assertion of ϕ on that basis.

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

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

¹¹Sorensen (2009) presents examples that purport to demonstrate that no adequate logic for epistemic modals can validate the characteristic axiom of S5. The cases he presents raise issues that are too complex to be discussed here. See Yalcin (2009a) for a response to Sorensen's critique.

¹²At least that is all I have to say without digressing too much. One thing that certainly deserves further discussion—especially when one is in a dynamic mood—is the assumption that *might* and *must* are duals. The fact that *might* and *must* behave differently when it comes to modal subordination (Roberts (1987, 1989)) speaks heavily against this view. But all this is irrelevant for current purposes, and the story could be extended to cover modal subordination, adopting some of the things said by Asher and McCreedy (2007).

In words, A will agree with ϕ if A 's information state already encodes the information encoded in ϕ . If A 's information is incompatible with ϕ , then we expect A to reject an assertion of ϕ . And finally, if A is agnostic about ϕ , then A might as well admit that ϕ is the case.

I have promised a simple and uniform perspective on Alex's and Mary's disputes about the keys, and I can now deliver. First example: Mary denies what Alex has asserted, as the content of her assertion—that the keys can't be in the car—is just the negation of what Alex has said. The difference in what is known 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 managed to develop 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. No other semantics can offer a comparably attractive perspective on disputes about epistemic modality and disputes about matters of fact. I conclude that the suggestion made in this paper is superior to alternative theories.

2.4 The Frege-Geach Problem

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. For each element of \mathcal{L}_0 we can define a set of indices at which it is true.

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

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

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

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

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 a state with such a sentence eliminates either all or no possibilities, and depending on global features of the state.¹³ Hence we will not find some proposition that mediates the kind of context change that is induced by an epistemically modalized sentence. And that is just how things should be. *Might*-statements are simply not designed to deliver some true or false description of reality: they have content, but not truth-conditional content.

What the current framework does, then, is to show how we can combine two quite distinct intuitions into one unified picture. On the one hand, it is hard to see how Stalnaker's view could be false when we look at assertions about matters of fact. On the other hand, the view also seems to be off track when it comes to discourse about what might or must be. Our dynamic framework shows how these two views can live happily with each other: some but not all context change is mediated by propositional content, and if the meaning of a sentence is its context change potential, that allows for a plausible but unified semantics for a language that includes epistemic modals.

Some people say that all merits aside, a non-truth-conditional semantics for epistemically modalized sentences fails to be viable. The plain reason is that such sentences embed (though not as freely as factual sentences), and thus a semantics for such sentences has to be truth-conditional to avoid the Frege-Geach problem.¹⁴ This is certainly a fair objection to the popular view that

¹³Another 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*: $\sigma \uparrow \phi \neq \bigcup_{w \in \sigma} \{w\} \uparrow \phi$.

¹⁴See, e.g., MacFarlane (2009), §4.

epistemic modals are used to express some kind of ‘comment’, indicating the degree or source of the speaker’s commitment to the embedded proposition.¹⁵ Such theories are clearly in bad shape when it comes to embeddings. But it should be already obvious from what has been said so far that the framework laid out here does not suffer from this problem—it gives clear answers to what happens when epistemic modals are embedded under negation, conjunction, etc. Not only that, but the treatment of negation, conjunction, etc. yields the right results. Specifically, it is straightforward to predict that ‘It might be raining’ and ‘It can’t be raining’ are inconsistent:

Fact 6 For all $p \in \mathcal{A}$: $\models \neg(\diamond p \wedge \neg\diamond p)$

This observation straightforwardly follows from our well-motivated treatment of negation: updating a state σ with ‘ $\neg\phi$ ’ leaves everything in σ that get eliminated by an update with ϕ . It follows that updating a state σ with ‘ $\diamond p$ ’ and then with ‘ $\neg\diamond p$ ’ yields the empty set. Hence updating σ with ‘ $\neg(\diamond p \wedge \neg\diamond p)$ ’ will always return σ , which just means that ‘ $\neg(\diamond p \wedge \neg\diamond p)$ ’ is a tautology.¹⁶ So there is no reason at all to believe that there is any embedding problem for the framework developed so far.

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

- (4) 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. Let me explain.

What is true is that if such verbs denote relations between individuals and truth-condition determining propositions, then every semantics will be in trouble which does not assign to epistemically modalized sentences such propositions. But we may as well understand attitude verbs as denoting relations between individuals and CCPs.¹⁷ More precisely, let s range over

¹⁵This approach to epistemic modality operators as comments on embedded propositions is certainly the standard non-truth-conditional view. See, among many others, [Kant \(1781\)](#), [Frege \(1879\)](#), [Price \(1983\)](#).

¹⁶This is a crucial difference between my account of negation and expressivist treatments of negation as developed by [Gibbard \(1992, 2003\)](#) and [Horgan and Timmons \(2000, 2006\)](#). As [Schroeder \(2008\)](#) has shown in his discussion of moral expressivism, none of the current expressivist accounts of negation can explain why ‘Murder is wrong’ and ‘Murder is not wrong’ are inconsistent.

¹⁷If 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.

indices of evaluation, e range over individual objects, t range over truth-values. If attitude verbs denote relations between individuals and propositions, they are of type $\langle\langle s, t \rangle, \langle e, t \rangle\rangle$. Things look only slightly different from the alternative perspective. States are, like propositions, of type $\langle s, t \rangle$, information states are of type $\langle\langle s, t \rangle, t \rangle$, and CCPs are of type $\langle\langle\langle s, t \rangle, t \rangle, \langle\langle\langle s, t \rangle, t \rangle, t \rangle\rangle$. Thus if attitude verbs denote relations between individuals and CCPs, they are of type $\langle\langle\langle\langle s, t \rangle, t \rangle, \langle\langle\langle s, t \rangle, t \rangle, t \rangle\rangle, \langle e, t \rangle\rangle$. The only notable 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.

One may nevertheless think that I have not solved the problem in its entirety. The issue is not one of pure logical form but rather of our common sense conception concerning belief. It is tempting to say that beliefs are essentially representational: there is a piece of reality that a belief is *about*. Hence we will leave out an important feature of beliefs about epistemic modality in case we do not provide them with truth-conditional content.

As a response, it is important to see that we do not have to think of belief as essentially representational. Instead we may think of belief *holistically*: what one believes is a matter of what the information state as a 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*:

$$(5) \quad \llbracket \text{believe} \rrbracket^i = \lambda R. [\lambda x. \exists \Sigma: x \text{ is in } \Sigma \text{ at } i \text{ and } \langle \Sigma, \Sigma \rangle \in R]$$

In other words, a sentence of the form ‘ S believes that ϕ ’ is true at an index i if and only if S is in a ϕ -supporting information state at i . The semantics of such sentences is then given in terms of their CCPs: they eliminate, from each information state, those worlds in which the subject is not in an appropriate information state. It follows that we have every reason to think that the dynamic framework developed in this article avoids the Frege-Geach Problem. The proposal gets the facts about logical connectives right and delivers a structural analysis of belief-attributions. Moreover it provides the basics of a holistic view of belief, which responds to the worry that truth-conditional content must be the hallmark of belief.¹⁸

¹⁸A worry that is closely related to the worry about belief: I have repeatedly said that one may *learn* that ϕ might be the case. This suggests that one may come to *know* that ϕ might be the case. The worry now is

3 Bonus

The framework developed in the previous section allows for a simple and uniform view on the two disputes between Alex and Mary. This is an achievement, for no other semantic theory currently on the market does the same. In addition to that, it can explain a range of puzzling data about epistemic modals. 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, 2009b) observes that the following sentence (6a) sounds terrible, and cannot even be supposed to be true:

- (6) 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 (6a) as a matter of pragmatics, analogous to Moorean paradoxical constructions of the form ‘ ϕ but I do not believe that ϕ ’. However, (6a) 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 (6a) must go deeper. The most straightforward explanation is that ‘It is raining and it might not be raining’ is a plain semantic contradiction. 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, as there is nothing contradictory about the claim that there are unknown truths.

that one cannot come to know something unless it is *true*, and hence epistemically modalized sentences must have truth-conditions. A proper analysis of learning and knowledge in the context of epistemic modality must be left to future research. It suffices to say at this point that the worry is poorly motivated. On my view, learning that ϕ might be the case should be understood as involving a cognitive enhancement that goes together with the acquisition of some new skills. This cognitive change does not consist in ruling out certain epistemic possibilities. Hence there is no reason to think that coming to know that ϕ might be the case consists in acquiring truth-conditional information about worldly facts.

3.1.2 Ignorance

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

- (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 having cancer is compatible with what they know that. It follows immediately from this that one may learn that so-and-so might be the case, in the same sense in which one may learn that so-and-so is the case. This conclusion is also supported by the following example:¹⁹

- (8) 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 is just another observation that a good theory of epistemic modals must be in the position to explain.

3.1.3 Must

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

¹⁹Remember also the discourse in (1), in which Mary seems to learn that the keys might be in the car.

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

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

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

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

Here Mary points out that the available evidence does not eliminate the possibility that Colin is out—he simply might have forgotten to turn off the lights when leaving the house. Notice, again, that Mary's reaction to Alex's 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 (10) may continue as follows:

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

Of course, conversations do not always proceed as in our toy examples. We expect there to be situations in which an agent will simply go along with an utterance of epistemic necessity, even though the recipient is aware of a possibility in which the prejacent is not true. This is not a special feature of epistemic claims: we very often revise our current beliefs in order to accommodate new information, especially when it comes from a source that we think knows much more about the issue than we do. But subtle issues such as trust and authority need not detain us here: even if we keep the degree of Mary's trust in Alex's reliability fixed, there is a clear distinction between the situation in which it is legitimate for her to reject Alex's claim of

epistemic necessity and the one in which she has no grounds whatever to reject it. The former is a situation in which Mary has grounds for believing that the prejacent might be false, the latter is a case in which the falsity of the prejacent is merely compatible with what she knows. This, I submit, is a distinction that a semantics for judgments of epistemic modality should be able to draw.

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.²⁰ So an attempt at updating one’s information state with (6a) 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 (6a) 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 (6a) is pragmatically infelicitous.

3.2.2 Ignorance

Remember the distinction between ϕ being compatible with what is settled in Σ and ϕ being a live epistemic possibility in Σ . This corresponds to the distinction between $\Sigma \triangleright \diamond\phi$ and $\Sigma \models \diamond\phi$. And this is all we need to make sense of the DeRose’s observation that agents may be ignorant of what might be the case. An agent who is in a ‘ $\diamond\phi$ ’-admitting information state is agnostic about that sentence—the information the agent possesses neither entails ‘ $\diamond\phi$ ’ nor ‘ $\neg\diamond\phi$ ’. We can make perfect sense of dialogues like the one in (7): it is compatible with what Alex and Mary know that John has cancer, but that alone does not make it a live epistemic possibility. That is, the possibility is admitted, but not supported, by the relevant information states.

²⁰But 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’.

The existence of information states that admit epistemically modalized sentences also guarantees that update with such sentences can be informative, as observed in (8). Let's first get clear on what it means for an update to be informative.

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

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

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

The following is a trivial consequence of the definitions of admission and support:

Fact 8 For all $\phi \in \mathcal{L}$, $\Sigma \in I$: If $\Sigma \triangleright \phi$, then $\Sigma[\phi] \neq \Sigma$ and $\Sigma[\phi] \neq \emptyset$

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

3.2.3 Must

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

Fact 9 For all $\phi \in \mathcal{L}$, $\Sigma \in I$: If $\Sigma \triangleright \diamond \phi$, then $\Sigma \triangleright \square \neg \phi$

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

Thus the contrast between admitting and supporting a certain epistemic possibility also allows us to capture the observations we made in connection with epistemic *must*. If an agent is in a $\ulcorner \diamond \phi \urcorner$ -admitting information state, what the agent knows is compatible with ϕ , but does not encode any grounds for believing that ϕ is a real epistemic possibility. Accordingly, the agent has no grounds for disbelieving $\ulcorner \square \neg \phi \urcorner$. Going back to the discourse in (9), it is merely compatible with what Mary knows that Colin is not in Chicago. Thus Mary's information admits Alex's claim that Colin must be in Chicago, and we predict that she does not—and in fact should not—reject Alex's utterance.

On the other hand, Mary’s information in (10) supports the epistemic possibility that Colin is out, the fact that the lights are on notwithstanding. So if we let q stand for ‘Colin is at home’, then Mary’s information state Σ_M is such that $\Sigma_M \models \diamond\neg q$ and, accordingly, $\Sigma_M[\Box q] = \emptyset$. Thus we predict that Mary may legitimately reject Alex’s claim of epistemic necessity—Colin *might* be out, and so he does not *have* to be at home.

Moving away from standard approaches by modeling information states as sets of sets of possible worlds does some genuine work. It leaves room for speakers to be agnostic about certain epistemic possibilities. This is the right diagnosis for the observations we made at the beginning of this section. Utterances of epistemically modalized sentences can be informative, and agents admit or reject claims of epistemic necessity depending on whether or not they have good reason to believe that the prejacent might be false.

We have seen that the dynamic framework developed here solves a range of problems that go beyond the phenomenon of modal disagreement. This, I submit, is another reason to believe that we have found the right semantics for *might* and *must*.

References

- Asher, N. (2007). “Dynamic Discourse Semantics for Embedded Speech Acts”. In Tsohatzidis, S. L., editor, *John Searle’s Philosophy of Language*, pages 211–244. Cambridge University Press, New York, NY.
- Asher, N. and Lascarides, A. (2009). “Another Look at Questions in Dialogue”. Forthcoming in Riester, A. and Solstad, T. editors, *Proceedings of SuB13*, Stuttgart.
- Asher, N. and McCready, E. (2007). “Were, Would, Might, and a Compositional Account of Counterfactuals”. *Journal of Semantics*, 24 (2): 93–129.
- Beaver, D. I. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford, CA.
- DeRose, K. (1991). “Epistemic Possibilities”. *Philosophical Review*, 100 (4): 581–605.
- Dever, J. (2009). “Living the Life Aquatic: Does Presupposition Accommodation Mandate Dynamic Semantics?”. <https://webpace.utexas.edu/deverj/personal/papers/aquatic.pdf>. Unpublished manuscript.

- van der Does, J., Groeneveld, W., and Veltman, F. (1997). “An Update on Might”. *Journal of Logic, Language and Information*, 6 (4): 361–380.
- Egan, A. (2007). “Epistemic Modals, Relativism, and Assertion”. *Philosophical Studies*, 133 (1): 1–22.
- Egan, A., Hawthorne, J., and Weatherson, B. (2005). “Epistemic Modals in Context”. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy*. Oxford University Press, New York.
- von Fintel, K. and Gillies, A. (2008a). “An Opinionated Guide to Epistemic Modality”. In Gendler, T. and Hawthorne, J., editors, *Oxford Studies in Epistemology 2*. Oxford University Press, Oxford.
- von Fintel, K. and Gillies, A. (2008b). “CIA Leaks”. *Philosophical Review*, 117 (1): 77 – 98.
- von Fintel, K. and Gillies, A. (2009). “‘Might’ Made Right”. <http://web.mit.edu/fintel/fintel-gillies-2008-mm.pdf>. Forthcoming in Egan, A. and Weatherson, B., editors, *Epistemic Modality*, Oxford University Press.
- Frege, G. (1879). *Begriffsschrift: Eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Louis Nebert, Halle.
- Gibbard, A. (1992). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press, Cambridge, MA.
- Gibbard, A. (2003). *Thinking How To Live*. Harvard University Press, Cambridge, MA.
- Gillies, A. (2001). “A New Solution to Moore’s Paradox”. *Philosophical Studies*, 105 (3): 237–250.
- Groenendijk, J. (2008). “Inquisitive Semantics”. In *Proceedings of SALT XVIII*, Amherst, MA. The University of Massachusetts at Amherst.
- Groenendijk, J. and Stokhof, M. (1991). “Dynamic Predicate Logic”. *Linguistics and Philosophy*, 14 (1): 39–100.
- Groenendijk, J., Stokhof, M., and Veltman, F. (1996). “Coreference and Modality”. In Lapin, S., editor, *The Handbook of Contemporary Semantic Theory*, pages 179–214. Blackwell, Cambridge, MA.
- Hacking, I. (1967). “Possibility”. *Philosophical Review*, 76 (2): 143–168.

- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst.
- Horgan, T. and Timmons, M. (2000). “Nondescriptivist Cognitivism: Framework for a New Metaethic”. *Philosophical Papers*, 29: 121–153.
- Horgan, T. and Timmons, M. (2006). “Cognitivist Expressivism”. In Horgan, T. and Timmons, M., editors, *Metaethics after Moore*, pages 255 – 298. Oxford University Press, New York.
- Kamp, H. (1981). “A Theory of Truth and Representation”. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Formal Methods in the Study of Language, Part I*, pages 277 – 320. Mathematisch Centrum, Amsterdam.
- Kamp, H. and Reyle, U. (1993). *From Discourse to the Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Press, Dordrecht.
- Kamp, H., van Genabith, J., and Reyle, U. (2009). “Discourse Representation Theory”. <http://www.ims.uni-stuttgart.de/~hans/Papers/hpl-drt.pdf>. Unpublished manuscript.
- Kant, I. (1781). *Critik der Reinen Vernunft*. Johann Friedrich Hartknoch, Riga.
- Kratzer, A. (1977). ‘What ‘Must’ and ‘Can’ Must and Can Mean’. *Linguistics and Philosophy*, 1 (3): 337–355.
- Kratzer, A. (1981). “The Notional Category of Modality”. In Eikmeyer, H.-J. and Rieser, H., editors, *Words, Worlds, and Contexts. New Approaches to Word Semantics*, pages 38–74. Walter de Gruyter & Co., Berlin.
- Kratzer, A. (1991). “Modality”. In von Stechow, A. and Wunderlich, D., editors, *Semantics: An International Handbook of Contemporary Research*, pages 639 – 650. Walter de Gruyter & Co., Berlin.
- MacFarlane, J. (2009). “Epistemic Modals Are Assessment-Sensitive”. <http://johnmacfarlane.net/epistmod.pdf>. Forthcoming in Egan, A. and Weatherson, B., editors, *Epistemic Modality*, Oxford University Press.
- Price, H. (1983). “Does ‘Probably’ Modify Sense?”. *Australasian Journal of Philosophy*, 61 (4): 396–408.

- Roberts, C. (1987). “*Modal Subordination, Anaphora, and Distributivity*”. PhD thesis, University of Massachusetts at Amherst.
- Roberts, C. (1989). “Modal Subordination and Pronominal Anaphora in Discourse”. *Linguistics and Philosophy*, 12 (6): 683–721.
- Schroeder, M. A. (2008). *Being For: Evaluating The Semantic Program of Expressivism*. Oxford University Press, New York, NY.
- Sorensen, R. (2009). “Meta-agnosticism: Higher Order Epistemic Possibility”. *Mind*, 118 (471): 777–784.
- Stalnaker, R. C. (1978). “Assertion”. In Cole, P., editor, *Syntax and Semantics 9*, pages 315 – 332. New York Academic Press, New York.
- Stephenson, T. (2007). *Towards a Theory of Subjective Meaning*. PhD thesis, Massachusetts Institute of Technology.
- Stephenson, T. (2008). “Judge Dependence, Epistemic Modals, and Predicates of Personal Taste”. *Linguistics and Philosophy*, 30 (4): 487–525.
- Teller, P. (1972). “Epistemic Possibility”. *Philosophia*, 2: 303–320.
- Veltman, F. (1985). *Logics for Conditionals*. PhD thesis, University of Amsterdam.
- Veltman, F. (1996). “Defaults in Update Semantics”. *Journal of Philosophical Logic*, 25 (3): 221–261.
- Yalcin, S. (2007). “Epistemic Modals”. *Mind*, 116 (464): 983–1026.
- Yalcin, S. (2009a). “More on Epistemic Modals”. *Mind*, 118 (471): 785–793.
- Yalcin, S. (2009b). “Nonfactualism about Epistemic Modals”. Forthcoming in Egan, A. and Weatherson, B., editors, *Epistemic Modality*, Oxford University Press.