Dynamic semantics for epistemic modality*

Matthew Mandelkern
New York University, mandelkern@nyu.edu

January 24, 2025

1 Introduction

This chapter motivates a dynamic treatment of epistemic modality, then briefly discusses some drawbacks of the classic dynamic approach. This chapter mostly comprises excerpts from my recent book (Mandelkern, 2024, Ch. 2, 4); in the rest of that book I also motivate a dynamic treatment of conditionals and anaphora, discuss its drawbacks, and develop an alternative approach.

2 The relational semantics

Epistemic modals are modal auxiliaries like 'might', 'must', 'can't', and so on (and their analogues in other languages), used as in (1)–(5), where they are interpreted in a broadly epistemic or informational sense (see 'Epistemic Modality', this volume, for an overview; here I focus exclusively on the features of epistemic modals which are of interest to motivating dynamic semantics).

(1) It might be raining.

(2) Bob must be in his office.

(3) If Sue might come, I'll make sure we have a vegan option.

(4) I've concluded that the gardener can't be the murderer.

(5) Either the butler is the murderer, or else it must be the gardener.

As a foil for the dynamic theory, I'll start by briefly explaining the standard *static* semantics for epistemic modals, which comes from Kratzer 1977, 1981, 2012b,a and treats 'might'

*To appear in the *Cambridge Handbook of Natural Language Modality*. Many thanks to the editors for the invitation to contribute this chapter.

and 'must' roughly as context-sensitive versions of the diamond and box of modal logic (I'll simplify Kratzer's view in some ways, ignoring some sources of complexity that need not concern us for present purposes; see 'The Construction of Modal Domains — the Kratzerian View', this volume, for more detail). Hence ⌜Must p⌝ is true just in case all contextually accessible worlds are p-worlds and ⌜Might p⌝ is true just in case some contextually accessible world is a p-world. The relevant accessibility relation relates a world w to any world w' which is compatible with the information that is relevant or salient in the context in question. I'll use E to denote the contextually relevant epistemic accessibility relation. So wEw' iff w' is compatible with the contextually relevant information; I'll also write $E(w)$ for the set of worlds accessible from w under E. I use $\Diamond$ for epistemic 'might', and $\Box$ for epistemic 'must'. Finally, I follow standard practice in static semantics of using double brackets for the interpretation function: given a context c, world w, and sentence p, $[\![p]\!]^{c,w}$ is the truth-value of p, assessed relative to the contextual parameters of c at the world w (relative to an implicit atomic valuation). $[\![p]\!]^c$ is, correspondingly, the set of worlds where p is true, as assessed at c—that is, $[\![p]\!]^c = \{w : [\![p]\!]^{c,w} = 1\}$. '1' and '0' stand for the semantic values 'true' and 'false', as usual. In practice, I will usually leave the contextual parameter c implicit for readability. Finally, the static version of the relational semantics which I'll take as my foil couples these treatments of epistemic modals with classical connectives. Putting all this together, we have:

- $[\![A]\!]^w = 1$ iff $\mathcal{I}(A, w) = 1$, when A is atomic
- $[\![\neg p]\!]^w = 1$ iff $[\![p]\!]^w = 0$
- $[\![p \wedge q]\!]^w = 1$ iff $[\![p]\!]^w = [\![q]\!]^w = 1$
- $[\![p \vee q]\!]^w = 1$ iff $[\![p]\!]^w = 1$ or $[\![q]\!]^w = 1$
- $[\![\Diamond p]\!]^w = 1$ iff $\exists w' \in E(w) : [\![p]\!]^{w'} = 1$
- $[\![\Box p]\!]^w = 1$ iff $\forall w' \in E(w) : [\![p]\!]^{w'} = 1$

3 There's something dynamic about epistemic modality

The key motivation for a dynamic approach to epistemic modals is the claim that the static approach misses important generalizations about how the interpretation of epistemic modals depends on its *local information*. I will present the case in two parts. First, in this section, I will draw out some high-level patterns involving the dynamics of epistemic modals, building on observations from Dorr and Hawthorne 2013 and Klinedinst and Rothschild 2012. These regularities of interpretation make for a prima facie, but inconclusive, case that there is something dynamic about epistemic modality. Then I will make the central case for a dynamic approach, which comes from the surprising incoherence of certain embedded epistemic modals.

The key observation in this section is that the interpretation of epistemic modals in right disjuncts seems to be influenced by the negation of the left disjunct. Thus consider (6):

(6) Either John is in the US or he must be in China.

(6) seems to be saying, roughly, that either John is in the US, or else the contextual information *updated with the information that John is not in the US* entails that he is in China. In other words, the 'must' in the right disjunct of (6) is intuitively restricted to a domain of quantification where John isn't in the US.

By contrast, an *unrestricted* reading for the 'must' in (6) is not particularly accessible. On this reading, (6) would say that either John is in the US, or else the information relevant in the context entails that he is in China. One way to see that this is not the prominent reading of (6) is that, in general, a disjunction should not be asserted unless both disjuncts are compatible with the speaker's beliefs (Stalnaker, 1975). So a sentence like (6) will generally be assertable only when the speaker's beliefs are consistent with John being in the US. But then the speaker should plausibly also believe that the contextually relevant information is consistent with John being in the US; hence that that information doesn't entail that John is in China. Hence the speaker believes the right disjunct is false. This kind of reasoning should thus make (6) pragmatically odd. But there is nothing odd about (6). So on the most prominent reading, the domain of quantification for 'must' in (6) seems to be restricted to worlds where the left disjunct is false.

Another way to bring out this shiftiness in interpretation, following Klinedinst and Rothschild 2012, is to consider Disjunctive Syllogism, the classically valid inference pattern which says that $p \vee q$ and $\neg q$ together entail p. Suppose we don't know where the dog is. Then (7) is intuitively true:

(7) Either the dog is inside or the dog must be outside.

It also seems true that the dog might be inside. But 'the dog might be inside' is (pretty much everyone agrees) equivalent to the negation of the right disjunct of (7), 'She must be outside'. We cannot, however, use Disjunctive Syllogism to conclude that, since the disjunction is true and the right disjunct is false, the left disjunct is true, that is, the dog is inside. This suggests that 'must' as we naturally interpret it in the right disjunct of (7) is restricted by the negation of the left disjunct, leading to a different interpretation from when it appears unembedded.

4 Incoherence

We find similar shiftiness in interpretation when it comes to modals in the consequents of conditionals (Kratzer, 1981, 1986) and in the scope of attitude predicates (Stephenson, 2007b,a;

Yalcin, 2007), which I omit here for reasons of space. Together, these bring out a suggestive pattern: epistemic modals seem to be sensitive to their local informational environment. But this doesn't yet make a strong argument for a dynamic approach as against the static view. A proponent of that view could try to account for those patterns in terms of salience or other broadly pragmatic factors, for instance by arguing that, when interpreting a modal in a right disjunct, we tend to do so with an accessibility relation which holds fixed the truth-value of the left disjunct (Dorr and Hawthorne (2013) develop a view along these lines).

A stronger argument for a dynamic *semantic* approach to epistemic modals comes from patterns of incoherence which are not predicted by the static view. These patterns are much harder to explain pragmatically, since in these cases, there are salient interpretations which would make the sentences in question coherent, but the sentences nonetheless have a strong feeling of incoherence.

4.1 Supposition

The key data which make this point involve sentences with the form $p \wedge \Diamond\neg p$ or $\neg p \wedge \Diamond p$, which Yalcin (2007) calls *epistemic contradictions*. What do epistemic contradictions show? At first blush, not very much. Consider a sentence like (8), which has the form $p \wedge \Diamond\neg p$:

(8) #It's raining and it might not be.

There is something wrong with an assertion of (8). But this can be explained on analogy with the incoherence of Moore sentences like (9) (Moore, 1942):

(9) #It's raining and I don't know it.

Whatever is wrong with Moore sentences is almost certainly pragmatic, and whatever explanation we give of the infelicity of asserting a Moore sentence can plausibly be extended to explain the infelicity of asserting an epistemic contradiction. For instance, many think that what is wrong with (9) is that it is not knowable by the person who asserts it (knowing (9) would entail knowing that it's raining, and also knowing that you don't know it's raining, and hence not knowing it's raining). Plausibly (8) isn't knowable either, for similar reasons: if 'it might not be raining' means something roughly like 'for all I know, it's not raining', then knowing (8) would, per impossibile, involve knowing that it's raining and also knowing that you don't know it's raining.[1]

Evidence that the infelicity of Moore sentences is indeed pragmatic comes from the fact that (9), while unassertable, is still clearly consistent. The evidence for this is that we can

[1] I'm not sympathetic to this standard story about Moore sentences, for reasons given in Mandelkern and Dorst 2022. But whatever story we tell, including the story we sketch there, will plausibly generalize to asserted epistemic contradictions.

imagine it to be true. Suppose that (9) is asserted by Sue; it's raining out; but Sue doesn't know that it's raining out. Then, even though (9) is not something Sue should have asserted, what it says is still *true*. Similarly, note that (10) is perfectly coherent:

(10) Suppose that it's raining and I don't know it's raining.

That is, we can easily imagine that a Moore sentence, even one indexed to the speaker, is true, even if we can't know that it is. So it seems clear that Moore sentences are consistent: whatever is wrong with asserting Moore sentences is pragmatic, not semantic, in nature.

But Yalcin (2007) observed that epistemic contradictions pattern strikingly differently in similar environments, as in:

(11) #Suppose it's raining and it might not be!

(12) #Imagine it's raining and it might not be!

There is something quite weird about these commands. What, exactly, are you being asked to imagine? It is hard to understand. This is in stark contrast to (10), where it is perfectly clear what is expected of you.

So epistemic contradictions embed, in these cases, not like Moore sentences, but rather more like ordinary contradictions. And this creates a puzzle for the static account. 'It's raining and it might not be' should, on the static account, be roughly equivalent to the corresponding Moore sentence 'It's raining and I don't know it', or, if you prefer, 'It's raining and for all we know, it's not', where the nature of the 'we' is determined by the context. Hence (11) and (12) should have a prominent interpretation on which they are roughly equivalent to (10). But they do not seem to have this interpretation at all. It *is* possible to use 'might' to access something like this, as in (13):

(13) Suppose it's raining and, for all we know, it might not be raining.

But this requires explicit relativization, or at least a lot of contextual priming. If the interpretation of 'might' depended just on broadly pragmatic matters, then it is hard to see why we *wouldn't* interpret 'might' as it appears in (11) in just the same way we interpret it in (13), that is, as tracking the group's information, which would yield a coherent interpretation of (11).

Of course, there *are* accessibility relations relative to which (11) is predicted to have the incoherent meaning it seems in fact to have. For instance, suppose that we have an accessibility relation under which all rain-worlds can access only rain-worlds (that is, we treat *rain* as "known" at all rain-worlds). Then (11) will be predicted to be incoherent; in fact, it will be equivalent to (14), since, under that accessibility relation, there is no world at all where it's

raining and might not be:

(14) #Suppose it's raining and it's not raining.

The puzzle for the static view is why, if the choice of accessibility relations were a broadly pragmatic matter, we would choose to interpret (11) in this incoherent way, instead of doing so relative to any number of other accessibility relations which are, intuitively speaking, equally salient, and which render the command in question perfectly coherent. Compare: 'John ran into Mark on the street. John was excited but he wasn't.' The second sentence does have an incoherent reading, where 'he' refers to John. But its more prominent reading is a coherent one, where 'he' refers to Mark. We naturally gravitate towards that reading in order to get a coherent interpretation; this sequence doesn't lead to the kind of bafflement that (11) does. The puzzle for static theories, then, is why we don't analogously gravitate towards the coherent reading of (11) which static views predict exist.

Hence embeddings like this provide a strong argument for encoding sensitivity to local contexts in the meaning of epistemic modals. Before turning to see how dynamic semantics achieves this, I will show that the surprising incoherence we've just seen reappears across a wide variety of embedding domains. This suggests that this phenomenon is not the product of the interaction of epistemic modals with attitude verbs in particular, but rather that epistemic contradictions are (in some sense) genuine semantic contradictions.

4.2 Other attitudes

A first, unsurprising point is that epistemic contradictions also lead to incoherence under attitude predicates other than 'suppose':

(15) Liam believes it's raining and it might not be.

(15) is felt to ascribe incoherent beliefs to Liam. The same point applies when p and $\Diamond\neg p$ are distributed across two attitude predicates, as in (16):

(16) Liam believes it's raining. He also believes it might not be raining.

While this is somewhat unsurprising, it is worth noting, since it shows that the puzzle is not peculiar to the interaction of 'might' with conjunction in particular. That is, it shows that the puzzle is not just about sentences with the form $p \wedge \Diamond\neg p$, which doesn't appear in (16).

The problem for the relational semantics is that, if 'might' can be associated with an accessibility relation tracking, say, the speakers' information, then (15) and (16) should be coherent, and roughly glossable as 'Liam believes it's raining and that the salient information leaves open that it's not raining', which of course is coherent; it's true when Liam thinks that

he has more information than is salient in the context.

4.3 Conditionals

Yalcin (2007) also observed that epistemic contradictions lead to incoherence embedded in the antecedents of conditionals, as in (17):

(17) #If it's raining and it might not be, then we won't bring the right clothes for the weather.

This is not surprising, given the infelicity of epistemic contradictions under 'suppose' and the close connection between conditionals and supposition. But it is surprising on a static approach to epistemic modals, on which (17) should have a salient interpretation where it means what (18) does:

(18) If it's raining and we don't know it's raining, then we won't bring the right clothes for the weather.

But (18) is straightforward enough to interpret, whereas (17) again feels somehow incoherent. The puzzle for the static approach is explaining why we do not simply interpret the 'might' in (17) relative to an accessibility relation which tracks something like the contextual information, which would render the two sentences roughly coherent.

4.4 Disjunction

So far we have looked at epistemic contradictions embedded under intensional operators. Focusing on these cases might lead to the conclusion that the incoherence arises because of an interaction between epistemic modals and intensional operators in particular.[2] But, importantly, epistemic contradictions also lead to incoherence under various operators which are, classically at least, extensional. This shows that the problem is not about the interaction of epistemic modals with intensional environments in particular, and instead provides an argument that we need something dynamic in our account of the connectives themselves.

Start with disjunction. Disjoined epistemic contradictions feel incoherent Dorr and Hawthorne 2013; Mandelkern 2019:

(19) #Either it's raining and it might not be; or it's sunny and it might not be.

[2]This was, broadly speaking, the conclusion of Yalcin 2007; given the points that follow, that is a problematic limitation of that approach and its descendants, as well as broadly similar approaches like that in Incurvati and Schlöder 2020.

This is unexpected from the point of view of the static theory, since disjoined Moore sentences are felicitous:

(20) Either it's raining and I don't know it, or it's sunny and I don't know it.

(20) is a (somewhat roundabout, but totally coherent) way of saying that you don't know what the weather is like. This is predicted by the standard pragmatic account of Moore sentences. On that account, what is wrong with Moore sentences is that they cannot be believed or known by the speaker. But to assert a disjunction, you do not need to know, or believe, either disjunct—you just have to know or believe the whole disjunction. And given their intuitive truth-conditions, disjoined Moore sentences are knowable.

The puzzle for the static theory is why we don't interpret (19) like (20). If the contextually salient accessibility relation tracked the speaker's or group's beliefs or knowledge, then (19) would be coherent: it would communicate roughly the same thing as (20), namely, that the contextually salient evidence leaves open both rain and sun. But that interpretation is unavailable for (19).

Importantly, the only ingredients here are conjunction, disjunction, negation; and 'might' (plus the prejacents of 'might', but there seems to be nothing interesting going on there). This suggests that the underlying puzzle is (in part) a puzzle about the connectives and their interaction with modals. And it's a puzzle that, on the face of it, requires some kind of non-classical treatment of the connectives.

To bring this out, note the contrast between (19), repeated here, and the minimal variant in (21), which is coherent, and says roughly the same thing that (20) does:

(19) #Either it's raining and it might not be; or it's sunny and it might not be.

(21) Either it's raining or it's sunny, and it might not be raining, and it might not be sunny.

This is an intriguing minimal pair, because (19) is *entailed* by (21), given classical assumptions about the connectives. (21) has the form $(p \vee q) \wedge (\Diamond\neg p \wedge \Diamond\neg q)$. By classical truth-functional reasoning, this classically entails $(p \wedge \Diamond\neg p) \vee (q \wedge \Diamond\neg q)$, which is the form of (19). In fact, a small variant makes this point even more dramatic. Compare:

(22) #Either it's raining and it might not be; or it's not raining and it might be.

(23) Either it's raining or it's not raining, and it might be raining, and it might not be raining.

Assuming that the accessibility relation for $\Diamond$ is reflexive (so that p entails $\Diamond p$), these two sentences are *logically equivalent* according to the static theory (thanks in particular to the

distributive law for $\wedge$ and $\vee$). But, again, they are intuitively very different: (22) feels incoherent, while (23) sounds like a statement of ignorance.

One thing you might have thought is that what's bad about disjoined epistemic contradictions has something to do with redundancy. They are, after all, very roundabout. But a sentence like (23) is even more redundant—after all, the first conjunct has the form $p \vee \neg p$, which is totally uninformative. But while it does indeed sound somewhat roundabout, it does not have the feeling of incoherence of (22). So the infelicity of disjoined epistemic contradictions goes beyond whatever feeling of redundancy they lead to.

More generally, these points show that there is something non-classical in the interaction of epistemic modals and the connectives. This is just what dynamic semantics says.

4.5 Quantifiers

A final argument for something dynamic in the semantics of epistemic modals comes from quantificational structures with the form $Q(p, \Diamond\neg p)$.[3] Consider:

(24) a. #Someone hiding in the closet might not be hiding in the closet.
b. #The winner might not be the winner.
c. #Every cat might be a dog.

These are all very hard to interpret. But, again, their Moorean counterparts are coherent:

(25) a. Someone hiding in the closet, is for all we know, not hiding in the closet.
b. The winner is, for all we know, not the winner.
c. Every cat is, for all we know, a dog.

Some caution is needed with these because there is a tendency to interpret 'for all we know' as taking scope over the quantified sentence, which of course leads to incoherence. But on reflection I think these can clearly be true. For instance, (25-c) is true in a case in which we are given a list of names of animals in a shelter, and we don't know which are cats and which are dogs. By contrast, (24-c) is very hard to interpret this way. But, once more, if the epistemic modals here were interpreted with contextually salient epistemic accessibility relations, then these should be interpretable in roughly the same ways.

[3] See Beaver 1994; Groenendijk et al. 1996; Gerbrandy 1998; Aloni 2000, 2001; Yalcin 2015; Moss 2018; Rothschild and Klinedinst 2015; Ninan 2018; Mandelkern 2019.

5 Dynamic epistemic modality

These incoherence data suggest that, rather than being a pragmatic phenomenon, the sensitivity of epistemic modals to their local informational environments is semantically encoded.[4] This is exactly the position that dynamic semantics takes, starting with Groenendijk et al. 1996, who incorporated the semantics for epistemic modals from Veltman 1996 into the architecture for dynamic semantics developed by Heim 1982.[5] (This is somewhat anachronistic, since the motivations I've given above were not, apart from the quantificational cases, in the literature at that time; but, as we will see, the account answers to the motivations nicely.)

Dynamic semantics was developed by Kamp 1981; Heim 1982, 1983 to explain phenomena involving anaphora and presupposition projection. While those phenomena need not concern us here, it is worth keeping that history in mind in thinking about the broader theoretical ambitions of dynamic semantics. In the dynamic architecture, sentence meanings are (possibly partial) functions which take a context to a context (*context change potentials*). These functions are standardly written in post-fix notation, so, where $[p]$ is the context change potential denoted by p, $c[p]$ is the result of applying p to the context c. Contexts for our purposes are simply sets of worlds. The Boolean fragment gets the following semantics:

- For A atomic, $c[A] = \{w \in c : \mathcal{I}(A, w) = 1\}$
- $c[p \wedge q] = c[p][q]$
- $c[\neg p] = c \setminus c[p]$
- $c[p \vee q] = c[p] \cup c[\neg p][q]$

Parentheses are omitted where there is no danger of ambiguity, so $c[p][q]$ is $(c[p])[q]$.

Note, crucially, that conjunction here is treated as successive update, first with the left conjunct, then with the right. This gives rise to a rough but natural notion of a *local context*: in a complex sentence p which contains q as a constituent, assessed in context c, the local context for q is the context which $[q]$ updates in the computation of the complex sentence's update effect. Thus, for instance, the local context of q in $p \wedge q$ is $c[p]$, since that is the context that $[q]$ updates in the course of updating c with $p \wedge q$. The key idea behind dynamic approaches to epistemic modality is that epistemic modals are, in effect, quantifiers over their local context, which check their input local context for compatibility with their prejacent.

More precisely, $\Diamond p$ checks its local context to see whether it is compatible with p. If so, the context remains unchanged; otherwise, there is a "crash" to the empty set:

[4] This is not to say that a pragmatic account is out of the question; there are incoherence phenomena that theorists aim to explain pragmatically, as in the literature on obligatory implicature stemming from Magri 2009; see Dorr and Hawthorne 2013 for a pragmatic account of the incoherence data.

[5] Modulo some small differences. The system I present here is essentially the propositional fragment from Groenendijk et al. 1996. There are many dynamic systems; my goal is to present a simple system that captures the key ideas and empirical successes.

- $c[\Diamond p] = \begin{cases} c & c[p] \neq \varnothing \\ \varnothing & \textit{otherwise} \end{cases}$

It is natural to say that c is *consistent* with p just in case $c[p]$ is non-empty. So this says that $[\Diamond p]$ takes a context c and leaves that context unchanged just in case it is consistent with p, and otherwise induces a crash to the empty set.

This is a bit puzzling from a pragmatic point of view, a point I'll return to below. What's crucial for now is that the context that $\Diamond$ checks is not the *global* context set, but rather the local context for the modal sentence—that is, the context that it encounters in the course of computing the update.

'Must' can be defined as the dual of 'might', just as in the static theory, so we have $c[\Box p] = c[\neg\Diamond\neg p] = c \setminus c[\Diamond\neg p]$. This simplifies to:

- $c[\Box p] = \begin{cases} c & c = c[p] \\ \varnothing & \textit{otherwise} \end{cases}$

In other words, 'must', like 'might', is a test—but a test for entailment rather than consistency: $[\Box p]$ checks its input context to make sure it entails p (that is, remains unchanged after update with p), and crashes otherwise.

5.1 Regularities

By tying the interpretation of epistemic modals to their local context, this approach captures both sets of patterns described above. Start with the regularities of interpretation from §3. As we saw there, a disjunction like (6) is typically interpreted so that the 'must' is restricted to a domain where the left disjunct is false:

(6) Either John is in the US or he must be in China.

This is just what is predicted by the present account. Given a starting context c, the 'must' in (6) will check whether c entails that John is in China *once c is updated with the negation of the left disjunct*, that is, with the information that John is not in the US. So (6) will be accepted by a context just in case that context entails that John is either in the US or in China. More generally, whenever c entails $p \vee q$, it entails $p \vee \Box q$. Correspondingly, Disjunctive Syllogism is not valid. Consider a context c which contains both p- and $\neg p$-worlds for some modal-free sentence p. c entails $\Diamond p$, since it includes p-worlds. And it entails $p \vee \Box\neg p$, since $c[p] \cup c[\neg p][\neg p] = c$. But, contrary to Disjunctive Syllogism, c does not entail p. The intuition behind this is that $\Diamond p$ as it appears unembedded checks the whole context for p-worlds. By contrast, $\Box\neg p$ as it appears in $p \vee \Box\neg p$ checks whether $\neg p$ is entailed by $c[\neg p]$, which, of course,

it is. So the usual classical reasoning that underlies Disjunctive Syllogism will not go through in cases like this, matching observation that Disjunctive Syllogism appears to systematically fail for disjunctions embedding modals.

5.2 Incoherence

More importantly, this approach also captures the incoherence data we reviewed. Start by considering epistemic contradictions. I'll focus on $p \wedge \Diamond\neg p$, but the reasoning for $\neg p \wedge \Diamond p$ is the same. Assume for now that p is Boolean (that is, free of modals and conditionals). Consider any context c, and suppose we update it with $p \wedge \Diamond\neg p$. $c[p \wedge \Diamond\neg p]$ is, by our entry for conjunction, $c[p][\Diamond\neg p]$. Given our entry for $\Diamond$, this is either equal to $c[p]$, if $c[p][\neg p]$ is non-empty, or else $\varnothing$, if $c[p][\neg p]$ is empty. $c[p][\neg p]$, in turn, is equal to $c[p] \setminus c[p][p]$. But now note that, in the Boolean fragment of our language, $c[p]$ is always equal to $c[p][p]$ (the modal fragment lacks this idempotence property, which is why we need to restrict our attention to Boolean sentences here, an issue we'll return to). Hence $c[p] \setminus c[p][p] = \varnothing$; hence $c[p][\neg p] = \varnothing$; hence $c[p][\Diamond\neg p] = \varnothing$; hence $c[p \wedge \Diamond\neg p] = \varnothing$, for any c and any Boolean p. In other words, Boolean epistemic contradictions are just contradictions.

The intuition behind this is that $\Diamond\neg p$, as it appears in an epistemic contradiction, checks a local context which has already been updated with p. It checks it, in particular, for consistency with $\neg p$. But, having been updated already with p, this check will always fail, since updating with p is a way of keeping only the p-worlds and getting rid of all the $\neg p$-worlds.

This is the heart of dynamic semantics' account of the incoherence of the sentences above. To see this more concretely, consider Yalcin's (11), repeated here:

(11) #Suppose it's raining and it might not be!

Without worrying about the details of the meaning of imperative mood, let's assume that (11), as addressed to a, amounts to a command to make it true that a supposes that it's raining and it might not be. And let's assume a semantics for 'suppose' parallel to Heim's semantics for 'believes' from Heim 1992; that is, where $S_{a,w}$ is the set of worlds compatible with a's suppositions at w, we have:

- $c[S_a p] = \{w \in c : S_{a,w}[p] = S_{a,w}\}$

By the reasoning above, the only $S_{a,w}$ such that $S_{a,w}[p \wedge \Diamond\neg p] = S_{a,w}$, when p is Boolean, is the empty set. So a command to make it true that $S_a(p \wedge \Diamond\neg p)$, like (11), will be equivalent to a command to make your suppositions inconsistent, naturally explaining its incoherence.

Things are similar for epistemic contradictions under other attitude predicates, as well as epistemic contradictions distributed across attitude predicates. For instance, $B_a p \wedge B_a\Diamond\neg p$ is

equivalent to $B_a(p \wedge \neg p)$, since the only state which entails both p and $\Diamond\neg p$, for Boolean p, is the empty set.

The dynamic account of epistemic contradictions in other environments will be identical. The point is simple: Boolean epistemic contradictions *are* contradictions in this framework, and moreover can be everywhere substituted for contradictions. So a conditional with an epistemic contradiction as its antecedent will be interpreted like the corresponding conditional with $p \wedge \neg p$ as its antecedent. Likewise for disjoined epistemic contradictions.

Moreover, this approach invalidates the distributive law, and hence accounts for the contrast between disjunctions with the form $(p \wedge \Diamond\neg p) \vee (\neg p \wedge \Diamond p)$ versus the classically equivalent, but intuitively very different, $(p \vee \neg p) \wedge (\Diamond p \wedge \Diamond\neg p)$. Indeed, for Boolean p, the latter is equivalent to $\Diamond p \wedge \Diamond\neg p$, just as desired, while the former is, again, contradictory.

Turning finally to quantifiers, recall that we want to account for the fact that sentences with the form $Q_x(p(x), \Diamond\neg p(x))$ are generally infelicitous. The standard dynamic treatment of quantifiers is somewhat complicated. But we can already see how the dynamic approach accounts for sentences like this with a highly simplified dynamic semantics for quantifiers, following Yalcin 2015. We simply relativize our interpretation function to a variable assignment g (a function which takes any variable to an individual), and add the obvious clause for predicates: where R is an n-place predicate, and $\mathcal{I}$ is a valuation which takes an n-place predicate and world to a set of n-tuples of individuals, we have:

- $c[R(x_1, x_2, \ldots x_n)]_g = \{w \in c : \langle g(x_1), g(x_2), \ldots g(x_n)\rangle \in \mathcal{I}(R, w)\}$

Then, where $g_{x \to a}$ is the variable assignment which takes x to a and otherwise is just like g, we can say:

- $c[some_x(p, q)]_g = \{w \in c : \exists a : w \in c[p \wedge q]_{g_{x \to a}}\}$
- $c[every_x(p, q)]_g = \{w \in c : \forall a : w \in c[p]_{g_{x \to a}} \to w \in c[p \wedge q]_{g_{x \to a}}\}$

Entries for 'the', 'most', and so on can be generated on this model. It is easy to see how an approach like this accounts for the infelicity of quantified sentences with the form $Q_x(p(x), \Diamond\neg p(x))$: just from the inconsistency of Boolean epistemic contradictions, it follows that, for Boolean $p(x)$, $some_x(p(x), \Diamond\neg p(x))$ will always take a context to the empty set. And $every_x(p(x), \Diamond\neg p(x))$ will take any context to the empty set, provided the restrictor is non-empty. Again, this is a simplified account of the quantifiers, but the simplification is irrelevant for the present points.

Finally, let me be explicit how this account predicts not only the incoherences explored above, but also the contrasts to the corresponding embedded Moore sentences—that is, how it predicts contrasts like that between 'Suppose it's raining and it might not be' versus 'Suppose it's raining and we don't know it'. The observation is simply that the semantics for attitude ascriptions we introduced above doesn't depend on the local context in an interesting way: it

is a 'worldly' semantics, in the sense that global properties of the input context don't matter for updates with attitude ascriptions. So consider a context c which contains only worlds w such that (i) it's raining in w; and (ii) the set of worlds compatible with the relevant agents' knowledge at w, $K_{a,w}$, contains worlds where it is not raining. Then c entails 'It's raining and we don't know it', $p \wedge \neg K_a p$. Hence there are non-empty (consistent) states that entail this sentence, and hence the command to suppose that it's raining but we don't know it is not a command to enter an incoherent supposition state. The explanation of the contrasts in other cases is parallel.

In sum: the dynamic approach to epistemic modality provides an elegant account of how epistemic modals are tied to their local contexts. And it builds naturally on a system that is independently motivated by considerations about presupposition and anaphora.

6 Objections

Having said that, there are reasons to think that the simple dynamic system I have sketched is not ultimately correct. In this section, I will highlight what seem to me the three most compelling worries about this approach.

6.1 Logic

The first point concerns the logic of this system, which is weak and non-classical, resulting in peculiar and empirically questionable predictions. In particular, the following classically valid schemas have counter-models in the dynamic semantics we've given so far (van Benthem 1996; Mandelkern 2020).

Non-Contradiction: $p \wedge \neg p \vDash q$

Excluded Middle: $\vDash p \vee \neg p$

$\vDash$ is semantic entailment. In a static system, we read $p \vDash q$ as saying that q is true in any model where p is. In dynamic systems, there are a few options. The most popular one is to interpret $p \vDash q$ as saying that, for any context c, $c[p] = c[p][q]$, provided both updates are defined. That is, $p \vDash q$ iff $c[p]$ entails q, when defined.[6] So, on this reading, in a dynamic framework, Non-Contradiction is equivalent to the claim that $c[p \wedge \neg p]$ is always $\varnothing$; and Excluded Middle says that $c[p \vee \neg p]$ is always c. There are other notions of entailment in dynamic semantics, but the choice between them won't matter much for present purposes, so I'll focus on this one.

[6] Without the definedness caveat (which is not always included in definitions of dynamic entailment) both Non-Contradiction and Excluded Middle would fail for a presuppositional fragment, but in ways that are not obviously problematic.

To be clear, one of the whole *points* of dynamic semantics is to capture inference patterns that are missed by classical static approaches. For instance, one way to formulate the problem posed by the incoherence data involving epistemic modals is that $p \wedge \Diamond\neg p$ appears to be a contradiction. But from the contradictoriness of $p \wedge \Diamond\neg p$ in classical logic, we would be able to infer $\Diamond\neg p \vDash \neg p$, which is obviously not valid ('It might not be raining' does not entail 'It is not raining'). So it is a signal success of dynamic semantics is that it *enriches* classical logic with the inference pattern $p \wedge \Diamond\neg p \vDash q$ (for Boolean p), while *weakening* classical logic by invalidating the pseudocomplementation rule, which says that from $p \wedge q \vDash \bot$ we can infer $q \vDash \neg p$, thereby blocking the inference from $p \wedge \Diamond\neg p \vDash \bot$ to $\Diamond\neg p \vDash \neg p$. So, again, divergence from classical logic is, in a way, the whole point of dynamic semantics. But the particular divergences involving Non-Contradiction and Excluded Middle are, I think, indefensible.

I'll focus on the former, but the issues raised by the latter are essentially the same. Recall that $c[p \wedge \neg p]$ is, by our entry for $\wedge$, $c[p][\neg p]$, which, by our entry for $\neg$, is $c[p] \setminus c[p][p]$. Non-Contradiction requires that $c[p \wedge \neg p] = \varnothing$. So, Non-Contradiction is valid just in case $c[p] \setminus c[p][p] = \varnothing$, that is, just in case $c[p] \subseteq c[p][p]$. In the system so far, all updates are eliminative: that is, for any c and p, $c[p] \subseteq c$ if defined, so in particular, $c[p][p] \subseteq c[p]$.[7] So given eliminativity, Non-Contradiction is valid iff $c[p]$ is always $c[p][p]$.

But this equality does not hold in general. For instance, consider a context s which contains two worlds, w and w', with A true at w and false at w'. Then $s[\Diamond A \wedge \neg A] = \{w'\}$: the update first checks whether s is compatible with A, which it is, so we get s back; then we update with $\neg A$, to get down to $\{w'\}$. But updating twice yields $\varnothing$. $s[\Diamond A \wedge \neg A][\Diamond A \wedge \neg A] = \{w'\}[\Diamond A \wedge \neg A]$ (by the calculation we just ran). But $\{w'\}[\Diamond A] = \varnothing$ since $\{w'\}$ is incompatible with A; so $\{w'\}[\Diamond A \wedge \neg A] = \{w'\}[\Diamond A][\neg A] = \varnothing$.

Putting this together, we can observe that $s[(\Diamond A \wedge \neg A) \wedge \neg(\Diamond A \wedge \neg A)]$ is not the empty set: we update s once with the left conjunct and get $\{w'\}$; then we update that result again and get the empty set, which we subtract from $\{w'\}$, leaving us with $\{w'\}$.

To see that this is an empirical problem, it suffices to look at English sentences like (26) and (27), which have the form $p \wedge \neg p$. These intuitively can never be true, but they are predicted by dynamic semantics to be consistent:

(26) #Latif might be sick and he isn't, and it's not the case that Latif might be sick and he isn't.

(27) #Ariel might be home, and Jane is home, and it's not the case that: Ariel might be home and Jane is home.

A corollary of this prediction is that $p \wedge \Diamond\neg p$ can also be consistent: it is easy to see that

[7] This is not true in standard dynamic approaches to anaphora. But of course, extending our language with non-eliminative operators will not let us avoid any of the failures I am bringing out here.

whenever $p \wedge \neg p$ is consistent, so is $p \wedge \Diamond \neg p$. $p \wedge \Diamond \neg p$ remains inconsistent whenever p is Boolean (i.e., free of modals). It is not clear that this limitation is correct, though the examples that would show this are perhaps too complex to be dialectically important.

6.2 Order

A second important objection to the dynamic approach is based on the observation that the incoherence phenomena that we have surveyed are order-insensitive, contrary to the dynamic picture. This point is distinct from the logical peculiarities I have just reviewed, but it is connected, since the order asymmetries predicted by dynamic semantics stem from its asymmetric treatment of the connectives, which also gives rise to the logical anomalies we have reviewed.

As we have seen, a central idea in dynamic semantics is that conjunction is consecutive update, *first* with the left conjunct, *then* the right. This kind of asymmetry can be motivated naturally from patterns involving anaphora and presupposition, which, again, were the starting point of dynamic semantics. But, by contrast, the phenomena which motivate dynamic treatments of epistemic modals appear to be symmetric in ways that, on the face of it, run counter to the dynamic account.

I use the name 'Wittgenstein sentence' to refer to sentences which conjoin p and $\Diamond \neg p$ (or $\neg p$ and $\Diamond p$) in either order, since Wittgenstein (1953) talked (briefly) about sentences with this form. I will continue to use 'epistemic contradiction' only for such conjunctions where the modal conjunct is the right conjunct, that is, sentences with the form $p \wedge \Diamond \neg p$. Dynamic semantics predicts a striking contrast between the two orders: for Boolean p, it predicts that $p \wedge \Diamond \neg p$ is inconsistent, while $\Diamond \neg p \wedge p$ is consistent.

Advocates of the dynamic approach have defended this contrast. Thus e.g. Groenendijk et al. (1996) point to the differences between the following sequences:

(28) It might be raining outside. [...] It isn't raining outside.

(29) It isn't raining outside. [...] It might be raining outside.

There is indeed an important difference between these sequences: (28) describes an ordinary, "monotonic" evolution of information, while (29), on the face of it, encodes a change of mind. Order does matter in the processing of epistemic modals. But, as Yalcin (2012b) rightly emphasizes, the existence of order effects *across* sequences does not show there are semantically encoded order effects *within* sentences, which is what dynamic semantics predicts. To find evidence for that—that is, for a dynamic *semantics* rather than (what everyone agrees we need) a dynamic theory of updating—we would need evidence that conjunctions with the modal in the left conjunct embed coherently by contrast to epistemic contradictions.

As far as I can tell, there is no such evidence: all Wittgenstein sentences embed in the same way. Here are representative pairs:

(30) a. #If it's raining but it might not be, then we won't have the right gear.
b. #If it might not be raining but it is raining, then we won't have the right gear.

(31) a. #Suppose it's not raining and it might be.
b. #Suppose it might be raining and it's not.

(32) a. #John believes it's not raining, and he believes it might be raining.
b. #John believes it might be raining, and he believes it's not raining.

(33) a. #Either it might not be raining and it is; or it might not be sunny and it is.
b. #Either it's raining and it might not be; or it's sunny and it might not be.

(34) a. #It could be that it's not raining but it might be.
b. #It could be that it might be raining but it's not.

(35) a. #Someone I know is sick but might not be sick.
b. #Someone I know might not be sick but is sick.

The variation in order doesn't seem to make a difference.

Although dynamic semantics predicts a simple binary distinction between $\Diamond\neg p \wedge p$ (consistent) versus $p \wedge \Diamond\neg p$ (inconsistent for Boolean p), its verdicts on these examples is actually mixed, and depends on the details of the embedding operator: in standard dynamic treatments, some embedding operators care about *coherence* rather than consistency, and neither form is coherent, in the sense that there is a context c which remains unchanged after update with that sentence. For instance, the standard dynamic semantics for the conditional predicts a stark contrast between $(p \wedge \Diamond\neg p) >_i q$ (never defined for Boolean p, where $>_i$ is the indicative conditional) vs. $(\Diamond\neg p \wedge p) >_i q$ (which is predicted to be consistent and coherent, and, indeed, equivalent to $\Diamond\neg p \wedge (p >_i q)$ wherever defined). By contrast, the standard dynamic semantics for attitudes predicts that order does *not* matter for epistemic contradictions embedded under attitude predicates, since, as we have seen, attitude predicates are given a semantics in terms of what they *entail*. When it comes to disjunctions, dynamic semantics again predicts a difference: $(p \wedge \Diamond\neg p) \vee (q \wedge \Diamond\neg q)$ is inconsistent (for Boolean p), while $(\Diamond\neg p \wedge p) \vee (\Diamond\neg q \wedge q)$ is consistent (but not coherent). For modals, dynamic semantics again predicts that $\Diamond(p \wedge \Diamond\neg p)$ is inconsistent (for Boolean p), while $\Diamond(\Diamond\neg p \wedge p)$ is both consistent and coherent, as is $some_x(\Diamond\neg p(x), p(x))$.

Of course, one could try to wash out the order sensitivities by writing semantics for all embedding operators in terms of coherence rather than consistency. The result might be empirically more adequate, but it would be puzzling why we introduced a semantical-logical distinction in the first place only to render it invisible.

6.3 Pragmatics

A final worry about the dynamic approach, in particular the idea of a *test semantics* for epistemic modality, concerns pragmatics: that is, how the dynamic theory of the meaning of epistemic modals can account for the way people *use* epistemic modal claims in conversation. A natural thought about the pragmatic theory that should go with dynamic semantics is simply that, when p is asserted and accepted at a context whose common ground is c, we apply $[p]$ to c, so that the posterior context is $c[p]$. If we had a simple pragmatic picture like this, then the proposed semantics for epistemic modals and conditionals is obviously untenable. Consider a conversation with context c in which $\neg p$ remains an open possibility. Then someone we trust asserts $\Box p$. Intuitively, this is an unproblematic kind of situation: one that represents ordinary growth of information. But according to dynamic semantics, $c[\Box p] = \varnothing$, since $c[p] \neq c$, so the result of updating with $\Box p$ will be just as if we had updated with $p \wedge \neg p$.

A little reflection shows that saddling dynamic semantics with such a simple pragmatics is unfair. Proponents of dynamic semantics need not abjure the many motivations for rich pragmatic theory, even if they are moving some of what is traditionally reserved for pragmatics into the semantics. It is reasonable for them to tell some bridging story which gets us from the *semantic* update effect of modal claims to the characteristic pragmatic effect we observe. But we still need an actual theory of what we do when someone asserts $\Box p$—and why an assertion like that strikes us so differently from $p \wedge \neg p$ even when they are contextually equivalent according to dynamic semantics. Likewise, we need a story about why $\Diamond p$ can feel informative. It is by no means impossible to see how such a story might go (see especially Willer 2013 for work on this), but it has not to my knowledge been developed in a fully satisfying way.

7 Conclusion

Problems like these mean that the classic dynamic approach to epistemic modality sketched above does not have many defenders left. Still, I think it remains an important system for study, because it illuminates the very substantial and intriguing motivations behind such a system, which must be accounted for in any reasonable approach. In addition to these, there are striking and similar motivations for a parallel dynamic approach to indicative conditionals (Dekker, 1993; Gillies, 2004; Boylan and Schultheis, 2022), and perhaps also for non-epistemic modals and subjunctive conditionals (though the motivations are more indirect; see Mandelkern 2024 for discussion of both cases).

Some recent work has tried to account for those motivations with relatively close variations on the dynamic system above which aim to avoid some of these objections (see e.g.

Willer 2013; Yalcin 2012a; Goldstein 2019; Gillies 2020). Other approaches try to recombine ingredients from broadly dynamic approaches in new ways (e.g. Klinedinst and Rothschild 2012; Mandelkern 2019, 2024). Others try to capture these data either on broadly pragmatic terms (Dorr and Hawthorne, 2013; Stojnić, 2016) or with semantic approaches that are far from dynamic semantics (e.g. Yalcin 2007; Hawke and Steinert-Threlkeld 2020; Incurvati and Schlöder 2020; Aloni et al. 2022; Holliday and Mandelkern 2024).

Hence there is no consensus on the correct account of the facts that motivate dynamic semantics, or even the general contours of the correct account of them; but there should be consensus that they constitute a central desideratum for any theory of epistemic modality—indeed, a domain which will be potentially revealing, not just about the correct semantics for modality, but also about how to think about content and (local) information in general.

References

Aloni, M. (2000). Conceptual covers in dynamic semantics. In Cavedon, L., Blackburn, P., Braisby, N., and Shimojima, A., editors, *Logic, Language and Computation*, volume III.

Aloni, M., Incurvati, L., and Schlöder, J. J. (2022). Epistemic modals in hypothetical reasoning. *Erkenntnis*.

Aloni, M. D. (2001). *Quantification Under Conceptual Covers*. PhD thesis, University of Amsterdam, Amsterdam.

Beaver, D. (1994). When variables don't vary enough. In Harvey, M. and Santelmann, L., editors, *Semantics and Linguistic Theory (SALT)*, volume 4, pages 35–60.

van Benthem, J. (1996). *Exploring Logical Dynamics*. Center for the Study of Language and Information, Stanford, CA.

Boylan, D. and Schultheis, G. (2022). The qualitative thesis. *The Journal of Philosophy*, 119(4):196–229.

Dekker, P. (1993). *Transsentential Meditations: Ups and Downs in Dynamic Semantics*. PhD thesis, University of Amsterdam.

Dorr, C. and Hawthorne, J. (2013). Embedding epistemic modals. *Mind*, 122(488):867–913.

Gerbrandy, J. (1998). Identity in epistemic semantics. In *Third Conference on Information Theoretic Approaches to Logic, Language and Computation*, pages 147–159.

Gillies, A. (2004). Epistemic conditionals and conditional epistemics. *Noûs*, 38(4):585–616.

Gillies, A. S. (2020). Updating data semantics. *Mind*, 129(513):1–41.

Goldstein, S. (2019). Generalized update semantics. *Mind*, 128(511):795–835.

Groenendijk, J., Stokhof, M., and Veltman, F. (1996). Coreference and modality. In Lappin, S., editor, *Handbook of Contemporary Semantic Theory*, pages 179–216. Oxford: Blackwell.

Hawke, P. and Steinert-Threlkeld, S. (2020). Semantic expressivism for epistemic modals. *Linguistics and Philosophy*.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. 2011 edition, University of Massachusetts, Amherst.

Heim, I. (1983). On the projection problem for presuppositions. In Barlow, M., Flickinger, D. P., and Wiegand, N., editors, *The West Coast Conference on Formal Linguistics (WCCFL)*, volume 2, pages 114–125. Stanford, Stanford University Press.

Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9(3):183–221.

Holliday, W. H. and Mandelkern, M. (2024). The orthologic of epistemic modals. *Journal of Philosophical Logic*.

Incurvati, L. and Schlöder, J. J. (2020). Epistemic multilateral logic. *Review of Symbolic Logic*.

Kamp, H. (1981). A theory of truth and semantic representation. In et al., J. G., editor, *Formal Methods in the Study of Language*. Mathematisch Centrum.

Klinedinst, N. and Rothschild, D. (2012). Connectives without truth tables. *Natural Language Semantics*, 20(2):137–175.

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. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. de Gruyter.

Kratzer, A. (1986). Conditionals. *Chicago Linguistics Society*, 22(2):1–15.

Kratzer, A. (2012a). The notional category of modality. In Kratzer, A., editor, *Modals and Conditionals: New and Revised Perspectives*, pages 21–69. Oxford University Press, Oxford.

Kratzer, A. (2012b). What 'must' and 'can' must and can mean. In *Modals and Conditionals*, pages 1–20. Oxford University Press.

Magri, G. (2009). A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural Language Semantics*, 17:245–297.

Mandelkern, M. (2019). Bounded modality. *The Philosophical Review*, 128(1):1–61.

Mandelkern, M. (2020). Dynamic non-classicality. *Australasian Journal of Philosophy*, 98(2):382–392.

Mandelkern, M. (2024). *Bounded meaning: The dynamics of interpretation*. Oxford University Press.

Mandelkern, M. and Dorst, K. (2022). Assertion is weak. *Philosophers' Imprint*, 22(19):1–20.

Moore, G. (1942). A reply to my critics. In Schilpp, P., editor, *The Philosophy of G.E. Moore*. Northwestern University, Evanston, IL.

Moss, S. (2018). *Probabilistic Knowledge*. Oxford University Press.

Ninan, D. (2018). Quantification and epistemic modality. *The Philosophical Review*, 127(2):433–485.

Rothschild, D. and Klinedinst, N. (2015). Quantified epistemic modality. Handout for talk at Birmingham.

Stalnaker, R. (1975). Indicative conditionals. *Philosophia*, 5(3):269–86.

Stephenson, T. (2007a). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525.

Stephenson, T. (2007b). A parallel account of epistemic modals and predicates of personal taste. In Puig-Waldmueller, E., editor, *Sinn und Bedeutung*, volume 11, pages 583–97, Barcelona. Universitat Pompeu Fabra.

Stojnić, U. (2016). One's *Modus Ponens*: Modality, coherence, and logic. *Philosophy and Phenomenological Research*, 93(3):1–48.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261.

Willer, M. (2013). Dynamics of epistemic modality. *Philosophical Review*, 122(1):45–92.

Wittgenstein, L. (1953). *Philosophical Investigations*. Wiley-Blackwell, 2001, 3rd edition.

Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.

Yalcin, S. (2012a). Context probabilism. In Aloni, M., Kimmelman, V., Roelofsen, F., Sassoon, G. W., Schulz, K., and Westera, M., editors, *The 18th Amsterdam Colloquium*, pages 12–21.

Yalcin, S. (2012b). Dynamic semantics. In Russell, G. and Fara, D. G., editors, *Routledge Companion to Philosophy of Language*. Routledge.

Yalcin, S. (2015). Epistemic modality *de re*. *Ergo*, 2(19):475–527.