

ASSERTION AND MODALITY

Fabrizio Cariani

[Forthcoming in *The Oxford Handbook of Assertion*, S. Goldberg (ed.), Oxford University Press.]

Abstract

This essay is an opinionated exploration of the constraints that modal discourse imposes on the theory of assertion. Primary focus is on the question whether modal discourse challenges the traditional view that all assertions have propositional content. This question is tackled largely with reference to discourse involving epistemic modals, although connections with other flavors of modality are noted along the way.

My assignment for this handbook — to write about assertion and modality — presupposes that there is an interesting connection between these two. That presupposition alone might, at first, seem surprising. We do not expect there to be important connections between speech acts and specific linguistic categories; this volume does not contain entries about, say, assertion and pronouns. Why is modal language any different? The answer I explore is that modality poses a distinctive, broad and relatively unified set of challenges to the development of a theory of assertion. The goal of this chapter is to acquire a clear view about the content and scope of these challenges. Although we aim for a clear view, we will not be able to do better than gaze down at the territory from 50,000 feet. My discussion will be extremely light on technical details, on matters of compositional implementation, and even on careful discussion of data. (All those things are at the heart of the project of natural language semantics, but there is no substitute for mastering the works surveyed here.) The case at the center of the paper is epistemic modality, but I

also note some important extensions and applications concerning other flavors of modality.

1 Content skirmishes.

Stalnaker's classic "Assertion" begins with three truisms about assertion. The very first of these is that "assertions have content; an act of assertion is, among other things, the expression of a proposition — something that represents the world as being a certain way" (1978, p.78). Stalnaker's truism rings true of a great many assertions. Consider, first, assertions performed by uttering ordinary factual sentences like:

- (1) There is chocolate in the cupboard.

Say that **C** is the content of such an assertion (in context). It is reasonably clear how **C** represents the world. Plausibly, representing the world is, at least in part, a matter of distinguishing possibilities. So, **C** must distinguish those possibilities that contain chocolate in the (contextually relevant) cupboard from the rest. A successful assertion of **C** has the effect of ruling out from consideration the latter set of possibilities. Related to this representational function, contents like **C** have distinctive cognitive and conative roles. One can have, and maintain, a degree of credence in **C**—perhaps a numerical one, like .5. One can also hope that **C** is true and prefer the scenario it represents to others.

A terminological break before we go any further: in the previous paragraph, I stuck to the phrase *assertion made by uttering sentence s*. That is my tribute to the standard ontology of assertions according to which what speakers assert are propositions, not sentences. Sentences are *vehicles*, not objects, of assertions. However, it will be convenient to abuse terminology and speak as if we assert sentences.

The pattern illustrated by (1) holds for many assertions of non-epistemic modal sentences. Consider:

- (2) I can reach the chocolate in the cupboard.

By asserting (2), I represent the world as one in which I have the ability to reach the chocolate in the cupboard.¹ The same cognitive and conative interaction pattern arises here as well: it makes sense to have, and

¹My having that ability might itself be a modal property. See Mandelkern *et al.* (2017) for an advanced proposal about the semantics of ability statements.

maintain, a degree of belief in the content of (2). It also makes sense to hope that one has the relevant ability, and to prefer that scenario to the alternatives.

Now for one of the interesting cases: how do assertions involving unembedded epistemic possibility modals represent the world? Consider:

(3) There might be chocolate in the cupboard.

What possibilities do I rule out by asserting (3)? This is not clear at all. Surprisingly, it is not even clear that I necessarily rule out *any* possibilities. In some cases, I mean to ensure that we keep some possibilities relevant. In others, I mean to expand the range of possibilities we are talking about, so as to include the possibility that chocolate is in the cupboard. These do not seem to be effects we can easily recover if asserting (3) is a matter of ruling out possibilities. The puzzlement seems even sharper when we consider future-directed epistemic possibility claims, like:

(4) The tax plan might hurt the economy.

What future outcomes would falsify (4)?

These epistemic possibility claims do not seem to have the cognitive and conative roles we assigned to non-modal contents. It is quite difficult to assign degrees of belief to the contents of (3) and (4).² And it is hard to see what sort of state of the world one hopes will result when one hopes that P is epistemically possible.

To be clear: these disanalogies are itches, not arguments. They are reasons to inquire into whether it makes sense to associate these sentences with propositions. They also are reasons to go into this inquiry with an open mind. But in the end, there may well be other explanations for why the analogy between non-modal and epistemic modal sentences is strained. I discuss proper challenges to propositionalist accounts of epistemic modality in the next section.

There is a canonical answer to the questions about the truth-conditions of (3) and (4). According to contextualists, (3) is true iff the contextually relevant information is compatible with there being chocolate in the cupboard (false, otherwise). Similarly, (4) is true iff the information

²With some effort, we can come with *recherché* contexts in which this is possible. But the difficulty in producing these contexts and their rather sparse distribution is already enough to sustain the point I am making here.

that is salient in the context of utterance is compatible with the tax plan's hurting the economy. What makes these views *contextualist* is the idea that $\diamond P$ (read: *it might be that P*) may have different truth-values in different contexts, which in turn is due to the fact that $\diamond P$ expresses different propositions in different contexts.³

We will consider more precise versions of contextualism in section 2. For now, it will help our view of the territory if we fix some terminology. Draw an initial distinction between *propositions* and *contents*. These terms have lots of uses, but as I will use them here, propositions are objects that, in Stalnaker's words, represent the world as being a certain way. Contents are whatever objects play the content role in our theory of communication and intentionality. Given this distinction, Stalnaker's first truism consists of two separate claims: one, that assertions have contents; two, that those contents are propositions. We will see some views that endorse the first but not the second claim.

Contextualism is the most prominent way to bring epistemic modals back in line with Stalnaker's first truism, but it is not the only option. Invariantist approaches have the same effect (Braun 2012). To have a name for the broader class of approaches, say that a *T* is a *propositional theory* of some sentences *S* if and only if *T* entails (i) that all of the *S*s are eligible vehicles of assertion and (ii) that all of the *S*s express propositions. Contextualism and invariantism yield propositional accounts of sentences involving unembedded epistemic modals, like (3) and (4). Any view that rejects either (i) or (ii) is *non-propositional*.

Non-propositionalism about epistemic modal sentences is close to incompatible with Stalnaker's first truism. After all, if the problematic utterances of $\diamond P$ are assertions and the non-propositional diagnosis is correct, some assertions do not express propositions. In fact, there are two ways of going down this "assertions without propositions" path. One is to accept that assertions must have content but deny that contents must be propositions (in the strict sense in which I am using the word *proposition*). I discuss this in section §4.3. The other way is to reject the idea that all assertions must have content—whether propositional or of a different kind. This sounds heretical but I am going to suggest (in §4.2) it is a possible way of interpreting the framework of *update semantics*.

Before moving on, let me note that other kinds of modals generate patterns that resemble what we saw with epistemic possibility claims. Consider:

³MacFarlane (2011a, 2014) calls this *indexical contextualism* so as to distinguish it from a non-indexical version of contextualism. That difference will not matter here.

- (5) There probably is chocolate.
- (6) You may buy chocolate.
- (7) You should donate to the local food depository.
- (8) There should be chocolate in the cupboard.

The modals in (6)-(7) have a deontic flavor; the modal in (8) is sometimes treated as epistemic modal (Cariani, 2013) but has been argued by Yalcin (2016) to be a modal that tracks normality (I'm happy to go along with Yalcin's proposal). The relation between these modal claims and epistemic possibility claims is that standard utterances of these sentences do not seem well understood as ways of describing the world. Yet, they also have some of the hallmarks of assertion.⁴

In the rest of the essay, I will call the class consisting of all these modal claims the *disputed class*. This unifying terminology is a convenient device to remind ourselves that some of the points we will make about epistemic possibility have tight correlates with other kinds of claims. It is not meant to signal that I think the arguments against propositional accounts have equally strong correlates.

2 Challenges against propositionalism.

Propositionalists must explain what possibilities are ruled out by assertions of unembedded epistemic modal sentences. As noted, contextualists maintain that an assertion of (3)—*there might be chocolate in the cupboard*—represents the world as one in which some contextually salient state of information is compatible with there being chocolate in the cupboard.

Different versions of contextualism vary in how they identify the relevant information state. A simple but unpopular option is *solipsistic* contextualism — the view that the salient information state is the speaker's. Most contextualists follow Kratzer (1977, 1981, 1991, 2012) in allowing the relevant information state to be some salient body of evidence, perhaps the information that is collectively possessed by a relevant group.⁵

⁴To the list exemplified by (5)-(8), we might add indicative conditionals. The literature on non-propositionalist theories of conditionals is a subject matter in its own right. For two recent contributions that directly bear on the topic of the present essay, see Stalnaker (2011) and Rothschild (2013).

⁵For variations on this theme, with different degrees of departure from Kratzer, see DeRose (1991); Dowell (2011); von Stechow and Gillies (2011); Yanovich (2013); Stojnić

If standard contextualism about the disputed class of modal sentences is right, there is no need to intervene on the theory of assertion. Indeed, part of what it is for standard contextualism to be right is that whatever general theory of assertion we accept for ordinary non-modal sentences, should apply to assertions of sentences in the disputed class.⁶ This is often touted as an advantage of the contextualist picture.

In this section, I provide a short review of some well-known challenges against standard contextualism along with some of the most salient replies. The main emphasis is on arguments that are at least in part empirical. The literature also contains several discussions that are more theoretical in character.⁷

2.1 Rejection.

One prominent argument against contextualism involves the correct modeling of rejection (Egan *et al.* 2005, Stephenson 2007, MacFarlane 2011a, MacFarlane 2014, §10.1). Suppose that Stephanie utters (3) while on the phone with me; unbeknownst to her, I am looking directly into the relevant cupboard and I know, as well as I know anything else, that it contains no chocolate. Then, the objection goes, it is legitimate for me to reject Stephanie's assertion. I might say: *that's wrong! there is no chocolate!* This kind rejection would not make sense had Stephanie asserted *I don't know that there isn't chocolate in the cupboard*. If so, solipsistic contextualism must be wrong. Moreover, if successful, this argument should generalize to flexible versions of contextualism. As long as the relevant body of information is fixed rigidly by the context of utterance, and as long as it leaves open whether there is chocolate in the cupboard, it should not be possible to reject it from a more informed point of view.

"Eavesdropper" variations amplify this point: disagreeing with an epistemic modal claim from a more informed point of view appears to be possible even for individuals who are clearly not salient in the context. Imagine eavesdropping on a conversation between Stephanie and me.

(forthcoming).

⁶Stojnić (forthcoming) argues that this still leaves contextualists free to develop rich and novel accounts of the interaction between context and content.

⁷Yalcin (2011) objects that contextualism gets the subject matter of epistemic claims wrong (the subject matter of *it might be raining* is the weather, not some body of information). Moss (2018, §2.2) defends her proposal that assertions have probabilistic contents on the basis of different sort of high-level virtues. For example, she argues that probabilistic contents help unify the "account of the communication of full beliefs and probabilistic beliefs" (p. 27).

The background, as usual, is that neither Stephanie nor I know whether there is chocolate in the cupboard; she utters (3). Unbeknownst to us, you are hidden under an invisibility cloak and staring directly into the empty cupboards. You lift your invisibility cloak, and reject her utterance by saying *that's wrong, there is no chocolate in those cupboards!*. If it is legitimate for eavesdroppers to reject an assertion in this way, contextualists must include in the relevant body of information the information of people who are not parties to the conversation. But that would make it a mystery of why we take ourselves to be warranted in making these assertions in the first place.

Arguments from rejection have been controversial. Dowell (2011) defends standard contextualism by suggesting that these disagreement cases are underdescribed. Expand one way, and the original utterance doesn't sound warranted; expand the other and the rejection won't sound great. Similarly, von Fintel and Gillies (2008) argue that the sort of rejection we saw in the previous paragraph is not always appropriate, especially if there is a large time lag between the assertion and the rejection. One aspect of the picture of epistemic discourse these critics have in mind can be illustrated with a nice quote from a popular book about the psychology of prediction:

[...] the Harvard economic historian and popular commentator Niall Ferguson told an interviewer in January 2012 "The Greek default may be a matter of days away". Was Ferguson right? [...] let's imagine that there was no default of any kind. Could we say Ferguson was wrong? No. He only said Greece "may" default and "may" is a hollow word. (Tetlock and Gardner, 2015, p.59)

The authors' main worry is the weakness and flexibility of epistemic possibility claims (the quote is extracted from an impassioned plea for replacing them with precise probabilistic language in prediction). Speaking for myself, I don't mind the slack of *may*. But the quote also offers a clean and natural data point about rejection: Ferguson's prediction reads like a hedged claim; the use of *may* makes it really unnatural to say he was wrong. We get an even stronger judgment of this sort if we consider a slight variation on this case involving numerical expressions of subjective uncertainty. Suppose that a counterfactual version of Ferguson, call him Schwerguson, had said *there is a 40% chance that Greece will default by the end of the year*. This presumably entails what Ferguson actually said. So, if it's appropriate to reject

Ferguson's claim, it should be appropriate to reject Schwerguson. And yet, it is intuitively inappropriate, at year's end, to say that Schwerguson was wrong. If we were really serious, we might score Schwerguson's prediction with a numerical measure like the Brier score. But even if Schwerguson doesn't get a Brier-score gold star, we are not in a position to reject his claim—certainly not by saying that he was wrong.

One of the most intriguing developments surrounding rejection-based arguments against contextualism is the development of a small but thriving experimental literature challenging some of these data. Knobe and Yalcin (2014) find experimental evidence that the pattern of rejection in these challenges is not one that people find plausible. Participants in Knobe and Yalcin's experiment are told that there is evidence for the conclusion *Fat Tony might be dead* (for short DEAD). They are also told that DEAD is false—Fat Tony is, in fact, alive. Knobe and Yalcin's finding is that participants who are given this assessment overwhelmingly judge DEAD to be false but mostly judge $\diamond(\text{DEAD})$ to be true. (Their task tested this in a somewhat indirect way: see their descriptions of experiments 1 and 2 for the relevant details.) Knobe & Yalcin note that the relativist should predict much higher endorsement rates for the rejection judgment. In a recent paper, Khoo and Phillips (ms.) have contested whether this experiment successfully undermines the relativist position. Khoo and Phillips note that relativists do not have to say that the body of information that matters in a context of assessment is the assessor's—they too could go flexible. Indeed, flexible relativism is MacFarlane's official position (see MacFarlane 2014, §10.4.4). Khoo and Phillips run a more elaborate set of experiments that they interpret as making trouble for both relativism and contextualism.⁸

2.2 Retraction.

Suppose again that Stephanie asserts (3) while in an information state that is compatible with chocolate in the cupboards. Shortly thereafter, she gets access to the cupboards. No chocolate in sight. In such a case, the objection goes, it is legitimate for her to retract her assertion. Reporting on her earlier utterance, Stephanie might say *I was wrong*. This is a problem for the contextualism: after all, her original assertion was warranted by her epistemic state at the time she made it. If her assertion was warranted, the permissibility of her retraction is not expected under

⁸For other recent experimental investigations of modal disagreement and rejection (sometimes beyond the epistemic domain), see Khoo (2015); Khoo and Knobe (2016).

any version of contextualism.

As in the case of rejection, these retraction arguments have been the object of criticism. In particular, von Fintel and Gillies (2008) deny MacFarlane's data concerning the permissibility of taking back utterances by saying *I was wrong*. Indeed, resisting the rejection arguments should naturally lead to resisting these retraction data as well, since retracting by uttering *I was wrong* is a way for a future self to reject an assertion by an earlier self.

In addition, Knobe and Yalcin's (2014) experimental studies also tested judgments about retraction. Participants in their retraction experiment were given an extended version of the Fat Tony story. In this extended version, the inquiring agent finds out the falsehood of P and retracts the original assertion by saying *I was wrong*. Participants in the experiment tended to agree with this self-assessment in the non-modal case (i.e. when the speaker had said *Fat Tony is dead*, but it turns out he isn't) but were evenly split in the modal case.

In my view, this is a node in the dialectic in which thinking in a more intuitive way about retraction offers a better view of the anti-contextualist point. One of main points of retracting assertions is to stop others from relying on them. This can be accomplished without having the speaker declare that they were wrong: they might just say *I take that back* or *I take back what I said*. Consider this case: on Monday I tell my roommate that he might have failed his chemistry class. On Tuesday, my roommate makes an important life decision on the basis of my assertion. On Wednesday, I learn that he hasn't failed it. At this point, it behooves me to take back my Monday utterance, if I can. This sort of retraction might not be best understood in terms of utterances of *I was wrong*. For this reason, I think that the real test of retraction data (including for experiments) ought to revolve around weaker phrasings such as *I take that back*. In fact, *I take that back* is the canonical form of retraction in the most guarded moments of MacFarlane's book (see e.g. §1.2.2). The change in phrasing should not affect the structure of the anti-contextualist argument: the pattern will be problematic for contextualism if there are cases in which the speaker was entitled to assert $\Diamond P$ in initial context, but ought to retract it in the more informed context.

2.3 Epistemic Contradictions.

Yalcin (2007, 2011) develops a different kind of argument against contextualism. His argument takes off from the observation that P entails

$\Diamond P$ but not vice-versa. If so, $\neg P$ and $\Diamond P$ should be consistent. Moreover, if we think in terms of possible worlds, there should be a possible world at which P is false but $\Diamond P$ is true. Both predictions seem problematic: for one thing, the conjunction $\neg P$ & $\Diamond P$ sounds bad to the point of inconsistency (Yalcin calls sentences that instantiate these schema *epistemic contradictions*). Moreover, it is hard to understand what a world would have to be like in order to make this conjunction true. Yalcin extends this argument by noting that, if the badness of epistemic contradictions was just a form of pragmatic infelicity—similar to what happens with Moore’s paradox—it should disappear when these statements are embedded under attitude verbs like *suppose*. But it is clear that the badness of epistemic contradictions persists.

(9) Suppose there is a fire but I don’t believe there is one.

(10) #Suppose there is a fire but there might not be one.

Unlike (9), (10) is clearly contradictory. Because the standard pragmatic effects involved in Moore’s paradox should be suppressed in these supposition environments, this observation should block the plausibility of the pragmatic account.

Epistemic contradictions pose two problems for a theory of epistemic modality. At the compositional level, they require the design of a semantic framework that predicts why asserting (and embedding) epistemic contradictions is bad, while also keeping the asymmetric entailment from P to $\Diamond P$. There are various ways of accomplishing this: (Yalcin, 2007, §3) favors a style of semantics he calls *domain semantics*. In domain semantics for epistemic possibility, sentences are evaluated relative to (at least) two parameters, an information state and a world. However, the epistemic modal $\Diamond P$ is only sensitive to the information-state coordinate. For instance, the domain semanticist might say that $\Diamond P$ gets the semantic value 1 relative to an information state i and a world w if and only if some worlds in i make P true. As the point is sometimes put, the semantic value of $\Diamond P$ consists of a global constraint on i and not a condition on which worlds are true.

While this has invited descriptions of Yalcin’s compositional semantics as ‘expressivist’, and so non-propositional, Yalcin is actually clear that the domain semantics by itself does not settle questions about content. In particular, there are ways of defining the content of epistemic modal claims within domain semantics that would be friendly to a propositionalist. One might, for instance, suppose that the information state i is

determined by the context of utterance. To evaluate $\Diamond P$, the semantic engine asks the context to provide its associated state i ; $\Diamond P$ evaluates as true if some worlds in i make P true, false otherwise. To complement this view, one might attempt to extract a viable content by diagonalizing (see Yalcin 2007, pp. 1009-1011). This complex of moves shows that it is not impossible to accept domain semantics and be a propositionalist. What Yalcin does argue (Yalcin 2007, pp. 1011-1012) is that, conditional on the adoption of a domain semantics, a non-propositionalist package is preferable to the propositionalist alternatives. This is partly because it does a better job of capturing some intuitive inferences involving epistemic modals.

Many other authors have developed alternative compositional accounts of the badness of epistemic contradiction that deploy the relational semantics framework Yalcin's argument seeks to overturn (Dorr and Hawthorne, 2013; Ninan, 2016; Silk, 2016a; Stojnić, forthcoming; Mandelkern, m.s.).⁹ Some of these authors are resolutely neutral on questions concerning the content of epistemic modal claims, while for others the defense of a relational semantics is an essential first step towards a defense of contextualism.

3 Theories of assertion and modality.

Let us set aside these content challenges. Instead, let us approach the main question from the point of view of the theory of assertion. There are several ways of conceptualizing the speech act of assertion. In his survey, MacFarlane (2011b) discusses four families of accounts. The speech act of assertion can be understood as the *expression of a belief* (Bach and Harnish, 1979), as a practice regulated by *constitutive norms* (Williamson, 2000), as a practice that generates distinctive *commitments* (Brandom, 1994) and finally as a practice that results in distinctive *updates on context* (Stalnaker, 1978). Within each of these families we can identify many non-equivalent accounts. Normative theories share the idea of characterizing assertion in terms of norms: one may assert C only if some standard involving C is met. But they diverge on the normative standard: it could be *knowledge* (Williamson, 2000);

⁹Mandelkern takes the puzzle of epistemic contradictions one step further by drawing attention to disjunctions of epistemic contradictions (e.g. *Either I might win but I won't or I might not win but I will*). Mandelkern argues that these disjunctions—which sound just as bad as ordinary epistemic contradiction—pose problems for virtually all extant semantic frameworks. (Mandelkern himself advocates a sophisticated relational semantics.)

truth (Weiner, 2005) or *reasonable belief* (Lackey, 2007; Douven, 2006). Similarly, different versions of the commitment view disagree on what sort of commitment assertions involve.

In this section, I showcase a couple of ways in which epistemic modality could constrain normative accounts and commitment accounts. The next section is reserved for a more extensive discussion of context update theories. I do not aspire to provide a complete treatment and I will set aside belief expression theories. In any case: MacFarlane's list was not meant to be an exhaustive catalog of accounts of assertion—as I am sure that this volume will demonstrate.

3.1 Normative Accounts.

My starting point is a brief but suggestive discussion by Swanson (2011, §5). Swanson notes that the knowledge rule (in a slogan: *assert only what you know*) appears to be too demanding when it comes to assertions of epistemic possibility claims. Intuitively, many such assertions are legitimate even if they come from a position of complete ignorance: one needn't know anything in particular to be in a position to assert that it might rain.

Some will no doubt resist this challenge by insisting that asserting $\Diamond P$ does require knowledge that $\Diamond P$. They would also emphasize that this knowledge is very weak: for contextualists, it amounts to the fact that one has to know that the relevant information state does not rule out P . To explain away the 'complete ignorance' intuition, these objectors could note the availability of a restricted interpretation of the last sentence of the previous paragraph. That is, they could grant that one needn't know *any non-modal* facts to be in a position to assert that it might rain.

However, a bit more can be said on behalf of Swanson's observation. The standard motivation for the knowledge account of assertion does not cleanly apply to assertions of epistemic possibility claims. A critical piece of evidence that supports the knowledge account of assertion is that if you assert that there is chocolate in the cupboard, I can challenge your assertion by asking *How do you know?* (Unger, 1975; Williamson, 2000). However, such challenges seem inappropriate when you just make an epistemic possibility claim—a significant disanalogy with non-modal claims.¹⁰ Another piece of evidence in support of the knowledge

¹⁰I do note that this observation fails to extend to many other modal claims in the disputed class. Focus on the critical case of deontic modality: if I say "You should climb that mountain", the *How do you know?* challenge seems much better.

norm is that it helps explain why we cannot assert a categorical proposition only on the basis of statistical evidence (e.g., in lottery cases). There seem to be no parallel cases in the disputed class (but see Moss 2018, p.235 for some case that might put pressure on this). Granted, not all of the arguments for the knowledge account fail. In particular, Moorean sentences involving epistemic modals still sound somewhat bad.

(11) #It might rain but I don't know that it might rain.

But perhaps it is the badness of (11) that is the odd phenomenon out. There might be an independent explanation for why (11) sounds bad, which would allow us to stick to the line that epistemic possibility claims pose problems for the knowledge account of assertion. To begin, *I don't know that it might be that P* is probably bad on its own. A google search for *don't know that it might* returns 0 hits (usual caveats on 'research-by-google-search' apply here). A more acceptable phrasing is

(12) I don't know whether it might be that P.

But even this is often used as a long-winded way of saying *I don't know that P*—and we already know from Moore's paradox that there is something defective about claims of the form *P and I don't know that P*. Moreover, it is (11), and not its *whether*-variant, that figures in the Williamsonian paradigm for the knowledge norm of assertion. Any argument that the argument could run through (12) requires special assumptions about the semantic contribution of *whether*. Overall, Swanson seems right that the generalization of the standard arguments for the knowledge norm to assertions of epistemic modals seems at best spotty.

In addition to this batch of arguments, it is important to notice that the knowledge account is particularly hard to sustain for those contextualist views of epistemic modals according to which the salient body of information is fixed by some group's aggregated body of evidence. In particular, the more flexible one's contextualism, the more demanding the claim that one knows the content of an epistemic modal. This is particularly puzzling for epistemic necessity claims: in the typical case, knowing that Alex is sick seems to be sufficient for knowing that Alex must be sick.¹¹ However, flexible contextualists are committed to extra

¹¹There might, of course, be restrictions on the type of knowledge. But those are not extra pieces of knowledge that are required to the knowledge of the embedded proposition.

requirements beyond these ones: one has to know that the Alex's being sick follows from the relevant body of evidence.

A theoretical package that is friendly to the knowledge account of assertion is the non-propositionalist account recently developed by Moss (2018, m.s.). There is no hope to do justice to Moss's sophisticated theory in this survey (although I will say a bit more about it in §4.3). The key point is that, for Moss, the contents of belief, assertion, and knowledge are probabilistic (they are modeled as sets of probability spaces). So, knowing that it might be raining is not understood as knowing that some salient state is compatible with rain. Instead, it is understood as an appropriate relation to a probabilistic content (the set of probability spaces that treat rain as possible); importantly, this relation is not mediated by a relevant body of evidence. Thus, Moss's theory entails (i) that knowing $\Diamond P$ is an extremely weak type of knowledge and (ii) that knowing P is sufficient to know $\Box P$. Both points can be used to address the earlier objections and overlay the knowledge account of assertion on top of Moss's theory of content.

The dialectic I just outlined involving the knowledge norm extends immediately to normative accounts that demand less than knowledge but still require some positive epistemic status. Consider the normative theory of assertion according to which one must assert only what one reasonably believes. Swanson's point about the knowledge account applies here too. One need not have pertinent reasonable beliefs to be licensed to assert epistemic possibility claims.

It might look however that the truth rule (slogan: *assert only what is true*, see Weiner 2005, MacFarlane 2014, §5.2) has potential to avoid these issues. Applied to epistemic modals, the truth-rule entails that when one asserts $\Diamond P$ in context c , one's assertion is permissible only if $\Diamond P$ is true in c . This approach can be applied to epistemic modal claims, as long as we accept that they have truth-values in context. The truth rule avoids the previous objections, because it does not require that permissible assertions of epistemic possibility claims be backed by any positive epistemic status.

Even so, the previous dialectic might resurface once truth rules are embedded within a more complex account of the normative constraints involved in assertion. Many defenders of the truth rule often combine it with the idea that assertions are also subject to a secondary norm.¹² The primary norm, they claim, is what constitutes the practice of assertion;

¹²This includes both Weiner and MacFarlane, to the extent that he is sympathetic to the normative picture.

the secondary norm is a general constraint on rule-governed behavior. The idea is that whenever one is subject to a rule R, one must act so as to know, or at least be reasonably confident that one's behavior complies with the rule. An analogy will help see the motivation for secondary norms. When I do my taxes, IRS norms only cover the correctness of my taxes. But if we are evaluating the responsibility of my tax-compiling actions, I ought to have good reason to believe that my taxes are correct. Out of the analogy, secondary norms are useful to the truth-norm theorist so she can predict that lucky guesses still count as irresponsible assertions. But, if so, we should expect assertions of epistemic modal claims to be evaluated on the basis of both the primary norm of assertion (truth) and the stronger combination of primary and secondary norm. That will make our earlier concerns resurface at the level of secondary propriety (i.e. at the level of whether the assertion has met the secondary norm).¹³

3.2 Indeterminate normative statuses.

One of the morals of this discussion is that contextualists face a three-way tradeoff between (i) holding on to the intuition that asserting $\Diamond P$ should be easy (ii) adding flexibility to the semantics and (iii) endorsing a norm on which truth is the normative standard for assertion. It is partly on the basis of this conflict that von Fintel & Gillies (2011) advocate *cloudy contextualism*. Cloudy contextualism rejects the presumption that there is a unique proposition that is determined from context when $\Diamond P$ is asserted. Instead, cloudy contextualists maintain that there is a set of propositions that are expressed by an assertion of a bare epistemic modal. An assertion of $\Diamond P$ is permissible if any one of the propositions in the "cloud" is true. In the typical context, the low-hanging fruit for a speaker to latch onto is the solipsistic proposition (*P is not ruled out by my knowledge*). Hearers however can disagree with any one of the propositions in the cloud.

The idea that it might be indeterminate which proposition is expressed by an epistemic modal seems attractive. But the account of assertion at the center of von Fintel & Gillies's proposal appears to be tailored to the specifics of the task of providing a theory of epistemic modal assertions. What we would want instead is a general purpose theory of assertion

¹³Lackey (2007) challenges the very idea of secondary propriety. For reasons I am not able to get into here, I am unconvinced by Lackey's critique. However, it is worth noting that, if Lackey is right, the truth-norm theorists are independently in trouble. So *they* have to push back.

under indeterminacy of content. Moreover, we would want such a theory to result from combining a theory of assertion for determinate contents (call this T) with an interface module that aggregates the verdicts of T on each candidate proposition. The cloudy contextualist's pragmatics violates one plausible desideratum for how to build this interface. This is how such a desideratum might be stated:

If (i) it is indeterminate what proposition an assertion a expresses, but (ii) it is determinate that the eligible propositions are members of the cloud $\mathcal{C} = \{A_1, \dots, A_n\}$ and (iii) the propositions in \mathcal{C} vary with respect to the standard that is tracked by the norms of assertion (e.g. some but not all are true), then it is indeterminate whether a satisfies the norm of assertion.

Here is an example that does not involve modals: suppose that it is indeterminate which of two people is my future self. Suppose also, for the sake of illustration, that this metaphysical indeterminacy is reflected in a content indeterminacy for assertions of *my future self wears glasses*. Suppose finally that only one of my eligible future selves wears glasses. Then, the constraint entails that the normative standing of my assertions of that content is indeterminate.

To safeguard this constraint while keeping relatively close to von Fintel & Gillies' picture, one might consider a close alternative roughly in the vein of a proposal in Cariani and Santorio (2018). The idea is simply to supervaluate on normative statuses across the cloud. The assertion determinately meets (/violates) the norm if for every way of resolving the content indeterminacy, the norm is met (/violated). Consider an assertion of $\Diamond P$ associated with a cloud of propositions \mathcal{C} and suppose that the propositions in \mathcal{C} vary in truth value (with truth being the norm of assertion). Then it is indeterminate whether the speaker is in compliance with the norm. The indeterminacy at the level of which proposition is expressed is transformed into an indeterminacy at the level of the correctness of the assertion (just as required by our constraint). Under this proposal, the status of epistemic modal assertions is short of determinate compliance with the norm of assertion. At the same time, however, these assertions are not determinate violations either. Depending on how the conversational context resolves or constrains the indeterminacy, our judgments about the permissibility of the assertion might fluctuate.

3.3 Commitment Theories.

Commitment theories of assertion (in the style of Brandom 1994) share with normative theories the idea that assertion is to be understood in terms of the normative statuses it gives rise to. However, while normative accounts focus on normative statuses that depend on what is upstream from the assertion, commitment views focus on downstream normative statuses. In particular, Brandom proposes that assertions are, in the first place, commitments to the truth of the assertion. As MacFarlane (2004, 2014) argues, we need to say more than this. Specifically, we must explain what sort of action is demanded by such a commitment and under what circumstances. At a minimum, MacFarlane says, the commitment involves withdrawing the assertion if it is shown to be untrue.

Swanson's observation about the knowledge account suggests a parallel challenge for commitment theories: assertions of $\Diamond P$ signal a lack of commitment—not the presence of one. For example, it is possible to assert $\Diamond P$ with the aim of withdrawing a previous commitment. The invited conclusion seems to be that assertions of epistemic possibility claims are not undertakings of a commitment. This argument breaks down, however, for commitment theories that are based on MacFarlane's minimal proposal. It seems reasonable to maintain that assertions of epistemic modal claims do, indeed, indicate a commitment to withdraw the assertion if its content is shown to be untrue. That might not register as a substantive commitment when we compare it to the much stronger commitment to the truth of P but it is a commitment nonetheless. Moreover, adopting that commitment might have the effect of suspending some other prior commitments.

That said, proponents of commitment views must tread carefully if they go beyond MacFarlane's minimal commitment. This is because the dialectic surrounding Swanson's observation might resurface again if the speaker's commitment involves the requirement to *justify* her assertion should she be challenged (as Swanson 2011, p.262 notes). This is, again, proof that thinking about epistemic modality—and indeed thinking about the statements in the disputed class—can help constrain the range

of viable views of assertion.

4 Context update theories.

Let us circle back to views according to which the speech act of assertion is distinctively associated with an effect on context. For Stalnaker, an assertion is a proposal to add a proposition to the common ground. As we have already seen, this view entails that assertions of epistemic modal claims, if there are any, must eliminate possibilities that are compatible with the common ground. This is fine for propositionalists but problematic for those who think that utterances of epistemic modals do not subtract from the range of possibilities that are compatible with the common ground.¹⁴ In this final section, I want to review three strategies available to the non-propositionalist for dealing with this situation. The first denies that sentences in the disputed class can be vehicles of assertion (§4.1). The second approach accepts that they can sometimes be asserted, but goes in for an account of assertion in which not all assertions update the context in the same way (§4.2). The final approach (§4.3) accepts that epistemic modals can sometimes be asserted and also that assertions are always eliminative but achieves this by adopting a non-standard theory of content.

Ahead of these developments, I will note that the choice between these options (especially between the first and the second) is at least in part a terminological one. I will not attempt to disentangle which bits of the argument are terminological in nature—or even whether it matters that they are terminological as opposed to substantive—but the issue will be evident in a couple places.

4.1 Assertion barriers.

Perhaps, unembedded epistemic modals sentences fail to be eligible vehicles of assertions. Call this the *assertion barrier* view. One very strong way of making this thought more precise is as the claim that no assertion can ever be performed by uttering $\Diamond P$ (and similarly with the other disputed sentences). A weaker assertion barrier view might claim that $\Diamond P$ has two uses: one in which it serves to eliminate possibilities and one in which it does not.

¹⁴Indeed, MacFarlane (2011*b*) uses the discourse effects of epistemic modals as an objection to Stalnaker's account of assertion.

Yalcin (2007) hints at an assertion-barrier thesis. Here is a passage that implies a distinction between "saying $\diamond\varphi$ " and assertions.

To say $\diamond\varphi$ is not to propose to add some informational content, some proposition, to the common ground, as with assertions. Rather, it is to make explicit that φ -possibilities are compatible with the common ground (Yalcin, 2007, p.1010)

Yalcin's full picture is a bit more complex than this passage suggests. Yalcin (2018) distinguishes between illocutionary force and dynamic force. In particular, Yalcin suggests that there are two ways of thinking about assertion that are not equivalent, but nonetheless not in competition as the "correct account of assertion". On the one hand, we can think of assertion as a speech act that is associated with uttering declarative sentences (in the right conditions—whatever those might be); on the other, as a type of act that is associated with a distinctive update on the score of a conversation. The natural interpretation of the passage I quoted is that "sayings" of $\diamond\varphi$ are not assertions in the "dynamic force" sense, but that leaves it open that might be assertions in some other sense.

Perhaps unsurprisingly, Stalnaker himself develops an assertion-barrier proposal (Stalnaker, 2016, ch. 6). Stalnaker argues that we ought to distinguish a few different ways of updating a context. In addition to the standard assertive update—an intervention on the context prior to the assertion—we could have *prospective* updates. The effect of some claims is to propose that we move to a new context in a way that is constrained by the truth-conditions of what is being claimed. For instance, saying $\diamond P$ can be construed as a proposal to move to a posterior context in which P is compatible with the common ground. Stalnaker models his account of epistemic modal claims on his interpretation of the language game in Lewis's "A Problem about Permission" (Lewis, 1975). Much like epistemic modal claims can expand the possibilities under consideration, permission claims can have the effect of expanding the range of permissible actions for salient agents (Lewis calls these *spheres of permissibility*). For example, when I say to my daughter *you may have one square of chocolate* I say something that has certain truth-conditions. Perhaps, it is true in context c if and only if my daughter's sphere of permissibility in c includes some possibilities in which she eats one square of chocolate. But, Stalnaker suggests, the typical point of making this claim in a context c^- is not to add the proposition associated with

those truth-conditions to the common ground. It is instead to transition to a new context c^+ in which those truth-conditions are satisfied.

There are some important difficulties for the assertion barrier view. Swanson (2011, §5) argues—persuasively in my view—that simple versions of the assertion barrier view cannot work, because they do not make sense of seemingly assertive utterances of mixed claims. Consider, for instance:

(13) There might be chocolate and there is wine.

Is (13) a possible vehicle of assertion? Prior to any argument, we might entertain three possibilities:

option 1: all of (13) is asserted.

option 2: no part of (13) is asserted.

option 3: the second but not the first conjunct of (13) is asserted.

The first option is not compatible with the assertion barrier view: surely if you can assert a conjunction of the form $\Diamond P \ \& \ _$, then you can assert the unembedded $\Diamond P$. The second option makes a mystery of what should be an obvious fact: it is possible, maybe even typical, to utter (13), and have the content of the second conjunct be treated as asserted content. (We update the common ground with its content, and we hold speakers responsible for it in whatever relevant ways.) The third option either does not make sense or, at the very least, requires quite the revolution in the theory of speech acts (to allow us to categorize an utterance of s as a complex act consisting of a non-assertion followed by an assertion).

Dever (ms.) also notes that the assertion barrier view obscures the fact that epistemic modal sentences can be part of ordinary information exchanges. When I say to you *Jonathan will probably come*, I am calling for you to revise your information state (if it currently lacks the relevant structure). Although we might, once again, be in terminological dispute territory, Dever is pointing that it is intuitive to group together under the label "assertion" all attempts to affect the information states of conversational participants.

4.2 Variable effect theories.

Another possible reaction is that Stalnaker's truism goes wrong because it presupposes that the category of assertion is unified by the particular

type of update it is associated with. Maybe we should instead adopt a picture on which assertions are associated with a variety of possible effects.

This response is particularly natural if we think from the perspective of update semantics (Veltman, 1996). The driving slogan of update semantics (and of dynamic semantics more generally) is that meanings are context-change potentials and not contributions to truth-conditions. For our purposes, however, the interesting feature of these systems is that context-change potential is not a uniform notion. Not every declarative sentence updates the context in the same way. Instead, the updates corresponding to various sentences, even in a simple propositional modal language, are defined by distinguishing a variety of possible cases.

Suppose that we only think of context as a set of possibilities i . Let us use the notation $i[P]$ to denote the update of state i with the information associated with sentence P . Here is the standard update semantics for a propositional modal language (this is Veltman's system with a standard tweak on conjunction due to Heim 1983):

If P is atomic, $i[P] = i \cap P$.

$i[\neg P] = i - i[P]$

$i[P \ \& \ Q] = i[P][Q]$

$i[\Diamond P] = i$, if $i \cap P \neq \emptyset$, else $i[\Diamond P] = \emptyset$.

Atomic sentences perform the standard eliminative update. In fact, the full system guarantees that updates with any Boolean combination of atomic sentences is intersective. The update corresponding to $\Diamond P$ is what Veltman calls a *test* on the state (in this case, since we are momentarily identifying states and contexts, it is a test on the context).

One potential problem with Veltman's original approach is that it does not seem to have any model for the natural thought that epistemic possibility claims can add to the set of possibilities we treat as live, or take seriously. Many theorists have found it attractive to think of some epistemic modal utterances as doing just that (von Stechow and Gillies 2007 §5, MacFarlane 2011b, Willer 2013, Lennertz 2014). Instead, on the test conception, $\Diamond P$ either returns the initial state or crashes the state into the contradictory state.

A particularly nice model of this phenomenon is developed by Willer (2013). Willer's key idea is to view the standard update model as part of

a two-step theory of update. Instead of thinking of the informational component of context as a single set of possibilities, Willer proposes thinking of it as a set of sets of possibilities. Setting aside the general formal representation, Willer's central idea can be illustrated with an example. Suppose that the only relevant event is a past flip of a coin. The conversational record is completely silent about this flip— nothing has been said about the coin. There are three (qualitative) information states we might be in with respect to this subject matter: *certain of heads* (**H**), *uncertain* (**U**), and *certain of tails* (**T**). (Each of these is modeled as a set of possibilities.) A context that has no commitment whatsoever could be represented as the set of all such states: {**H**, **U**, **T**}. (Note that we now have a distinction between context and states.) A possibility is *live* in a context *c* if it is compatible with every state in *c*. It is *compatible* with *c* if it is compatible with some state in *c*. In our example, the possibility of tails is not live in the initial context (even if it is compatible with it) because there is a state that excludes it. The effect of asserting any sentence in a context is to update each state in the context (this is where the standard dynamic theory gets used). If a state in *c* gets updated to contradiction, it is eliminated from the context. So, asserting *the coin might land tails* in the context represented by {**H**, **U**, **T**} has the effect of eliminating **H** (because the Veltman-style test on this state with that sentence crashes) and returning {**U**, **T**}. In the new state the coin's landing tails is a live possibility.

One interesting feature of the non-uniform approach is that it is easy to write new entries for different types of expressions. There are well-developed proposals for adding counterfactuals (Veltman, 2005), probability operators (Yalcin, 2012), and deontic modals (Cariani, 2016; Silk, 2016*b*; Starr, 2016; Willer, 2016) to the update semantics framework.

These dynamic approaches embody the idea that there are many ways of updating a context. But does this fact show that there are many possible effects for *assertions*? Well, it does not *follow*, for sure. But there certainly is an attractive way to think about this. Following a lead by Dever (ms.), we can say that every assertion is associated with an update, but not every assertion updates the context in the same way. Non-modal sentences update the context in the familiar, eliminative ways. Epistemic modal claims have their own distinctive way of updating the context. Deontic claims might have yet another.

This move is not available to all non-propositionalists. In section §1, I noted a fault-line among non-propositional views. Some demand that the disputed sentences, e.g. $\diamond P$, have no propositional uses. Others

suggest that *sometimes* these sentences can be used non-propositionally. The update semantics approaches I have discussed in this section (along with my overlays about assertion) seem only compatible with the first approach. Perhaps, however, the same move can be replicated by the *dynamic pragmatist* in the style of Stalnaker's (2016) view that I described in the last section.

It is natural to worry that a big question looms for this view: what is the more general conception of assertion? What unifies an assertion of P , with an assertion of $\Diamond P$? Here it is easy to take a few wrong turns. We might start by identifying assertions with updates on context. The immediate difficulty is that, when we start modeling a broader fragment of language, there will be different updates that have nothing to do with assertion (e.g. questions, commands). A preliminary proposal to fix this problem is by Dever (ms.), who suggests that we might think of assertions as updates on the *informational* coordinate of context. So, we might classify commands as non-assertions because they intervene on an agent's sphere of permissibility. One problem with this proposal is that some of our motivation for classifying utterances of the disputed modal sentences as assertions extends to deontic claims. The problem is that we might want to give a semantics for these claims that does not only involve updates on the informational coordinate of context. Additionally, there is a converse problem: there are speech acts that are not assertions but nonetheless affect the informational coordinate of context: one of the key points in (Stalnaker, 1978, p. 87) is that we cannot identify assertion with eliminative update because suppositions have the same effect on context as assertions.

While I do not quite see how to come up with a general characterization of assertion in terms of update, I am not sure that such a characterization is necessary. Without going all in for quietism, it seems reasonable to hope that a serviceable, if not completely correct, characterization might come from the theory of illocutionary forces. And even if that project failed, it is not clear what indispensable philosophical value would be provided to the theory of meaning from having such a characterization.

4.3 Uniform effect, complex contents.

According to the last option I want to consider, assertions of non-modal sentences and assertions of epistemic sentences all have the same kind of content, but not propositional content. Moss (2018) proposes that epistemic claims generally be understood as denoting sets of probability

spaces. A probability space is a triple $\langle W, S, P \rangle$ in which W is a set of worlds, S is a set of subsets of W , and P is a probability function defined over S .

To a first approximation, we might say: the content of *it is .6 likely that there is chocolate* is the set of probability spaces that assign .6 to the proposition that there is chocolate (in the relevant place). The content of *there might be chocolate* is the set of probability spaces that are compatible with the proposition that there is chocolate.¹⁵ Crucially, Moss proposes that non-modal sentences *also* have probabilistic contents: the content of *there is chocolate* is the set of probability spaces that are certain that there is chocolate.¹⁶

This picture invites a natural lifting of the standard Stalnakerian story about the effects of assertion. First, treat the common ground of a conversation as determining a set of probability spaces (Rothschild 2012, Yalcin 2012, Swanson 2016, Moss 2018, §2.3). Next, because of the unity in content, treat update as the addition of a constraint to the probabilistic common ground.

One possible shortcoming of this approach is that, as presently constructed, it singles out for special treatment the fragment of the disputed class that involves epistemic language. For instance, there is no parallel provision for deontic statements. Because the story thrives on its ability to treat non-modal and modal information uniformly, it is also not possible to adopt the piece-meal approach we followed in the case of update semantics. There are, I think, three options for development in this direction. First, we could claim that some or all of the non-epistemic statements in the disputed class are analyzed along contextualist lines. The problem is that much of the motivation for rejecting propositionalism carries over at least to the deontic case (for a non-exhaustive sample, see Ninan 2005; MacFarlane 2014, ch. 11; Charlow 2016; Starr 2016). At the opposite end of the spectrum, someone who is in broad agreement with Moss's framework might appeal to a more general object than sets of probability spaces. In addition to the technical task of figuring out what objects might work, this strategy would also compromise some of the elegance of the framework. The intermediate course would be to try to tell a story on which the contents of deontic modal claims are also sets of probability spaces.

¹⁵Compatibility is spelled out in a way that is subtle and involves a touch of context-dependence via partitions. See the semantic proposal in the appendix to Moss's book and in Moss (2015).

¹⁶This hypothesis is combined with an explanation for why we don't often take ourselves to assert such extreme contents when we assert non-modal sentences.

It is not hard to see how one might get started on this project (but it is hard to tell how much danger lies ahead). Deontic modals are already arguably sensitive to probabilistic information (Cariani, 2016). So, consider a context in which various bets are being placed on a coin flip and the odds of the bet are known to be even and the addressee cares about winning. In such a context,

(14) You should bet that the coin will land heads.

denotes the set of probability spaces that assigns more than .5 to the proposition that the coin will land heads. If we were in a context in which the addressee intended to lose, the same sentence might instead denote the set of probability spaces that assigns less than .5 to that same proposition. (Moss's framework is designed to be able to lift and exploit very standard ideas about context sensitivity.) If this approach could be generalized for all the deontic claims in the disputed class, it would preserve the unity and explanatory power of the framework.

Acknowledgments

Thanks to Nate Charlow, Simon Goldstein, Sarah Moss, Paolo Santorio, Eric Swanson and Malte Willer for conversations and exchanges. Thanks to Sandy Goldberg for feedback on the paper and for the invitation to contribute.

Bibliography

- BACH, Kent and HARNISH, Robert M. (1979), *Communication and Speech Acts* (Harvard University Press).
- BRANDOM, Robert (1994), *Making it explicit: Reasoning, representing, and discursive commitment* (Harvard University Press).
- BRAUN, David (2012), 'An invariantist theory of 'might' might be right', *Linguistics and philosophy*, 35(6), 461.
- CARIANI, Fabrizio (2013), 'Epistemic and Deontic *Should*', *Thought*, 2(1), 73–84.
- (2016), 'Deontic Modals and Probabilities: One Theory to Rule Them All?', in N. Charlow and M. Chrisman (eds.), *Deontic Modals*, 11–46 (Oxford University Press).
- CARIANI, Fabrizio and SANTORIO, Paolo (2018), 'Will Done Better: Selection Semantics, Future Credence, and Indeterminacy', *Mind*, 127, 129–165.
- CHARLOW, Nate (2016), 'Decision Theory: Yes! Truth Conditions: No!', in Nate Charlow and Matthew Chrisman (eds.), *Deontic Modality* (Oxford University Press).
- DEROSE, Keith (1991), 'Epistemic Possibilities', *The Philosophical Review*, 100, 581–605.
- DEVER, Joshua (ms.), 'Assertion and the Future', University of Texas, Austin.
- DORR, Cian and HAWTHORNE, John (2013), 'Embedding epistemic modals', *Mind*, 122(488), 867–913.
- DOUVEN, Igor (2006), 'Assertion, Knowledge, and Rational Credibility', *Philosophical Review*, 115(4), 449–485.
- DOWELL, Janice J. L. (2011), 'A Flexible Contextualist Account of Epistemic Modals', *Philosophers' Imprint*, 11, 1–25.
- EGAN, Andy, HAWTHORNE, John, and WEATHERSON, Brian (2005), 'Epistemic Modals in Context', in G. Preyer and P. Peter (eds.), *Contextualism in Philosophy* (Oxford University Press).

- VON FINTEL, Kai and GILLIES, Thony (2007), 'An Opinionated Guide to Epistemic Modality', in Tamar Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* 2, 32–62 (Oxford University Press).
- (2008), 'CIA Leaks', *The Philosophical Review*, 117, 77–98.
- (2011), 'Might' made right', in Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*, 108–130 (Oxford University Press).
- HEIM, Irene (1983), 'On the projection problem for presuppositions', in Flickinger D. et al. (ed.), *Proceedings of the Second West Coast Conference in Linguistics*, 114–125 (Stanford University Press).
- KHOO, Justin (2015), 'Modal Disagreements', *Inquiry : An Interdisciplinary Journal of Philosophy*, (5), 1–24.
- KHOO, Justin and KNOBE, Joshua (2016), 'Moral disagreement and moral semantics', *Noûs*.
- KHOO, Justin and PHILLIPS, Jonathan (ms.), 'New horizons for a theory of epistemic modals', MIT and Harvard University.
- KNOBE, Joshua and YALCIN, Seth (2014), 'Epistemic modals and context: Experimental data', *Semantics and Pragmatics*, 7(10), 1–21.
- KRATZER, Angelika (1977), 'What 'Must' and 'Can' Must and Can Mean', *Linguistics and Philosophy*, 1(3), 337–355.
- (1981), 'The Notional Category of Modality', in B. Partee and P. Portner (eds.), *Formal Semantics: the Essential Readings* (Blackwell).
- (1991), 'Modality', in A. von Stechow & D. Wunderlich (ed.), *Semantics: An International Handbook of Contemporary Research* (De Gruyter).
- (2012), *Modals and Conditionals* (Oxford University Press).
- LACKEY, Jennifer (2007), 'Norms of Assertion', *Noûs*, 41(4), 594–626.
- LENNERTZ, Ben (2014), 'Taking *might*-communication seriously', *Analytic Philosophy*, 55(2), 176—198.
- LEWIS, David (1975), 'A Problem about Permission', in *Papers in Ethics and Social Philosophy* (Cambridge University Press).
- MACFARLANE, John (2011*a*), 'Epistemic Modals are Assessment Sensitive', in Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*, 144–178 (Oxford University Press).

- (2011*b*), ‘What Is Assertion?’, in Jessica Brown and Herman Cappelen (eds.), *Assertion* (Oxford University Press).
- (2014), *Assessment Sensitivity* (Oxford University Press).
- MANDELKERN, Matt (m.s.), ‘Bounded Modality’, m/anuscript, University of Oxford.
- MANDELKERN, Matthew, SCHULTHEIS, Ginger, and BOYLAN, David (2017), ‘Agentive modals’, *Philosophical Review*, 126(3), 301–343.
- MOSS, Sarah (2015), ‘On the semantics and pragmatics of epistemic vocabulary’, *Semantics and Pragmatics*, 8(5), 1–81.
- (2018), *Probabilistic Knowledge* (Oxford University Press).
- (m.s.), ‘Full Belief and Loose Speech’, manuscript, University of Michigan.
- NINAN, Dilip (2005), ‘Two Puzzles About Deontic Necessity’, in Gajewski J., Hacquard V., Nickel B., and Yalcin S. (eds.), *New Work on Modality* (MIT Working Papers in Linguistics).
- (2016), ‘Relational semantics and domain semantics for epistemic modals’, *Journal of Philosophical Logic*, 1–16.
- ROTHSCHILD, Daniel (2012), ‘Expressing Credences’, *Proceedings of the Aristotelian Society*, 112, 99–114.
- (2013), ‘Do indicative conditionals express propositions?’, *Noûs*, 47(1), 49–68.
- SILK, Alex (2016*a*), *Discourse Contextualism: A framework for contextualist semantics and pragmatics* (Oxford University Press).
- (2016*b*), ‘Update Semantics for Weak Necessity Modals’, in *Deontic Logic and Normative Systems: 13th International Conference, DEON 2016*, 237–255.
- STALNAKER, Robert (1978), ‘Assertion’, in *Context and Content* (Oxford University Press).
- (2011), ‘Conditional Propositions and Conditional Assertions’, in Andy Egan and B. Weatherson (eds.), *Epistemic Modality* (Oxford University Press).
- (2016), *Context* (Oxford University Press).

- STARR, William B. (2016), 'Dynamic Expressivism About Deontic Modality', in Nate Charlow Matthew Chrisman (ed.), *Deontic Modality*, 355–394 (Oxford University Press).
- STEPHENSON, Tamina (2007), 'Judge Dependence, Epistemic Modals, and Predicates of Personal Taste', *Linguistics and Philosophy*, 30(4), 487–525.
- STOJNIĆ, Una (forthcoming), 'Content in a Dynamic Context', *Noûs*.
- SWANSON, Eric (2011), 'How Not to Theorize About the Language of Subjective Uncertainty', in Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*, 249–269 (Oxford University Press).
- (2016), 'The Application of Constraint Semantics to the Language of Subjective Uncertainty', *Journal of Philosophical Logic*, 45(2), 121–146.
- TETLOCK, Philip E and GARDNER, Dan (2015), *Superforecasting: The art and science of prediction* (Random House).
- UNGER, Peter K. (1975), *Ignorance: A Case for Scepticism* (Oxford University Press).
- VELTMAN, Frank (1996), 'Defaults in Update Semantics', *Journal of Philosophical Logic*, 25, 221–261.
- (2005), 'Making counterfactual assumptions', *Journal of Semantics*, 22(2), 159–180.
- WEINER, Matt (2005), 'Must We Know What We Say?', *The Philosophical Review*, 114, 227–251.
- WILLER, Malte (2013), 'Dynamics of Epistemic Modality', *The Philosophical Review*, 122, 45–92.
- (2016), 'Dynamic Foundations for Deontic Logic', in Nate Charlow and Matthew Chrisman (eds.), *Deontic Modals*, 324–354 (Oxford University Press).
- WILLIAMSON, Timothy (2000), *Knowledge and its Limits* (Oxford University Press).
- YALCIN, Seth (2007), 'Epistemic Modals', *Mind*, 116(4), 983–1027.
- (2011), 'Nonfactualism About Epistemic Modality', in Egan A. and Weatherson B. (eds.), *Epistemic Modality* (Oxford University Press).

- (2012), ‘Context Probabilism’, in M. Aloni et. al. (ed.), *Logic, Language, and Meaning: Proceedings of the 18th Amsterdam Colloquium. Lecture Notes in Computer Science, Vol. 7218*, 12–21 (Springer).
- (2016), ‘Modalities of normality’, in Nate Charlow and Matthew Chrisman (eds.), *Deontic modals*, 230–255 (Oxford University Press).
- (2018), ‘Expressivism by Force’, in Daniel Fogal, Daniel Harris, and Matt Moss (eds.), *New Work on Speech Acts* (Oxford University Press).
- YANOVICH, Igor (2013), ‘Standard contextualism strikes back’, *Journal of Semantics*, 31(1), 67–114.