How To Do Things With Modals*

Matthew Mandelkern

October 4, 2018

To appear in *Mind & Language*

Abstract

In a brief discussion of epistemic modals, Wittgenstein (1953, II.x.110) warns against 'regard[ing] a hesitant assertion as an assertion of hesitancy'. A modal claim like 'It might be raining', the thought goes, should not be regarded as an assertion of the speaker's uncertainty as to whether or not it is raining, but rather as something quite different in kind: a proposal to treat the possibility of rain as live. Wittgenstein's admonition has, in recent years, been at the heart of arguments that, in order to make sense of the dynamics of epistemic modal claims, we must reject the contextualist framework for analyzing communication—a framework on which assertions of epistemic modal claims, like all assertions, convey information. In this paper, I argue that, on the contrary, taking Wittgenstein's admonition seriously does not require abandoning the contextualist framework: we can capture the fundamental dynamics of epistemic modality within the contextualist framework, provided we take the assertoric content of unembedded epistemic modal claims to be determined by the prospective common attitudes of the conversants in question.

Keywords: epistemic modals; semantics; pragmatics; models of communication; philosophy of language

1 Introduction

In a brief but trenchant discussion of epistemic modals, Wittgenstein warns against 'regard[ing] a hesitant assertion as an assertion of hesitancy'.[1] An epistemic modal claim like 'It might be raining', the thought goes, should not be regarded as an assertion of the speaker's uncertainty as to whether or not it is raining, but rather as something quite different in kind: a proposal to treat the possibility of rain as live.

Wittgenstein's admonition provides a helpful lens for viewing the subsequent debate about the meaning of epistemic modal claims, and the system of communication that these claims fit into. On the one hand, the standard account treats 'It might be raining' as expressing an ordinary piece of information, namely that the proposition that it's raining is compatible with some contextually salient piece of evidence. On the face of it, the standard account thus treats modal claims as assertions of hesitancy, in Wittgenstein's phrase; because of this, it has been taken to task for failing to make sense of how speakers do what they do with modal

*Special thanks to Kai von Fintel, Justin Khoo, Robert Stalnaker, Stephen Yablo, and two anonymous referees for this journal for very helpful and extensive feedback on this paper. Thanks also to audiences at MIT, Rutgers, and UMass Amherst, and to David Boylan, Fabrizio Cariani, Daniel Drucker, Andy Egan, Peter van Elswyk, Cosmo Grant, Irene Heim, Matthias Jenny, Brendan de Kenessey, Rose Lenehan, Pengbo Liu, Milo Phillips-Brown, Agustín Rayo, Kevin Richardson, Bernhard Salow, Paolo Santorio, Ginger Schultheis, Kieran Setiya, Bradford Skow, Shane Steinert-Threlkeld, and Brian Weatherson for very helpful discussion.

[1] 'Betrachte nicht die zaghafte Behauptung als Behauptung der Zaghaftigkeit' (Wittgenstein, 1953, II.x.110).

claims.[2] On the other hand, an array of heterodox accounts have taken Wittgenstein's admonition to heart, re-engineering not only the standard theory of the meaning of epistemic modals, but also the contextualist theory of communication—on which communication is the transfer of information—in order to make sense of the dynamics of modal language.

In this paper I will show that there is a way to walk a line between these two options. It is indeed wrong to view a 'might'-claim as an assertion of uncertainty, as a flat-footed interpretation of the standard theory would have it. This approach fails to capture a fundamental observation about the conversational dynamics of epistemic modality, namely, that an assertion of ⌜Might p⌝ is in general a proposal to make p compatible with the *common ground*:[3] a proposal which, if accepted, ensures that p is compatible with the common ground, and, if rejected, ensures that it is not.[4] I call this observation about the fundamental dynamics of epistemic modality *the guiding observation*.

It is overhasty, however, to conclude—as many in the recent literature have—that there is *no* content that we can assign to 'might'-claims which will capture the guiding observation. I show, to the contrary, that we can capture the guiding observation, and therefore the fundamental dynamics of epistemic modality, within the contextualist framework. We can do so by assigning an assertion of ⌜Might p⌝ a content which is *about the conversation's common ground* itself. In particular, we maintain that ⌜Might p⌝ says that p is compatible with the common ground *after the claim in question has been made and negotiated*. I show that this theory of the assertoric content of modal claims (which I call *prospective contextualism*), together with the contextualist theory of communication, guarantees that an assertion of ⌜Might p⌝ just is a proposal to make p compatible with the common ground, thus capturing the guiding observation.

Prospective contextualism thus shows that we can make sense of the fundamental dynamics of epistemic modality within the contextualist framework. The main goal of this paper is to show just this: that, *pace* much recent work, a contextualist approach to epistemic modality cannot be rejected on the putative basis that it cannot capture the guiding observation. Whether prospective contextualism is the *right* way to model epistemic modality is another, much broader, question, one which I will not attempt to answer decisively here. I will try to say enough, however, to show that prospective contextualism should be taken seriously as a contender theory of epistemic modality. In particular, I will argue that some natural objections to prospective contextualism can be answered, by bringing out a parallel between the structure of modal claims, on my theory, and the structure of *performative* uses of language; and that prospective contextualism has some at least *prima facie* advantages over some competitor views, to a degree which makes it worth serious further exploration. I conclude by suggesting that the present approach to the dynamics of modality can be extended to give a contextualist model for a wide variety of other speech acts which at first blush appear to go beyond the bounds of contextualism.

[2] I will often use 'modal claim' for 'epistemic modal claim'; my focus in this paper will be exclusively on modal claims interpreted in a broadly epistemic sense.

[3] The common ground is the set of propositions commonly accepted in the conversation. 'In general' is included here to cover purely descriptive uses of modal claims of the kind discussed in §5.2 below.

[4] I use roman letters to stand for sentences and italic letters to stand for the corresponding propositions, with $\bar{p}$ standing for p's complement; I leave implicit relativization to contexts.

2 Contextualism and its discontents

I begin by sketching the contextualist framework, and the challenge that epistemic modals pose to it.

2.1 Contextualism

'Contextualist' is sometimes used to describe a theory of how a particular term works; for instance, any plausible theory of the meaning of 'I' will be contextualist, insofar as it will say that the denotation of 'I' varies with the context (it always denotes the speaker *of the context*). I will use 'contextualist' here in a more general way, to describe a wholesale model of communication and division of labor between semantics and pragmatics.[5] In the contextualist framework, semantics is responsible for assigning contents to well-formed declarative sentences, relative to contexts;[6] this process will be compositional, and the output of our linguistic cognitive system. The contents which are the output of this process of interpretation will be units of information, or *propositions*, which we model as sets of possible worlds—intuitively, the set of all and only those worlds which verify the proposition in question.[7]

Pragmatics, by contrast, concerns the application of domain-general reasoning to linguistic communication. Most importantly for our purposes, the propositions which are the output of the interpretation of sentences at contexts will interface with a pragmatic system, which does things with these contents. In particular, contextualism commits to a simple theory of the basic thing speakers do with these contents: they assert them in order to coordinate on their information. To model this, contextualism tracks the common commitments in a conversation at a given time—the conversation's *common ground*—and then says that an assertion of a declarative sentence is a proposal to add the proposition expressed by the sentence at that context to the common ground.[8] An assertion, in other words, is just a proposal for the interlocutors to come to commonly believe the content asserted.

I make all this more formal in Appendix A, but this informal exposition suffices for present purposes. Contextualism provides an elegant model of the division of labor between semantics and pragmatics, and an elegant model of the dynamics of conversation. And, since it views conversation as the mutual exchange of information, contextualism provides a fully explicit model of the evolution of conversation, at least once we couple it with an independently motivated theory of how agents assimilate new information: conversation just is belief change, and so the evolution of conversation just is the evolution of (common) beliefs.

[5] My theory of modals is contextualist in both senses. See especially Stalnaker 1970, 1978, Karttunen 1974a.

[6] More precisely, to logical forms. A *context* provides information about anaphora resolution, etc.; standardly, we can think of a context just as a variable assignment.

[7] This latter modeling choice is strictly independent of the contextualist framework; all that is really crucial is that propositions can be regarded as units of information.

[8] More explicitly, the common ground is the set of propositions accepted by all conversants, accepted to be accepted, and so on, where acceptance is an attitude like belief or knowledge. I will often conflate acceptance and belief in what follows for the sake of concreteness. Note that the common ground is defined for any number of conversants, including just one (in that case the common ground will be identical to the transitive closure of the individual's state of acceptance). A natural objection to the view I will give below is that it cannot make sense of what we are doing when we use modal language when thinking or talking to ourselves, since it references the common ground; but this objection is premised on the mistaken assumption that the common ground is an essentially interpersonal notion.

2.2 Discontents

Despite its attractions, however, it is easy to get in a mood in which contextualism looks too simple. Conversation is replete with uses of language which, at first blush, go far beyond the simple exchange of information. In the rest of this paper I will make a case study of one use of language which poses this problem in a particularly distinct and challenging way: epistemic modal claims.

What do speakers do with epistemic modal claims? Focus for now on 'might'-claims, like 'It might be raining' or 'John might be in his office'. Intuitively, a 'might'-claim is a way of proposing that some possibility should be taken seriously in inquiry: e.g. 'It might be raining' says 'Let's take seriously the possibility that it's raining'. In the contextualist framework, we can, at least as a first pass, make this precise by saying that a claim of ⌜Might p⌝ is a proposal to make p compatible with the common ground, and to make this fact itself common ground.[9] Thus an assertion of 'John might be in his office' is a proposal to make it compatible with the common ground that John's in his office, and to make this fact itself common ground.

This description of the dynamics of modality, which I will refer to in what follows as *the guiding observation*, makes sense of intuitions about what speakers are trying to do when they assert 'might'-claims: namely, to ensure that we treat its prejacent (the proposition it embeds) as a live possibility in our investigation of what the world is like. It also makes sense of intuitions about what speakers are doing when they agree or disagree about 'might'-claims. If you say that John might be in his office and I object that he can't be, what we are arguing about is whether to treat as live the possibility that John is in his office: you are saying we should, I am saying we should not. If we treat 'must' as the dual of 'might', then the guiding observation also makes sense of the dynamics of 'must'-claims, predicting that an assertion of ⌜Must p⌝ is a proposal to ensure that p is *entailed* by the common ground. This, again, conforms to intuitions and makes sense of agreement and disagreement about 'must'-claims, though this will not be our focus here.

Now that we have the guiding observation in hand, we can formulate the challenge epistemic modals pose to contextualism. The issue is that an assertion, on the contextualist theory, is a proposal to accept some piece of information. But, if the guiding observation is right, a 'might'-claim does not look like a proposal to accept a piece of information, but rather a proposal to make some piece of information *compatible* with the common ground: in Wittgenstein's phrase, a hesitant assertion, not an assertion of hesitancy.

To make this worry more pressing, consider the standard theory of the meaning of epistemic modal claims. On this theory, ⌜Might p⌝ means that p is compatible with a body of evidence or attitude state made salient by the context of utterance.[10] This is modeled by associating 'might' with a contextually

[9] See Stalnaker 1970 and many since for this idea. The latter part of this is non-trivial, since something can be compatible with the common ground without this fact itself being common ground. This is one way, in the present framework, that we make sense of the way in which a 'might'-claim is used not only to make its prejacent compatible with the common ground, but also to bring this fact to salience. As an anonymous referee for this journal helpfully points out, there is no doubt more that is required for an adequate model of the attention-getting role of 'might'-claims, which is what makes this gloss on this intuition only *prima facie*. Such a model will presumably involve something more fine-grained than the Hintikkan approach I am assuming here: e.g. QUDs/partitions (Roberts, 2012; Yablo, 2014; Yalcin, 2012), inquisitive contents (Ciardelli et al., 2009), or covers of attitude states (Willer, 2013). I will stick with a simple Hintikkan model of attitudes here for the sake of simplicity; while the resulting model of the attention-getting role of 'might'-claims is certainly partial, it nevertheless does go some way, I think, towards accounting for the role of epistemic modal claims in coordinating common information.

[10] E.g. Moore 1962, Hacking 1967, Teller 1972, Kratzer 1977, 1981, and many since.

provided accessibility relation between worlds, on which w accesses w' just in case w' is compatible with the contextually salient evidence or attitudes at w. Then ⌜Might p⌝ is true at context c and world w just in case p is true at some world accessible from w.

At first blush this looks reasonable enough, but, when we try to say more about how the accessibility relation in question is actually determined at given contexts, it is hard to see how to bring this theory in line with the guiding observation.[11] A natural first thought would be to take the accessibility relation to be determined by the knowledge of the speaker. Then ⌜Might p⌝ will be equivalent to ⌜For all I know, p⌝. But this clearly doesn't conform to the guiding observation: a claim with the form ⌜For all I know, p⌝ is not a *proposal* to treat p as compatible with the common ground, but a statement that p *already is* compatible with what you know (and, therefore, with the common ground). This model thus fails to account for how speakers negotiate about modal claims. If you say 'John might be in his office' and I reply, 'No, he's in England', I am not disputing that it is compatible with *your* antecedent knowledge that John is in his office (after all, you should know); instead, I'm disputing whether this is a good thing for us to leave open in *our inquiry*. Likewise, if I agree, I am not agreeing that it is compatible with your knowledge that John is in his office (again, you should know), but rather that it's a good thing to leave open in inquiry.

A natural second attempt would take the accessibility relation to be determined, not by the speaker's knowledge, but rather the knowledge of the *group* of conversants: the knowledge state that would result if all the individuals in the group *pooled their knowledge*.[12] But this approach still fails to conform to the guiding observation. Suppose we're in a big group. You see John's office light on. I ask you where John is. You can say, 'He might be in his office': you are proposing to treat as live the possibility that John is in his office. You are *not* asserting that, if everyone in the group pooled their knowledge, it would remain compatible with our knowledge that John is in his office. Indeed, you may think that, if we pooled all our knowledge, we might well discover that John was not in his office. More succinctly: in cases like this, 'John might be in his office' is intuitively assertable, whereas 'The strongest thing we together know leaves open that John is in his office' is, intuitively, not. These thus cannot have the same content in this context.

These two options do not exhaust the space of possibilities—a point I'll make crucial use of in the next section—but it is tempting at this point to echo Wittgenstein: the standard theory, and the contextualist framework in which it is embedded, makes the mistake of trying to analyze modal assertions—a special kind of speech act whose aim is to coordinate on which possibilities to treat as live—as assertions of hesitancy. It is a mistake to analyze epistemic modals as any kind of assertion at all, the thought goes; modal claims do not aim to transfer information, but rather constitute a different speech act, something which takes us altogether out of the bounds of the contextualist framework. This pessimistic line is implicit or explicit in much of the recent work on epistemic modals; here is a characteristic statement:[13]

[11] See Price 1983, MacFarlane 2011, von Fintel and Gillies 2011, a.o.

[12] In standard terminology, the group's *distributed* knowledge: i.e., where $K_{S,w}$ represents the set of worlds compatible with S's knowledge in w, and I is a group, $\bigcap_{i \in I} K_{i,w}$. A weaker approach, which takes the accessibility relation to be determined by the *weakest* thing known—i.e., by $\bigcup_{i \in I} K_{i,w}$—would fail to make sense of disagreement.

[13] I should note that Rothschild (2013) is not clear about whether he endorses this view. Another characteristic statement comes in Swanson (2011, p. 251): 'Construing subjective uncertainty about whether φ in terms of near certainty about some other proposition seems wrongheaded. But unless the truth conditional theorist can find such propositions, there is no reason to suppose that an assertion of a doxastically hedged sentence will inculcate the appropriate partial belief in the addressee. And the project of

> In general, there is no single proposition one can accept such that one doesn't rule out p if and only if one accepts that proposition (p itself is too strong). So if our language provides the resources to simply express that we don't rule out p, then it seems we will not be able to understand that bit of language if we confine ourselves to a framework where sentences express propositions. (Rothschild, 2013, p. 50)

3 Prospective contextualism

But this pessimistic conclusion is too fast. *Pace* this pessimistic conclusion, there *is* a proposition we can assign as the content of ⌜Might p⌝ which captures the guiding observation. The basic idea is to preserve the structure of the standard theory, but to have the accessibility relation determined *by the common ground itself*. In this section I show that, if we do this in the right way, then, thanks to certain properties of the logic of common ground, an assertion of ⌜Might p⌝ will amount to a proposal to make p compatible with the common ground (and to make this fact itself common ground).

As a first pass at implementing this idea, we could take 'It might be raining' to mean, roughly: the proposition that it's raining is compatible with the common ground.[14] But this wouldn't work. On this view, modal claims would merely *describe* the common ground. Like the views discussed in the last section, this cannot make sense of the way that speakers use 'might'-claims to *negotiate* about the common ground. If I say that it might be raining, and you say it is not, we are not arguing about whether the common ground is *presently* compatible with the proposition that it's raining. We know this in part because if you think that the common ground is compatible with p, then it follows from the logic of the common ground that the common ground *is* compatible with p.[15] We are, rather, arguing about whether the common ground *should be* compatible with the proposition that it's raining. A parallel point extends to 'must'. On the present account, ⌜Must p⌝ would say that p is already entailed by the common ground. That would mean that 'It must be raining' could not be used to *inform* one's interlocutors that it's raining; it could only be truly asserted in a context where it is already accepted that it's raining. This again fails to capture the way in which 'must'-claims are used to negotiate about the common ground.

A close variant on this approach, however, avoids these problems. On this view—which I call *prospective contextualism*—'It might be raining' means that the proposition that it's raining *will* be compatible with the *prospective* common ground: the common ground as it stands *after the assertion is made and negotiated*,

finding such propositions looks quixotic if not impossible.' Accounts which follow this pessimistic line in rejecting the contextualist framework to capture the guiding observation include dynamic accounts such as Veltman 1996; Beaver 2001; Willer 2013; expressivist accounts such as Yalcin 2007, 2012, 2011; Rothschild 2011; Swanson 2015; Moss 2015; revisionary contextualist accounts such as von Fintel and Gillies 2011; Stalnaker 2014, and relativist accounts such as Egan et al. 2005; Stephenson 2007a; Kölbel 2009; Egan 2011; MacFarlane 2011, 2014. I discuss some of these approaches in §§5.1 and 6 below. A different defense of contextualism from the one I give below comes from 'flexible' contextualist accounts such as Dowell 2011; Khoo 2015, as well as the coherence theory version of contextualism in Stojnić 2016; for reasons of space, I cannot discuss those interesting views here.

[14] For ideas in this direction, see Groenendijk and Stokhof 1975, Yalcin 2007, MacFarlane 2011.

[15] Models for common ground for a group I can be constructed by taking the transitive closure of the accessibility relations for acceptance for each $i \in I$ (more on this in Footnote 17). That means that if i accepts that p is compatible with the common ground, then a p-world will be accessible in the accessibility relation for common ground, and thus p will be compatible with the common ground.

i.e. after the assertion has been either accepted or rejected by all interlocutors.[16] According to prospective contextualism, ⌜Might p⌝ will mean that p is compatible with the prospective common ground, ⌜Must p⌝ that p is entailed by the prospective common ground.

On this theory, an assertion of ⌜Might p⌝ is just a proposal to make p compatible with the common ground (and to make this fact itself common ground). Here's why. Suppose that ⌜Might p⌝ is asserted at time t; let t' be the time when the assertion has just been either accepted or rejected. Suppose first that ⌜Might p⌝ is accepted. In the contextualist framework, that means, again, that its content will be added to the common ground. Thus, at t', the proposition expressed by ⌜Might p⌝ will be in the common ground. If ⌜Might p⌝ has the meaning attributed to it by prospective contextualism, it follows that it will be common ground at t' that the common ground at t' is compatible with p. This is not yet what we want, but it turns out that in the logic of common ground, 'It is common ground that p is compatible with the common ground' *entails* 'The common ground is compatible with p'. In other words, when the common ground thinks something is compatible with it, it is always right. Importantly, this does not follow from any (non-trivial) assumptions about the logic of the attitudes which constitute the common ground, but rather from the iterative structure of the common ground, on which the common ground is what is accepted, accepted to be accepted, and so on.[17] Given this fact about the logic of the common ground, it thus follows that an assertion of ⌜Might p⌝ is a proposal which, if accepted, ensures that p is compatible with the common ground (at the time after the assertion has been negotiated), and that this fact itself is common ground.

Suppose second that ⌜Might p⌝ is rejected by all parties to the conversation; in the strong sense of rejection I have in mind, it follows that its negation is common ground at t'. According to prospective contextualism, the negation of ⌜Might p⌝ says that the prospective common ground, i.e. the common ground at t', is not compatible with p; that is, that it entails $\bar{p}$. Thus it will be common ground at t' that $\bar{p}$ is common ground. This, again, is not yet what we want, but, once more, the logic of common ground helps us close the gap. It follows in the logic of common ground that, if the common ground entails that the common ground entails some p, then the common ground entails p. In short, when the common ground thinks that it entails something, it is always right.[18] Thus if ⌜Might p⌝ is commonly rejected at t', then, at t', $\bar{p}$ will be common ground.

In sum: an assertion of ⌜Might p⌝, according to prospective contextualism, will amount to a proposal

[16] What the prospective common ground actually amounts to is of course a vague matter; in this respect the notion is on a par with standard notions in semantics and pragmatics, like the notion of a context. Importantly, the prospective common ground differs from the common ground at the point just after it has been updated with the fact that the claim has been made (see Stalnaker 1998; von Fintel 2008. Quantifying over that common ground does not avoid the problems just discussed for the non-prospective view.

[17] Again, models for common ground for a group I can be constructed by taking the transitive closure of the accessibility relations for acceptance for each $i \in I$, where the transitive closure of a set of relations $R_i : i \in I$ is R^*, where xR^*y just in case there is a sequence of worlds $w_1 \dots w_n$ such that $w_1 = x$, $w_n = y$, and, for each w_i and w_{i+1} for $i \leq n-1$, $\exists k \in I : w_i R_k w_{i+1}$. Then, interpreting $\Box_t$ as 'it is common ground at t that', and $\Diamond_t$ as its dual, this claim is just the claim that $\Box_t \Diamond_t p \rightarrow \Diamond_t p$ is a theorem in the logic of common ground. To prove this, we need only to assume that the underlying attitude of acceptance is always consistent, in which case the system validates $\Box_t p \rightarrow \Diamond_t p$ (corresponding to the serial constraint). Substituting $\Diamond_t q$ for p, we get $\vdash \Box_t \Diamond_t q \rightarrow \Diamond_t \Diamond_t q$. Any system for common acceptance will validate $\Box_t p \rightarrow \Box_t \Box_t p$, thanks to its iterative structure (corresponding to transitivity); contraposing and substituting q for $\neg p$, we get $\vdash \Diamond_t \Diamond_t q \rightarrow \Diamond_t q$. It follows that $\vdash \Box_t \Diamond_t q \rightarrow \Diamond_t q$.

[18] This amounts to the claim that $\Box_t \Box_t p \rightarrow \Box_t p$ is a theorem in the logic of common ground. Known as $C4$, this corresponds to the *density* constraint on accessibility relations ($\forall a \forall b : aRb \rightarrow (\exists c : aRc \wedge cRb)$), but it is perhaps more intuitive to derive it from the quasi-reflexivity constraint (every world which is accessed by a world accesses itself, corresponding to the constraint that attitudes of acceptance represent themselves as being veridical). Then the logic of common ground will validate $\Box_t(\Box_t p \rightarrow p)$. By the K axiom and modus ponens we have $\vdash \Box_t \Box_t p \rightarrow \Box_t p$.

which, if commonly accepted, makes p compatible with the common ground (and makes this fact itself common ground); and which, if commonly rejected, makes $\bar{p}$ common ground. If we define 'must' as the dual of 'might', exactly parallel reasoning shows that ⌜Must p⌝ is a proposal which, if accepted, adds p to the common ground; and, if rejected, ensures that p is not common ground, and thus that $\bar{p}$ is compatible with the common ground.[19]

Prospective contextualism, together with the contextualist model of conversation dynamics, thus predicts that assertions of ⌜Might p⌝ and ⌜Must p⌝ will have exactly the update properties ascribed to them by the guiding observation. In this framework, when we accept ⌜Might p⌝, we accept a proposition which says that p is compatible with *our own (prospective) common ground*. And to accept such a proposition *just is* to render p compatible with our common ground. Likewise, when we accept ⌜Must p⌝, we accept a proposition which says that p is entailed by our own (prospective) common ground. And to accept such a proposition *just is* to make p entailed by our common ground. Prospective contextualism thus shows that the guiding observation can be reconciled with the contextualist framework: we can view modal assertions as proposals to add a proposition to the common ground, and still make sense of the fundamental dynamics of modal language, provided that we take modal claims to express propositions which are themselves about the common ground.

Before moving on, let me note two important features about prospective contextualism. First, prospective contextualism is a theory about the *assertoric content* of unembedded modal claims: what proposition they contribute to the common ground when asserted. In §5.3 and the appendices, I'll sketch a compositional implementation of this theory, as well as a formalization of the main claims of this section. But my main concern in this paper is not with how to semantically generate prospective contextualism—and, therefore, not with the question of how modals embed—but rather how this theory of the assertoric content of modal claims answers to the guiding observation. Let me emphasize here, though, that while prospective contextualism is committed to the claim that ⌜Might p⌝ and ⌜p is compatible with the prospective common ground⌝ generally have the same *assertoric* content—i.e., contribute the same proposition to the common ground when asserted—it is no part of prospective contextualism that they have the same compositional *semantic* content, and thus no part of prospective contextualism that they will embed in the same way—an important fact, since these plainly embed in different ways. I return to this point in §5.3.

Second, what is essential to prospective contextualism is that it assigns to a modal claim a content which has certain introspective properties with regard to the common ground.[20] Making epistemic modals about the prospective common ground is an obvious way of accomplishing this, but there may be other ways—for instance, by making them about the prospective common *evidence*, and adopting introspective constraints regarding evidence. What is really essential to prospective contextualism, then, is not that the content of modal claims be itself about prospective attitudes, but that it be related to those attitudes in the right kind of

[19] An assertion of ⌜Must p⌝ is thus predicted to have the same main update effect as an assertion of p alone. As von Fintel and Gillies (2010) argue, this is broadly plausible, *modulo* some subtle but important differences; see Mandelkern pear for an account of these differences within the present framework. One issue which von Fintel and Gilles make much of is whether ⌜Must p⌝ entails p. On my account, the answer to this question will depend on whether we treat the common ground as factive or non-factive. It does not matter for the purposes of this paper which way we go.

[20] In particular, where t is the time of assertion and t' is the prospective time, it has the properties that ⌜$\Box_{t'}$(Might$_t$ p)⌝ entails ⌜$\Diamond_{t'} p$⌝; and that ⌜$\Box_{t'}$(Must$_t$ p)⌝ entails ⌜$\Box_{t'} p$⌝.

way. Making the content of modal claims about those attitudes is a simple and natural way to accomplish this, but the basic insights of this approach can thus be extended to any variants which preserve these logical principles. Having said this, I will set aside exploration of such variations for present purposes, in favor of the simple version of the view which I have presented here.

4 Performativity

Prospective contextualism shows that, *pace* much recent work, we can capture the guiding observation within the contextualist framework. It does not follow, of course, that we *should* do so. The main goal of this paper is simply to show that contextualism *can* capture the guiding observation, and thus that we can dismiss any rejection of contextualism based on its putative inability to do so. Whether prospective contextualism provides the correct theory of epistemic modality is not something I will take a stand on here; this question turns on a wide array of issues which I cannot hope to address in this paper. In the rest of the paper, however, I will try to dispel some obvious concerns about prospective contextualism, and bring out some *prima facie* attractions vis-à-vis some of its competitors; I hope to say enough, not to show that prospective contextualism is the correct theory of epistemic modality, but rather to show that we should take prospective contextualism seriously as a contender theory.

I will begin by addressing a natural worry about prospective contextualism, which goes like this: There is something very *weird* about prospective contextualism. It captures the guiding observation in a way which is both structurally and normatively implausible. According to prospective contextualism, an assertion of an epistemic modal claim is a proposal about how the common ground *should* evolve which is made by saying something about how it *will* evolve. This is roundabout: the *truth conditions* in this account are idle wheels. And this account is normatively implausible: speakers can assert modal claims without having any idea of what will actually happen to the common ground, and thus without knowing or believing the content which prospective contextualism assigns to them. Given that you generally have to believe what you say, prospective contextualism can't be the right theory of epistemic modals.

I will address this worry by showing that the method of negotiation which I have argued epistemic modals exploit—making a proposal about how some contextual parameter *ought* to be set by making a truth conditional claim about how it *will* be set—is in fact widespread in natural language. In particular, epistemic modal claims, in this approach, recapitulate the structure of a wide variety of performative assertions.[21] This fact shows that there is nothing suspicious about the structure which prospective contextualism attributes to epistemic modal claims. It also puts us in a position to address the normative issue for prospective contextualism, by arguing that performatives are governed by norms which do not require one to know or believe their content—explaining how speakers can assert modal claims, despite not believing the content ascribed to them by prospective contextualism.

[21] This treatment of epistemic modals as performatives can be seen as an extension of the approach in Lewis 1979, which treats some uses of deontic modals—syntactic and semantic siblings of epistemic modals—in this way.

4.1 Performatives in natural language

Performative assertions are assertions which aim not just to describe the world, but rather to change it. There are many different kinds of performative, and there has been much debate about their structure.[22] I will focus here on one kind of performative assertion which involves a particularly simple formulation, and thus is, I think, particularly revealing about the structure of at least some performatives. These are performatives which are made using sentences which appear to have perfectly ordinary truth conditions: sentences which lack any markers distinctive to performatives (like 'hereby', on which more below), and which can be used purely descriptively, simply to describe the world; or performatively; or both at once.

To see the kind of thing I have in mind, suppose Mark tells Mary, 'This afternoon, John will be cleaning the rabbit cage'. There is nothing unusual about this claim, and it has straightforward truth-conditions: it's true just in case John will be cleaning the rabbit cage this afternoon. Now suppose that Mark says the same thing[23] to John: 'This afternoon, you will be cleaning the rabbit cage'. Assume that Mark has the right kind of authority over John to tell him what to do. In that circumstance, this sentence can be used not (just) to inform John about the future, but to *make* that future come about: to make it the case that John *will* clean the rabbit cage, by establishing normative facts about what John *may* do. And, crucially, the assertion in question does this *by* making a claim about what he *will* do. When Mark tells John that he will be cleaning the rabbit cage this afternoon, Mark (according to the contextualist model I am assuming) communicates to John that Mark believes John will be cleaning the rabbit cage this afternoon. If John had no pre-existing plan to clean the rabbit cage, but John thinks that Mark has the authority to determine what John will do this afternoon, then John will reason that Mark would only believe that John will clean the rabbit cage if Mark is requiring that he do so; and thus John (assuming he recognizes Mark's authority to impose requirements on him) will change his plans to accord with Mark's beliefs.

What is nice about performatives like this one is that they seem to wear their structure on their sleeves. There is no *prima facie* reason to think that Mark's sentence in the second context is semantically different from his sentence in the first context. A more parsimonious thought is, instead, that these sentences mean the very same thing in the two contexts: it is just that, in the second context, Mark's sentence serves not just to describe the world, but also to bring about the state of affairs which it describes—by describing what the future will be like in a normative setting which is conducive to bringing about that future by describing it.

One could still maintain that sentences like this are ambiguous between a purely descriptive meaning, on which they have ordinary truth-conditions; and a performative meaning, on which they do not. One argument that this is wrong comes from the fact that, in many cases, a single assertion of a sentence like this can serve to *both* describe the world and change it. Anscombe (1963) gives an example along these lines. Imagine a doctor telling a patient, in the presence of a nurse, 'The nurse will now take you to the operating theater'. In this context, the doctor's claim is both a description of what is going to happen—a claim that the world will be one in which the nurse takes the patient to the operating theater—and a means by which she makes it the case that the world will be this way—given the background normative conditions in play in the context, the

[22] 'Performative' is sometimes reserved for more specific kinds of speech act, but the terminology doesn't matter here. For some of the debate, see e.g. Austin 1962; Cohen 1964; Strawson 1964; Searle 1968, 1969; Bach 1975; Bach and Harnish 1979; Davidson 1979; Bach 1994.

[23] *Modulo* a change of 'John' to 'you', which presumably doesn't change the meaning of the sentence.

nurse, in hearing the doctor's claim, will act so as to bring it about that her claim is made true. If the doctor's sentence was ambiguous between a descriptive and a performative meaning, then we would not be able to derive both meanings from a single assertion (in general ambiguous sentences get one disambiguation per assertion). There is just one thing going on here: an assertion, about the future, which both describes the future and thereby brings it about.[24] Examples with a dual nature like this one abound. Consider a sign at the front of a train track which says 'The front of the train will stop here' (Robert Stalnaker, p.c.). This sign serves both to describe the train's stopping point, and to instruct train drivers where to stop.

What is helpful about examples like this is that they attest the same performative structure that I am attributing to epistemic modal claims. They make claims with ordinary truth conditions. But these assertions amount to proposals about how to change the world by way of making a claim with ordinary, but prospective, truth conditions; a claim which, if accepted, ensures that the change in question takes place. And so these kinds of assertions show that the performative structure that prospective contextualism attributes to modal claims is, in fact, widely attested in natural language.[25] This should allay fears that that structure is implausibly roundabout.[26]

4.2 Conversational norms: The perspective from performatives

It also helps us answer a natural objection concerning norms of assertion. The concern, again, is that, if modal claims have the content I am ascribing to them, then speakers will often be able to assert epistemic modal claims without knowing, or even believing, that their content is true. I can propose, for instance, to make the common ground entail that Sue is in her office, by asserting 'Sue must be in her office', without knowing or even believing that the common ground will come to entail this: for all I know, you may fail to accept my proposal. If epistemic modal claims were governed by a norm along the lines 'Only assert what you believe (or know)', prospective contextualism would leave it puzzling how people could assert epistemic modal claims in ordinary circumstances. I will argue in this section that a similar puzzle arises for performatives in general, which shows that performative assertions are not governed by a belief or knowledge norm of this form in the first place, but rather by a norm oriented towards what agents do with their words. This provides a principled explanation of how, and when, we can reasonably assert epistemic modal claims even if we do not know or believe their contents.

Let us consider first performatives like those considered in the last subsection, such as 'This afternoon, you'll be cleaning the rabbit cage'. Could a performative like this be governed by a norm along the lines: Assert p only if you know (or 'believe', or 'justifiably believe') p?[27] Clearly not. Mark may well tell John that he'll be cleaning the rabbit cage as an *attempt to get him to do it*, even though he knows this attempt may

[24] We cannot, of course, assume that all performatives have the same structure, and thus these examples do not show that *all* performatives have truth-conditions, or more generally share the structure of these performatives (though I think these performatives suggest that it would be fruitful to look for an account along those lines); it is consistent with everything I say here that some performatives lack truth conditions.

[25] See Truckenbrodt 2009, Eckardt and Schwager 2009, Eckardt 2009 and Condoravdi and Lauer 2011 for similar ideas about performatives.

[26] It may be objected that these cases are explicitly marked out as being about the future, whereas epistemic modal claims are not. This is true. But note that prospective contextualism does not claim that modal claims are interpreted *in the future tense* (a claim which would be farfetched), but simply that they are assigned accessibility relations in a way which references the future.

[27] As in e.g. Williamson 2000; Lackey 2007. It doesn't matter for present purposes which of these mental states is the relevant one.

well not be successful. One way to see this point is to compare a sentence like 'You'll be cleaning the rabbit cage' to 'You have to clean the rabbit cage'. In many contexts, these sentences play the same role: to try to get John to clean the rabbit cage. But it would be ludicrous to think that the latter of these is accompanied by a norm which holds that it can be asserted only if the speaker knows the addressee will clean the rabbit cage. And it seems no more plausible to hold that such a norm governs the former. This, of course, is in stark contrast to non-performative assertions: if Mark asks John what he was doing in the morning, then it is generally *not* permissible for John to say that he was playing outside, if he doesn't believe this to be true.[28]

This is *prima facie* puzzling. But the puzzle, it seems to me, is not very deep: these observations are not hard to make sense of. Performatives are used to *do* things, and so the norms that govern their production (and negotiation) will concern in the first instance what the performative is being used to *do*, and only derivatively about the speaker's doxastic relation to its content. In short, I propose that performatives are governed by a very general norm along the lines: Assert p only if the action which you aim to accomplish with your assertion is permissible. Call this the *Speech Act Norm*. This norm rightly predicts that whether or not Mark may tell John that he will be cleaning the rabbit cage depends on whether it is permissible for him to require John to clean the rabbit cage—*not* whether he knows, or believes, that John will clean the rabbit cage. In general, I propose that the best way to make sense of the observations just adduced is to hold that performatives are in the first instance governed only by the *Speech Act Norm*, and are governed by more specific doxastic norms only when those follow from the *Speech Act Norm*. In some cases, more specific norms will indeed follow. But in the cases under consideration, nothing in the *Speech Act Norm* entails that speakers must know or believe the contents in question.

This should, I think, be fairly uncontroversial. Here I will take a somewhat bolder stance, and propose that what goes for performatives of the kind under discussion goes also for assertions in general. All assertions, of course, are speech acts. As such, I propose that the *Speech Act Norm* governs assertions in general, not just performative assertions. That is, I propose that the fundamental norm governing assertions cares in the first instance about what is done with those assertions. Assertions are speech acts which can have diverse goals; it is natural to think that the norms which govern assertions should be sensitive to that diversity, and should in the first instance target the basic kind of thing that assertions are—speech *acts*.

More specific norms, however, may follow from the *Speech Act Norm*, given ancillary assumptions about the kinds of acts that are permissible in general or in a particular situation. For instance, in most situations, it is not permissible to take actions that mislead others. In those situations, it follows from this

[28] One reaction to these observations is to argue that performative 'assertions' are not really assertions in the fullest sense: we should reserve 'assertion' for the class of speech acts which are governed by doxastic norms. This dispute seems essentially terminological. It seems to me that there are different, cross-cutting taxonomies which we can bring to bear on speech acts and which may be useful for different purposes. I have no quarrel with this alternate terminological choice; the important point is that, either way, we will make sense of how speakers assert epistemic modal claims by assimilating them to performatives. I will continue to use 'assertion' here to refer to the class of speech acts which are attempts to update the common ground with a given content; this class includes both non-performative assertions and performative ones, since performatives of the kind I have been discussing, both epistemic and non-epistemic, are at least in part proposals to update the common ground with their content. (For instance, 'This afternoon, you'll be cleaning the rabbit cage' is a bid to get John to clean the rabbit cage, but it is an attempt to do this which goes by way of an attempt to update the common ground with its content.) In the present taxonomy, this speech act thus counts, *inter alia*, as an assertion (it is of course also many other things: most prominently, an attempt to get John to clean the rabbit cage). Again, this taxonomic approach can co-exist with other, cross-cutting taxonomies, which may be more useful for different purposes.

normative fact, together with the *Speech Act Norm*, that an assertion will be permissible only if it does not mislead others. Indeed, it is plausible that in many situations a more demanding norm is in play, which requires that an assertion somehow contributes to coordination on common knowledge or belief. Call this norm, which follows from the *Speech Act Norm* plus independent assumptions about the general norms of communication, the *Common Ground Norm*.

The *Common Ground Norm*, in turn, entails a limited norm of the kind familiar from the literature on norms of assertion, again given plausible ancillary assumptions. Consider a sentence whose content is exclusively about a state of affairs which does not depend on the assertion of the sentence. Given the *Common Ground Norm*, you plausibly should assert a sentence like this only if you know and believe it: since an assertion is a proposal to update the common ground with the asserted content, an assertion of something you do not know, or do not believe, will not generally be conducive to the end of coordinating on common knowledge. It follows that you should know what you assert (or at least believe it, if we take a more subjective approach to norms), when what you assert is not a performative.

But when it comes to performative assertions—assertions which aim to bring about a state of affairs, not just describe one—nothing in the *Common Ground Norm* entails that one must know or even believe the content of the assertion when asserting it. Indeed, it is easy to see that in many cases knowing or believing the asserted content may be positively inconsistent with the *Speech Act Norm*: sometimes the best thing to do—even in cases in which a conversation is generally speaking aiming at the truth—is to assert something which one does not believe or know, in the hopes (or the knowledge) that it will *become* true in virtue of your assertion (as e.g. in the rabbit cage or nurse case above). The *Common Ground Norm* does entail that one's claim should generally contribute, or at least not detract, from the good epistemic standing of the conversation. It follows that performatives should generally be such that, *if they are accepted,* they will be commonly known.[29] But it doesn't follow that performatives have to be known—or believed—*antecedently*: one can assert risky performatives, performatives the speaker doesn't know will be accepted (and thus doesn't know will be made true), provided the aim of the assertion is a permissible one. This is particularly clear, I think, in the case of the practical performatives we considered in the last section. But, from the perspective of the *Speech Act Norm* and prospective contextualism, it is also very plausible in the case of epistemic performatives of the kind I am arguing epistemic modal claims might be. Sometimes an epistemically useful assertion can be one whose content is not antecedently known, provided that its content will be known if it is accepted; and that the assertion serves the purpose of coordinating on the conversants'

[29] Among other things, this suffices to explain the infelicity of performatives which command the addressee to do something that isn't in their power, like 'You will win the lottery', said as a command to buy a lottery ticket. It also goes some way towards explaining the infelicity of Moore sentences containing performatives as a conjunct: 'You'll clean the rabbit cage but I don't know you will' cannot amount to common knowledge, even if it is accepted, provided the time interval relative to which 'know' is evaluated includes the prospective time. In theory, this doesn't rule out performative Moore sentences if care is taken to ensure that the interval relative to which 'knows' is evaluated does not include the prospective time. Sentences like this seem to be ruled out on independent grounds, however: you shouldn't tell someone to do something, and then say you aren't sure if they will. To see this, note the oddity of 'You have to clean the rabbit cage, but I don't know if you will', or 'Clean the rabbit cage! I don't know if you will'. (See Silk 2015, 2018 for this observation. As a referee for this journal helpfully points out, the data here are subtle, but there does seem to generally be something marked about these conjunctions. The infelicity of these conjunctions can be brought out by contrasting these with corresponding weaker modal claims, like 'You should clean the rabbit cage, but I don't know if you will'.) It seems as though, in trying to get someone to do something, you must act as if you are confident that they will. *Why* this is true is another question, one whose answer is not obvious to me.

common knowledge. Epistemic modal claims, in the framework of prospective contextualism, have just this profile. One need not antecedently know these claims in order to assert them; one need only know (or have sufficiently good reason to believe) that the proposals these claims make will be epistemically valuable.

There is much more to do in exploring this perspective on the norms of conversation, but I will leave the discussion at this sketch for the present. The key point for our purposes is that this approach to the norms of conversation makes sense of intuitions both about when ordinary, non-performative claims can be asserted, as well as when broadly performative claims can be asserted. And, if epistemic modal claims are performatives, as prospective contextualism maintains, then, given the normative regime that governs performatives in general, we can resolve the normative concern about prospective contextualism. According to prospective contextualism, one can often assert modal claims without knowing their content to be true. But the fact that, according to prospective contextualism, modal claims are performatives, allows us to give a disciplined account of why this should be so: they are in the first instance used to *do* things—make proposals about what entailment and compatibility properties the common ground should have—and thus that they are governed by norms which look, in the first instance, at those actions, rather than at agents' relation to the truth conditional content of those claims.

This also helps us say a bit more about what prospective contextualism predicts regarding when a modal claim *can* be made. According to the *Common Ground Norm*, a principle normative concern regarding assertions is whether they contribute to the epistemic coordination of the agents in conversation. Thus an assertion of ⌜Must p⌝ —a proposal to make p common ground—will generally only be acceptable if one knows (or at least believes) p. This seems broadly plausible.[30] And an assertion of ⌜Might p⌝ —a proposal to ensure that the common ground is compatible with p—will, in general, be acceptable only if it contributes to the conversants' coordination on their knowledge. It is somewhat harder to say just what this amounts to, but it will be something like: the speaker must know that p is a possibility which should be treated as live in the conversation—perhaps because she has some evidence in favor of p, or in any case knows that her evidence is compatible with p, and believes that her interlocutors should know this too. Again, these constraints make sense of our intuitions about when a modal claim can be made, and (therefore) what considerations we take into account when deciding whether to make modal claims.[31]

The performative structure prospective contextualism attributes to epistemic modal claims thus furnishes a principled response to the normative challenge outlined above.

It may be objected at this point that, although the prospective contextualist framework is technically within the bounds of contextualism, it has abandoned one of the principle attractions of contextualism: namely, that in a contextualist framework, we can view assertions as coordination on the speakers' knowledge and beliefs; if you can assert propositions without believing them, the thought goes, we lose this attractive feature of contextualism. But this is mistaken. The normative framework I have sketched here still views conversation as (among other things) aiming at the coordination of speakers' knowledge and beliefs. It just turns out that—when the truth of a given content depends on the speech act of asserting that content—

[30] But see Lassiter (2016) for some objections.

[31] Apparent exceptions to these generalizations come from 'exocentric' uses of modals. I will not discuss those uses at length here, but I assume that those uses are somewhat exceptional, involving either free indirect discourse of some kind, or a tacit modal operator, either of which would suffice to explain our use of them. See Appendix B for brief further discussion.

you need not know or believe that content in advance of it being accepted, in order for your assertion to contribute towards coordination of the speakers' information.

5 Loose ends

This completes the core of my presentation of prospective contextualism. In this final section, I will briefly address a few remaining issues.

5.1 Comparisons

The discussion so far has shown that prospective contextualism's treatment of epistemic modal claims is more plausible than it might first seem: provided we pay careful attention to the phenomena of performativity in natural language in general, we can make sense of the way that speakers assert, and negotiate about, epistemic modal claims within the prospective contextualist framework. A natural question to ask at this point is whether prospective contextualism does a better job of this than competitor views, like expressivism[32] or dynamic semantics.[33] I will not try to answer this question here. I will instead briefly compare prospective contextualism with just one other view in the literature, namely that given in Stalnaker 2014. My view, which owes much to Stalnaker's, is similar in many ways to his, but has some important differences.[34] The comparison with his proposal will help bring out what is distinctive about prospective contextualism.

Stalnaker, following orthodoxy, proposes that epistemic modals have as their meaning a simple quantificational structure. Rather than building prospectivity into the meaning of epistemic modal claims, however, Stalnaker proposes that epistemic modals are coupled with a special force rule. This force rule specifies that an epistemic modal claim is to be interpreted as a proposal to make the prospective context set (the set of worlds compatible with the common ground) verify the modal claim: i.e. as a proposal to render the prospective context set such that, if the modal claim is interpreted with that set as the modal's domain of quantification, the claim comes out true. Stalnaker's proposal nicely conforms to the guiding observation. But the proposal also diverges from the contextualist framework in an interesting way, by associating epistemic modal claims with a distinctive force rule. Stalnaker's proposal and mine are essentially routes to the same end, using quite similar tools. But where prospective contextualism fully conforms to the contextualist framework laid out above, Stalnaker's approach diverges from it: instead of assimilating modal updates to other kinds of propositional updates, Stalnaker introduces a special force rule that applies only to modal updates.

This move, in my view, has some drawbacks. First, the multiplication of force rules in Stalnaker's account adds complexity to the pragmatic framework which looks unmerited from the point of view of prospective contextualism, which has only one force rule—the standard assertoric update rule—for both modal and non-modal claims. This multiplication of force rules also entangles the pragmatic theory with

[32] See Yalcin 2007, 2011, 2012; Rothschild 2011; Swanson 2015; Moss 2015.

[33] See Veltman 1996; Groenendijk et al. 1996; Aloni 2001; von Fintel and Gillies 2007; Willer 2013; Yalcin 2015.

[34] Stalnaker's own presentation is brief, and there are different ways to interpret it; the present interpretation provides a useful foil for my own view, whether or not it is precisely what Stalnaker had in mind. For an earlier intimation of Stalnaker (2014)'s approach, see Stalnaker 1993.

lexical semantics in an unattractive way: on Stalnaker's view, what triggers the distinctive prospective force rule is a claim's being an epistemic modal claim. This pragmatic rule thus makes reference to specific lexical items—modal words—rather than to the output of the semantic operations.[35] A closely related, more technical, and, I think, more serious point concerns sentences which involve both modal and non-modal constituents. Consider a disjunction with the form ⌜p or might q⌝. What update rule should apply to these sentences? The prospective update rule simply does not make sense in application to the first disjunct; but the ordinary update rule will bleach out the proposal-like quality of the second conjunct. It looks like to make sense of assertions like this (and similar conjunctions) we need a single update rule which applies to both modal and non-modal sentences.

This is not to say that there is not something right at a *descriptive level* about Stalnaker's characterization of the force of modal claims. A comparison with performatives like 'This afternoon, you'll be cleaning the rabbit cage' is, again, helpful. There are prescriptive uses of this sentence, and also descriptive uses. These uses have different effects, and are appropriate in different cases; and, at the level of descriptive taxonomy, it is helpful to distinguish them. But I don't think there is reason to say that these uses are associated with distinctive *force rules* from the perspective of formal pragmatics (or semantics). Instead, it seems preferable to adopt a unified force rule which predicts that which of these different uses (prescriptive and descriptive) is brought out depends on varying background conditions (the normative authority of the speaker, their relationship to the addressee, and so on). Things are parallel for epistemic modals. While there is, at a descriptive level, something right in saying that modal claims have a distinctive force, in that they have a characteristic update effect of making a proposal about how the prospective common ground will look, I am inclined to think that we should not encode this in our pragmatic system; in particular, if we are to have any hope of making sense of disjunctions and conjunctions of epistemic modals with non-modal language, then, from the perspective of our general pragmatic system, we need a single unified force rule for both modal and non-modal claims. Together with an appropriate account of the assertoric content of modal claims, the relevant descriptive generalization of the force of modal claims will then fall out as a consequence, rather than being stipulated in the system from the start.

This is not meant to be a decisive argument against Stalnaker's approach, but these points help bring out the differences between our proposals. Nor is this the place for a careful examination of the other alternatives to prospective contextualism. Maybe one of them will turn out in the end to be preferable. But, in every case, we should explore whether those revisionary approaches retain the simplicity and predictive power of the contextualist framework which prospective contextualism exploits. Let me emphasize in closing one attraction of that framework. By staying within that framework, prospective contextualism gives a fully deterministic account of the dynamics of conversation, at least once we have coupled it with a theory of belief revision (see again Appendix A). In the prospective contextualist framework, modal claims, like all assertions, express propositions. So updating with a modal claim *just is* updating with an ordinary proposition—in other words, coming to accept a piece of information. Once we have a theory of belief revision in place, we can thus advert to general considerations about belief revision to give a fully deterministic,

[35] See Szabó 2016 for related criticism. Importantly, this is in part because the meaning of modal claims will be *structurally* just like other propositions; compare the situation with questions or imperatives, which many believe are associated with distinctive force rules, but also have contents with distinctive semantic types.

and independently motivated, theory of how such updates will go. By contrast, in most revisionary theories, modal claims do not express propositional contents, or else (as we have just seen in Stalnaker's account) we do not update with the expressed proposition in the ordinary way; and thus modal updating cannot be assimilated to belief updating in general. Those theories can, of course, propose rules of modal updating which will have downstream effects on belief updating, and they can posit a variety of connections between modal updating and ordinary belief revision; I will not explore or attempt to judge those attempts here. My point here is the simple and *pro tanto* one that those theories must give an account of the connection between modal updating and belief revision, whereas on prospective contextualism, this question simply does not arise: modal updating simply *is* belief revision, and so a theory of the latter will suffice for a theory of the former.[36]

5.2 Descriptive uses

According to prospective contextualism, epistemic modal claims have truth-conditions; it is in virtue of those truth conditions' prospective nature, and the interlocutors' privileged position with respect to determining features of the common ground, that epistemic modal claims have the force they have. One feature of the truth-conditional performatives we looked at in §4.1 was that, in the right contexts, those sentences had clearly distinguishable descriptive content, and could be used in a purely descriptive, and not performative, fashion: thus, for example, if the nurse has already been told what to do and is out of earshot, 'The nurse will now take you to the operating theater' can serve in a purely descriptive fashion. This leads us to expect that. if prospective contextualism is correct, epistemic modals, even in their performative uses, have purely descriptive content; and, furthermore, that they will have purely descriptive uses (cf. Lewis (1979)'s claim that deontic modals likewise can have both a descriptive and a performative function). In this section I will try to make the case that this is the correct prediction, though I will conclude that things are not altogether clear here.[37]

Let me first make a case that epistemic modals do have descriptive content which can be distinguished from their performative force. The first part of the case is based on the 'stand your ground' cases discussed

[36]To see this point in relation to other heterodox theories, take dynamic semantics. Consider an update with ⌜Must p⌝, in a context c which does not entail p or $\overline{p}$. The most natural update rule we could associate with dynamic semantics simply says that, when a sentence q is asserted in context s, we move to context $s[q]$. This is a fully deterministic update rule. But if we applied this rule in the present case, an assertion of ⌜Must p⌝ would take us to the empty set: asserting ⌜Must p⌝ in this context would have the effect of moving us to a contradictory context. This is of course not what it in fact does. What we want is an update rule that predicts, roughly, that an update with ⌜Must p⌝ takes our context to one which entails p. Presumably dynamic semanticists will invoke some kind of pragmatic rescue mechanism to get this result. (Willer's variant on dynamic semantics avoids this particular problem, but faces a related version of the problem: in a context in which ⌜Might p and might not p⌝ has been accepted, any subsequent assertion of ⌜Must p⌝ will take the context to the empty set; and so Willer still needs some kind of pragmatic rescue mechanism.) The point is not that there is nothing that can be said here, just that something must be said that goes beyond the simple update rule under consideration. A similar point, *mutatis mutandis*, can be made in the case of Yalcin's expressivism. On that view, an assertion of a sentence is a proposal to move to a context that accepts that sentence. This means in particular that an assertion of ⌜Might p⌝ in a context c is a proposal to move to a context c' that is compatible with p. But which context? Nothing in the system answers that question, and there are, of course, arbitrarily many contexts that are compatible with p, so this is a substantive question. Again, there are responses available here; my point is simply that *something must be said* to fill out these theories to cover the case of modal updates. By contrast, prospective contextualism does not incur any similar debt: on that view, modal claims express propositions; once we have a story about how we update with propositions (i.e., a general theory of belief updating), we thereby have a fully specified story of both non-modal and modal updating.

[37]Thanks to two anonymous referees for this journal for very helpful comments on all these points.

in von Fintel and Gillies 2011. Suppose Ann says 'The keys might be in the car'; Bill accepts her claim. Then, sometime later, they both discover that the keys are not in the car. If Bill takes Ann to task for not knowing where the keys were, Ann might reply: 'Look, I didn't say they were in the car. I only said they might be there—and they might have been.' Here, we can interpret Ann as standing by the truth-conditions of her claim—namely, that the prospective common ground in her conversation was compatible with the keys being in the car—even if she no longer stands by the proposal that she made to leave it open that the keys were in the car (at least from some objective vantage point).

Cases like this thus seem to provide some evidence that epistemic modal claims do indeed have a descriptive content which can be distinguished from their performative function, and which plays an independent role in conversational practice. In this sense they seem parallel to the truth-conditional performatives I discussed above. In particular, compare overtly epistemic performatives, like 'Ok, we now all recognize that we're in agreement that the keys are in the car'. Suppose that Ann says this as a way to make it the case that the group recognizes that they agree that the keys are in the car. If Mark then discovers the keys are not in the car, he can say: 'You were wrong about the keys!' Ann might concede that the group was wrong (in some objective sense) to agree that the keys were in the car, but she could also fall back on the descriptive content of her claim: that is, she could point out that it *was* true that the group recognized they agreed that the keys were in the car. This case seems parallel to the stand your ground case: in both cases, speakers can stand their ground on the descriptive content of their claim, even if they no longer endorse the proposal it was used to make. In this respect, modal claims and truth-conditional performatives indeed seem parallel.[38]

Do modal claims also have purely descriptive uses, like other truth-conditional performatives? Consider a claim of the form 'Ok, we now all recognize that we as a group leave it open that the keys are in the car'. This has performative uses, but it also has purely descriptive uses: Ann could assert this as a way, not of getting her interlocutors to recognize that it is compatible with the common ground that the keys are in the car, but rather as a way of simply stating something that is already commonly known. Moreover, Ann need

[38] An anonymous referee for this journal points out that the parallel looks more limited when we look at a broader range of 'stand your ground' cases. For instance, suppose that Bill rejects Ann's claim, and shows her that the keys are in the kitchen. It still seems like Ann can say 'I only said the keys might be in the car'. But in this case, the descriptive content ascribed to Ann's claim by prospective contextualism is false, since the prospective common ground did not end up being consistent with the keys being in the car. Presumably we want a unified explanation of this kind of case with the kind of case introduced in the main text. I can see two ways of responding to this worry. One option sticks with the line taken in the main text, and holds that in cases like the one just described, Ann is insisting that we interpret her claim as not being prospective at all, but rather as describing the context as it stood when she made her claim. A worry with this line of response is that 'stand your ground' responses seem available in most cases, so this seems to undermine the claim of this paper that a prospective interpretation is the default one. But we could couple the present response with something along the lines of von Fintel and Gillies (2011)'s cloudy contextualism, maintaining that what proposition modal claims express is typically underdetermined; that, when all interlocutors are content with it, the prospective interpretation is the default one; but that interlocutors may fall back on a weaker, non-prospective interpretation when pressured to do so as in cases like this one. On this way of thinking, both 'stand your ground' cases given here do indeed point to the descriptive component of epistemic modal claims, even though, after the interlocutors' negotation, only the first case has a prospective interpretation. A second response is to explain all the 'stand your ground' cases in a different way—as drawing attention, not to some true descriptive content, but rather to the relative weakness of the proposal they made (namely, as proposals to leave a content open, not to accept it). If we go this way, then of course 'stand your ground' cases do not have the dialectic force I ascribe to them in the main text of pointing to a descriptive content of epistemic modal claims separable from their performative force. However, essentially the same point is made with the Mordecai cases and subsequent variations given immediately below, so if we took this line, we could still rely on the latter to show that epistemic modal claims, like truth-conditional performatives, have descriptive content which can be distinguished from their performative force. I will not try to decide between these two lines of response here.

not herself think there is any chance that the keys are in the car. Do we find parallels in the modal domain?

I think we do. Von Fintel and Gillies again supply a useful illustration, in von Fintel and Gillies 2008. Pascal and Mordecai are playing Mastermind, and Mordecai has given Pascal some hints. It is compatible with all these hints that there are two reds. Pascal says, 'There might be two reds'. Mordecai can accept Pascal's claim, even if he knows that there are not two reds. Likewise, Mordecai could himself assert 'There might be two reds', as a way of summarizing the hints he has given to Pascal. In both cases, it seems that Mordecai is simply *describing* the information jointly available to the interlocutors; in light of Pascal's limited information, Mordecai knows it will remain compatible with the common ground that there are two reds, even though Mordecai himself knows that there are not two reds. This seems like a good candidate for a purely descriptive use of an epistemic modal claim.

On the other hand, in this case, although it is intuitive to describe Mordecai as simply describing the group's limited information, it is not as though Mordecai *objects* to leaving it open that there are two reds: he rather recognizes that it is part of the norms of the game that he keep Pascal's information limited. So—while this case provides clear further support for the claim that epistemic modal claims have descriptive content which can be distinguished from their performative force—it is somewhat less clear whether this case provides support for the claim that epistemic modal claims have uses which are purely descriptive and not at all performative.

To really distinguish a purely descriptive use from a performative one, we should see whether speakers can use epistemic modal claims while positively objecting to the proposal that they make. Suppose that Ann is absolutely convinced the keys are in the car, and Mark is convinced they might be upstairs. Ann tries to convince Mark that they are in the car, but she recognizes that Mark simply won't rule out the possibility that they are upstairs. Could Ann say, 'Ok, fine. The keys might be upstairs. Let's check'?

It certainly seems acceptable for Ann to say this. It is a bit hard to diagnose exactly what is happening when she does, however. One perspective on this is that Ann is indeed simply describing the common ground, and thus that this is indeed as a purely descriptive modal claim. On the other hand, we might view this as a concession by Ann to Mark—a simple change of mind—and thus as having performative force. I don't see a clear way to decide between these two options; but the pattern here at least looks consistent with the predictions of our account.[39]

One way to try to distinguish these options is to make Mark's claim a totally outlandish one, to make it clearer that Ann is not just conceding the point. Suppose that Mark is convinced that there might be monsters under his bed and wants Ann to check for him. Ann knows that there are not monsters under Mark's bed, and does not even entertain this possibility. It seems that she might, however, say 'Fine, so there might be monsters under your bed. What do you want me to do about it?' This kind of unendorsed 'might'-claim—which seems to me to have much in common with Egan et al. (2005)'s 'exocentric' modal claims—seems like a clear candidate for a clearly descriptive 'might'-claim.

In some more outlandish cases, a purely descriptive reading is at least *prima facie* harder to get. To take a case suggested by an anonymous referee for this journal, suppose that Ann and Mark take a pill which

[39]We might try to distinguish these by noting that Ann *can't* say, 'Ok, well the keys aren't upstairs, but they might be'. We might expect this to be acceptable if a purely descriptive reading were available. But this test won't work, because, if we adopt the semantics sketched in the Appendix, this sentence will be a contradiction, which will independently account for its infelicity.

in ten seconds time will change their beliefs, so that they will leave it open that they now live on Mars. In this context, Ann could truly assert 'It will shortly be compatible with the common ground that we live on Mars'. But it does not seem that Ann could truly assert here 'We might live on Mars'.

I am not entirely sure what to say about this case, but let me sketch two possible responses. The first is to note that timing is quite important. The prospective time of Ann's assertion is by definition whatever time negotiation finishes, and it is that time that 'might'-claims target on my account. In ordinary conversations, agreement is immediate or at least very fast—extended negotiation is the exception. So if Ann says 'We might live on Mars' and Mark immediately agrees before the drug kicks in, then her claim will be descriptively false. It only has a shot at being descriptively true if the drug kicks in before the negotiation is finished. And if we think about the situation this way, it's less clear to me that Ann's claim is unacceptable. Suppose that Ann knows the drug will kick in at exactly 10 pm. She is watching the second hand on the clock, and, just before it hits 12, she says 'We might live on Mars starting. . .now!' (with 'now' asserted at exactly 10 pm). Here her assertion sounds much less strange to me. In short, the prospective time is not just any time in the near future: it is usually a very proximal, perhaps immediate, time, and so it will not generally be easy to use modal claims to talk about the future. When we adjust the case so that the prospective time really is a time at which the interlocutors leave it open that they live on Mars, the modal claim sounds much less strange.

A second response to cases like this one and the preceding one would be to argue that epistemic modal claims, *unlike* the truth-conditional performatives above, have some kind of mandatory performativity built in. Given the parallels surveyed so far in this section, I am somewhat disinclined to go this way. However, let me note that, if we did go this way, we could do so in a way that left intact many of the features of prospective contextualism. Consider again performatives like 'I am commanding you to clean the rabbit cage'. A sentence like this can be used in a performative sense, as a way of getting someone to clean the rabbit cage; or in a purely descriptive sense—say, if you have just ordered someone to clean their rabbit cage in a language they don't understand, and you are expaining to them what you have just done. But adding 'hereby' rules out the purely descriptive sense: 'I am hereby commanding you to clean the rabbit cage' cannot be a description of an order you have just given in another language; it must itself be the performance of an order. There is much to say about how 'hereby' brings this about (see Eckardt 2012 and citations therein for discussion). But one thing which I think is natural to think is that the fundamental dynamics of a performative like 'I am hereby commanding you to clean the rabbit cage' are not all that different from those of a performative without 'hereby', like 'I am commanding you to clean the rabbit cage' (used as a performative): I think it is plausible (though of course not uncontroversial) that these operate in similar ways, are governed by similar norms, and so on.[40] So if we did want to say that epistemic modal claims encode a kind of mandatory performativity, one option would be to do so by building on the truth conditions of prospective contextualism and then saying that epistemic modals also encode something like a 'hereby' operator. Such an account could leave intact much of the picture put forward here, in particular the truth-conditions of epistemic modal claims, the mechanism by which they can change the context, and the norms that govern them, and would simply augment this account with the stipulation that they must be

[40] Although 'hereby' seems not to sit well with truth-conditional performatives like those considered above, other adverbials seem to have the same effect in those cases, e.g. 'You will *in virtue of this utterance* clean the rabbit cage'.

used in a performative way. Since we know that we already have operators like 'hereby' in natural language which have the effect of ensuring that a sentence is used performatively, this move does not seem implausible or unparsimonious. Having said that, it would indeed complicate the picture; and I am not convinced it is necessary, given the parallels discussed above which provide at least some evidence that epistemic modal claims do have purely descriptive uses. So I will not pursue this kind of emendation further here, but note again that it is available, and would remain in the neighborhood of prospective contextualism.

5.3 Semantic implementation

Prospective contextualism, again, is a theory about the assertoric content of modal claims: what proposition they contribute to the common ground when asserted and accepted. Prospective contextualism as I have presented it remains agnostic about the *compositional semantics* of those claims: i.e. about how modal claims end up with the assertoric content they have, as a result of compositional semantic machinery interacting with pragmatic considerations.[41] This is not my main interest here; there are a number of different semantic routes to prospective contextualism, and the choice between them will depend on considerations involving embedded modals which are beyond the scope of this paper and largely irrelevant to its main claims. But I will briefly sketch what I take to be the most promising route to take on these questions.

First, let me note that prospective contextualism is broadly compatible with the standard "relational" semantics for modality sketched at the outset, on which modal claims quantify over a set of accessible worlds, where the accessibility relation is determined by context. The main innovation in my view is about *what set of worlds counts as accessible*. However we capture this fact, we should not need to depart much from the standard semantics for modals.

The most obvious approach would be to hardwire the prospectivity into the content of modal claims, by stipulating in the semantics that a context c determines an accessibility relation which takes every world to the set of worlds compatible with the prospective common ground, at that world, of the actual conversation's counterpart there. This approach would serve our purposes, but it would go wrong when it comes to embedded modals. This approach would maintain that ⌜Might p⌝ and ⌜The prospective common ground is compatible with p⌝ have *the same semantic content*. But this is plainly false, as can be seen, for instance, by embeddings of the kind discussed in Yalcin 2007: sentences with the form ⌜Suppose p and might not p⌝ are felt to be infelicitous, but not so for sentences of the form ⌜Suppose p and the prospective common ground is compatible with $\bar{p}$⌝. This means that, even though we are maintaining that ⌜Might p⌝ and ⌜The prospective common ground is compatible with p⌝ have the same *assertoric* content, we cannot identify their compositional semantic contents; we need a more sophisticated approach.

A second natural approach would be to take the standard modal semantics as our compositional semantics, and then assume that the prospective interpretation happens entirely at the pragmatic, 'post-semantic' level. This does not seem implausible to me, but it again leaves unexplained embedding data.[42]

A more plausible approach, I think, walks a line between these two. The idea is to start with the standard

[41] On the distinction between assertoric and semantic content see e.g. Ninan 2010.

[42] This includes Yalcin's data, as well as similar data involving quantifiers (sentences like Aloni's 'The biggest flea might be the smallest flea' (Groenendijk et al., 1996; Aloni, 2001)) and connectives ('I might win and I won't, or I might lose and I won't' (Mandelkern, 2019)).

semantics, but augment it by assuming that modals come with a constraint that ensures that only worlds in the modal's *local context* are accessible. The notion of a local context comes from the literature on presupposition projection, and amounts to the notion of what information is locally available at the point of processing a given part of a sentence (see Stalnaker 1974; Karttunen 1974b; Schlenker 2009). The local context for unembedded claims is just the set of worlds compatible with the common ground. Since the worlds compatible with the prospective common ground will always be presupposed to be a subset of the worlds compatible with the present common ground (since it is generally presupposed that conversants will gain, not lose, information), they will constitute an admissible set of accessible worlds. General pragmatic considerations would be responsible for ensuring that, for unembedded modals, this is the particular subset selected. Absent other clues, this will be an obvious default, since it provides a strong and natural interpretation of what speakers might be trying to do with modal words. The benefit of this approach over hardwiring prospectivity into the semantics is that it gives us a better framework for making sense of embedding data; the local context can be shifted by embedding operators, explaining a variety of data involving embedded epistemic modals. I sketch this framework in more detail in Appendix B, and, in Mandelkern 2019, argue that this approach provides a satisfying account of the embedding behavior of epistemic modals across the board.[43]

In particular, this approach predicts that, while an unembedded assertion of ⌜Might p⌝ will generally be interpreted as having the same content as an unembedded assertion of ⌜The prospective common ground is compatible with p⌝, these two expressions will *embed* differently, since they have different compositional semantic values; in particular, for instance, we predict that ⌜Suppose p and might not p⌝ will be felt to be equivalent to a command to make one's suppositions inconsistent, since ⌜S supposes p and might not p⌝ will be true just in case S's suppositions are inconsistent, thanks to the constraint that a modal's domain of quantification is limited to its local context. I leave it to the reader to formally verify this fact, based on the semantics in Appendix B.

This approach, further, predicts that prospective contextualism will be a default interpretation, but that there may be situations in which other bodies of evidence are sufficiently salient that we interpret modal claims differently. This seems like a plausible way of walking the line between, on the one hand, the guiding observation that modal claims are generally used to negotiate about what possibilities to treat as live; and, on the other, the observation that, provided suitable contextual set-up, they may be used simply to describe someone's state of mind, as in the 'stand your ground' cases discussed in §4.1, or in 'exocentric' readings of modals.[44]

Although this is not the place to discuss embedded modals, let me note that prospectivity will still play an important role in their interpretation. Thus e.g. ⌜If p, then must q⌝ will, as a default, be interpreted as

[43] As a referee for this journal helpfully points out, on this approach, epistemic modals end up being semantically distinguished by their locality constraint. As I discuss in greater length in Mandelkern 2019, this distinction seems to me to be required by the differences in embedding behavior of epistemic modals from other kinds of modals. Contextualist approaches to epistemic modals have sometimes been motivated on the basis of considerations about unity across different modal flavors. On my approach, while there is less disunity than in a dynamic or expressivist approach (since epistemic modals maintain their core relational meaning), there is semantic disunity across flavors. I should emphasize, then, that my arguments for contextualism here come not from the attractions of unity of different modal flavors, but rather from the attractions of the theory of conversation that contextualism fits into.

[44] See also e.g. Egan et al. 2005; von Fintel and Gillies 2008; Kratzer 2012, for related cases like this.

saying that q is true in all the p-worlds of the *prospective* context. Just as with unembedded modals, this allows us to interpret modals embedded under connectives in a way which makes sense of what speakers are doing with them (in this case, negotiating about whether to make all the p-worlds in the context q-worlds); thus in particular an assertion of ⌜If p, then must q⌝ will have the same update effect as an assertion of ⌜Not p or q⌝ (despite having a different semantic value).

Spelling out these ideas in detail goes well beyond the scope of this paper, whose goal is to show what assertoric content we must assign to modal claims in order to conform to the guiding observation; but this should suffice to give a sense of how we might implement prospective contextualism in a semantically respectable way.

5.4 Disagreement

The second issue concerns certain kinds of reactions to modal claims of the sort that have motivated relativism about epistemic modals. If I overhear Sue saying, 'John might be in his office', and I know that John isn't in his office, then relativists claim that it is appropriate to respond to Sue by saying 'No' or 'That's not true', *even if I am not part of Sue's conversation*. This is a *prima facie* puzzle for any brand of contextualism, including prospective contextualism, which predicts that Sue's assertion is a proposal to leave it compatible with *her* common ground that John is in his office, not with *my* common ground. Puzzles like this have led some to reject contextualism, arguing that truth is determined relative not only to a world but also to a judge or information state supplied by the context of assessment.[45]

This challenge is complicated, both theoretically and empirically. I will, again, not attempt to fully address it here, but let me make two brief remarks about it. The first is that there are theoretical reasons to resist a move to a relativist framework, reasons similar to those I gave above for resisting the move to dynamic and expressivist frameworks. When we move to a relativist framework, we must abandon the view of conversation as figuring out, together, which world we are in, since, from a relativist perspective, different conversants are in different "worlds"—different relativist points of evaluation. This makes it hard to see what the point of conversation is and how we should model its dynamics; although there have been some attempts to answer these challenges, I do not think they have been successful, though there is not space here for an adequate discussion.[46] Again, it may be possible to answer these objections in a satisfying way, but these

[45] See e.g. Egan et al. 2005; Stephenson 2007b,a; Lasersohn 2009; MacFarlane 2011.

[46] Cf. Stojanovic 2007 for related criticism. See Egan 2007; Stephenson 2007a,b for attempts to answer these challenges, neither of which I find satisfying. Egan proposes that ⌜Might p⌝ is evaluated relative to world-individual pairs, and true just in case p is compatible with what's in the 'epistemic reach' of the individual at that world. He models the common ground as a set of world-individual pairs, and models assertions as adding information to the common ground by intersection. For this to make sense, Egan shows that there must be a presupposition in place that all conversants are relevantly the same with respect to their epistemic reach (otherwise, assertions will end up 'stranding' conversants). Egan argues that this is plausible, since it is just part of the notion of epistemic reach that if I'm talking to someone for whom p is entailed by what's in their epistemic reach, then p is entailed by what's in *my* epistemic reach. But such a strong notion of epistemic reach leads to trouble. Egan argues that you must believe what you assert; thus to assert ⌜Might p⌝ in Egan's system, a speaker would have to believe that p is compatible with what's in her epistemic reach; it follows from this strong notion of epistemic reach that she would have to believe p is compatible with what *everyone* in the group knows (since whatever someone knows is, presumably, entailed by their epistemic reach). But, like the group contextualism considered and rejected in §2.2, this makes the assertion of 'might'-claims implausibly demanding. Stephenson's account avoids this issue by arguing that the content one must *believe* in order to assert ⌜Might p⌝ is substantially weaker than the content that gets added to the common ground: '[I]n order for a speaker A to assert a sentence S, it must be the case that for all of A's doxastic alternatives $\langle w', t', x\rangle$, S is true at the index $\langle w', t', x\rangle$ . . .this means that A must believe that S is

issues at least make it well worth our while to explore whether the move to relativism can be resisted within a contextualist framework.

The second point is that, by giving a successful contextualist model of the basic intra-contextual dynamics of modal claims, prospective contextualism may provide a theoretical foundation for a successful model of the dynamics of the cases which have motivated relativism. Whether it can do so depends on what exactly the empirical picture is, which turns out to be much more complicated that relativists first suggested. Recent empirical work in Knobe and Yalcin 2014; Khoo 2015 suggests that, contrary to the main claim in the relativist literature, subjects are not actually much inclined towards cross-contextual truth-value contestations: that is, if I overhear Sue saying that John might be in his office, subjects do *not* find it to be particularly appropriate for me to reply 'That's not true' (they find this to be a markedly less appropriate response than in corresponding non-modal cases). That finding is in line with the predictions of prospective contextualism, and contextualist approaches to epistemic modals more broadly speaking.[47] But that same work shows that expressions of disagreement ('No, Joe is at home') and subsequent retraction ('Scratch that, Joe is at home') are generally felt to be reasonably acceptable in contexts like the one described, which remains a *prima facie* challenge for contextualism.

Khoo (2015), however, has provided a persuasive account of how a contextualist account can make sense of these facts. The basic idea is that, *if* a contextualist account can predict (what we have called) the guiding observation, then 'might'-claims will be felt to be proposals to make their prejacents compatible with the common ground, and cross-contextual or retrospective disagreement can target exactly this proposal. To reject or retract a 'might'-claim is, on this approach, to signal that one does not stand behind the proposal made by the 'might'-claim in the first place. Thus a response of the form 'No, Joe is at home!' can be interpreted as, essentially, saying that it is a bad idea to leave open the possibility that Joe is in his office—what Sue is proposing in saying 'Joe might be in his office'. Khoo provides a careful exposition of this idea; I refer readers to the paper for details.

Crucially, Khoo notes that this kind of response is contingent on finding a version of contextualism which predicts the guiding observation, observing that 'No contextualist theory has attempted to predict

true as judged by A, but does not need to believe that S is true as judged by the whole group of conversational participants. Thus the norm of assertion is crucially weak in a certain sense. In order for A to assert that S, A only needs to believe that S is true as judged by A, but if A's assertion is accepted by the other speakers and added to the common ground, it has the same effect as adding the proposition that S is true as judged by the group of conversational participants' (Stephenson, 2007a, p. 509), because the common ground comprises triples which are all centered on 'the plurality of the group of participants in the conversation'. Thus what gets added to the common ground when ⌜Might p⌝ is asserted is something stronger than what one must believe to assert it, namely, not that p is compatible with the speaker's epistemic possibilities, but rather that p is compatible with what the whole group's epistemic possibilities. This mismatch between asserted content and updated content avoids the problem I raise for Egan, but it seems *ad hoc* to me: I can't see independent justification for this weak norm of assertion, and we lose the simple picture of the orthodox system in which what gets added to the common ground is the same proposition that is believed.

[47] To the degree that subjects still find it to be slightly appropriate, there remains something to be explained here. Two possibilities, both of which seem plausible to me: first, expressions like 'That's (not) true' can sometimes serve simply to register broad (dis)agreement, rather than to contest a truth-value. (Evidence for this comes from the fact that expressions like this are used reasonably frequently in response to questions. 'Have you done a follow-up examining cross-cultural variations in this result?' 'That's true, that's an excellent line to pursue.') Second, pronominal expressions like 'That' or 'What S said' may in some cases refer to something other than the proposition expressed by S. It is well known at this point that pronouns in general have "sloppy" uses (see Karttunen 1969), and it seems perfectly plausible that this goes for these pronominal expressions, too: they may be able to pick out something like a function from assignments of contextual parameters to propositions, rather than a proposition. There are obvious limits to the extent to which this is possible, limits which I think are fairly straightforward to make sense of within standard theories of the φ-features of pronouns, but that is a topic for another time.

the dynamic update effects of uttering epistemic modal sentences' (Khoo, 2015, p. 529). But prospective contextualism does just this, showing how we can make sense of the dynamics of epistemic modality within a contextualist framework. Thus, in concert with a story about (dis)agreement along the lines Khoo gives, prospective contextualism provides a promising platform for responding to the relativist challenge to contextualism.

There is, of course, much more to explore with respect to this complicated empirical domain, but this discussion shows that, insofar as it accounts for the guiding observation, prospective contextualism provides substantial new resources to the defender of a contextualist theory of conversation against relativist challenges.[48]

Let me close this section by addressing one kind of disagreement which appears to pose a particular challenge to my view, brought to my attention by an anonymous referee for this journal: persistent modal disagreement. Suppose that Sue is convinced that the keys might be under the bed. Louise has already looked under the bed, and didn't find them. Sue doesn't trust Louise's ability to find things, however, and Louise knows this. Sue says, 'The keys might be under the bed!' Louise says, 'No, they can't be. I searched carefully.' Sue stands her ground: 'They still might be. You might have missed them!' Louise stands her ground: 'No, you're wrong. They can't be.' And so on. If Sue and Louise are suitably conflictual, we could imagine this going on for a while. The question for my account is what Louise is doing here. If Louise *knows* that Sue is going to stubbornly stand her ground, then she knows that the prospective common ground will remain consistent with the keys being under the bed. But then she knows that what she is asserting is false. Why is it nonetheless acceptable for Louise to go on asserting that the keys can't be under the bed?

Here (following a suggestion by the same referee) it is helpful to compare the situation to one of practical negotiation. Suppose that Mark is convinced that John will not clean the rabbit cage: John never has in the past, is chronically disobedient, and so on. We may be inclined to say that in this situation he knows that John will not clean the rabbit cage. It somehow nevertheless seems permissible for him to say to John 'You *will* clean the rabbit cage tonight'. For the sake of practical negotiation, it seems that Mark is able to suspend his disbelief and operate under the assumption that there is at least a possibility that this time will be different (whether this is a rational strategy or pathology no doubt depends on the particular relationship). I suspect that something similar is going on in the case of Sue and Louise. Despite in some sense knowing that Sue will remain intransigent, it seems that Louise is able to suspend her disbelief and operate under the assumption that there is some chance that Sue will change her position—and thus some chance that what Louise is saying will turn out to be true. (I don't have a theory of exactly why this kind of suspension of disbelief seems to be acceptable in cases like this; but the comparison with the practical case suggests that it is indeed possible.)

6 Conclusion

Prospective contextualism makes sense of the fundamental dynamics of epistemic modality within the bounds of contextualism. This shows that, *pace* much recent work, the dynamics of epistemic modality do

[48] One area for further exploration comes from very interesting recent work in Beddor and Egan 2018, which suggests that the question under discussion in the context of assessment plays a key role in how speakers interpret epistemic modal claims.

not present an insoluble challenge to the contextualist framework: we can preserve contextualism's elegant model of communication as information transfer, and its corresponding division of labor between semantics and pragmatics, while making sense of the guiding observation about how speakers use modal claims to negotiate about what possibilities to treat as live.

That we *can* do so does not mean that we *should*. I have addressed an obvious objection to prospective contextualism, arguing that the performative structure which prospective contextualism attributes to modal claims is widely attested in natural language, a fact which helps us make sense of the norms of modal assertions within the prospective contextualist framework. I have also brought out some attractive features of prospective contextualism. But my main goal has not been to argue that prospective contextualism is the correct model of epistemic modality, but rather to show that contextualism cannot be rejected on the grounds of its putative inability to capture the guiding observation, and that the resulting version of contextualism—prospective contextualism—is well worth serious further exploration as a theory of epistemic modality. Whether prospective contextualism is the *correct* theory of epistemic modal claims depends on a range of further questions involving the semantics and pragmatics of epistemic modals. I have highlighted two already: the first concerns the subtle embedding behavior of epistemic modals; the second, the relativist challenge to contextualism. A third concerns the pragmatic role of 'must'. I have focused here on 'might'-claims, but 'must'-claims have peculiarities which are too often neglected in the discussion of epistemic modals, and which must be accounted for by any theory of epistemic modals.[49] I hope to take up these issues at greater length in future work.

In concluding, I would like to suggest that the framework I have sketched here can be generalized to analyze not just epistemic modals but also a wide variety of other constructions that we use to negotiate matters that depend, in part or in whole, on parameters of the context. The extension to probability modals like 'Probably', which have much in common with epistemic modals, will be straightforward: these can be seen as making claims about what the contextually provided measure structure will be like at the prospective time; provided the identity of that measure structure depends on what it is commonly accepted to be, then probability modals can be used to negotiate about what measure structures to coordinate on. Deontic modals can receive a similar treatment. The standard theory of deontic modals faces a similar puzzle to that raised at the outset for the standard theory of epistemic modals: how do speakers use deontic modals to negotiate about what to do? If I say we should go to Chinese, and you disagree, we are clearly not disagreeing about what norms are accepted by me, nor about what norms are accepted by everyone. Instead, it looks like we are disagreeing about what norms *to accept*. We can capture this by saying that we are making assertions about what norms will be accepted at the prospective time; provided the identity of those norms depends on what is commonly accepted about them, deontic modals can be used to negotiate about what norms to coordinate on.[50]

A similar treatment may be available for a variety of other phenomena, such as negotiation about stan-

[49] See e.g. Karttunen 1972; von Fintel and Gillies 2010; Lassiter 2016; Mandelkern 2017 for discussion.

[50] As an anonymous referee for this journal helpfully points out, Khoo and Knobe (2018) likewise advocate a contextualist approach to normative negotiation and disagreement in which it is crucial that 'certain assertions can be regarded as proposals to update the norm parameter'. I think that an approach like the one I am suggesting would fit naturally into the overall picture they advocate. They suggest a different approach, which essentially supervaluates over contexts instead of talking about future contexts; as far as I can tell, the resulting picture is very similar to the picture that would result from spelling out the idea I am suggesting here.

dards of vagueness, matters of taste, and performatives of the kind discussed in Austin 1979, like 'I promise'. It is tempting to think that phenomena involving negotiation of the sort these involve take us beyond the contextualist framework, since in these cases we seem not to be *describing* how things are vis-à-vis some contextual parameter, but rather *proposing* how they ought to be. But the contextualist framework can make sense of this if we take these constructions to describe the way that parameter *will be*. Provided that the identity of the parameter depends in the right way on what the interlocutors accept it to be, constructions which describe how the parameter will come to be will amount to performative proposals about how the parameter ought to be.

Appendices

A The contextualist framework

Here I give an explicit model of the contextualist theory of communication, and verify that this model, together with the prospective contextualist account of the assertoric content of modal claims, predicts that an assertion of ⌜Might p⌝ is a proposal to make p compatible with the common ground, and an assertion of ⌜Must p⌝ a proposal to make p entailed by the common ground. What follows is just one possible implementation of the contextualist model; subtle questions of implementation, however, do not matter for present purposes.[51]

We can model contextualism's common ground at t with a consistent, logically closed set of sentences Γ_t in a standard modal language. As above, I use $\Box_t$ as a modal operator interpreted as 'it is common ground at t that', and $\Diamond_t$ as its dual. We read off the common ground at a time t as follows: for any sentence p, p is common ground at t iff $\Box_t p \in \Gamma_t$. In order to encode the monotonicity assumption implicit in the contextualist model—the assumption that conversants proceed as if they will gain, rather than lose, information as conversation goes on—we stipulate that whenever $\Box_{t'} p$ is in $\Gamma_{t'}$, so is $\Box_{t''} p$ for all t'' after t'.

With this background, we can say that, in the contextualist framework, an assertion of p at time t is a proposal which, if accepted, has the effect of ensuring that $\Box_{t'} p$ is in $\Gamma_{t'}$, where t' is the time at which p is either accepted or rejected (i.e. the *prospective time*, relative to that assertion). We further stipulate that $\Gamma_{t'}$ is the smallest set which includes Γ_t, updated with the fact that the assertion in question has been made and anything that follows from this (i.e. with any information which is *accommodated*; see Stalnaker 2002; von Fintel 2008); which includes $\Box_{t'} p$; $\Box_{t''} p$ for all t'' after t'; and is logically closed. Some revision of the common ground may be necessary to ensure that this set is consistent (as when, for instance, p is common ground and then ⌜Not p⌝ or ⌜Might not p⌝ is asserted); I assume this revision comes in in the accommodation step. The mechanics of how such a revision goes is beyond our scope here: although this topic is of special interest for the theory of epistemic modals, since assertions of epistemic possibility claims are a standard tactic for creating such a clash, the underlying theory of how such a revision goes is just an instance of the general problem of belief revision, and need not be addressed by our theory of modals.[52] We assume, finally, that if an assertion is rejected (in the sense we are interested here), this is equivalent to its

[51] The contextualist framework is sometimes glossed as an *intersective* update model: when p is asserted at a context set c, the subsequent context set is $c \cap p$. This is a reasonable approximation, but it ignores a feature of updating that is crucial to our present purposes, and that is built into the orthodox framework: namely, that when one updates with p, one *also* updates with the proposition that p is commonly accepted, and so on; these updates are critical so that the logic of common ground is preserved across updates. This is accounted for in the present, slightly more complicated, implementation of contextualism.

[52] See C.E Alchourrón and P. Gärdenfors and D. Makinson 1985 for a classic reference.

negation being accepted.

Because of the meaning assigned to modal claims by prospective contextualism, we can immediately regiment them in our formal language: 'might' and 'must' will just be treated as $\Diamond_{t'}$ and $\Box_{t'}$, respectively, where t' is the prospective time. Thus our logic for modal claims—and our formal model for updates with them—will follow immediately from our model for updates in general. A claim of ⌜Might p⌝ as asserted at t will thus be regimented as $\Diamond_{t'}p$, where t', again, is the prospective time relative to the assertion (again, what time that actually amounts to will depend on concrete features of the actual speech situation). Suppose first that this claim is accepted. Then $\Box_{t'}\Diamond_{t'}p$ is in $\Gamma_{t'}$. As proved in §3, it is a theorem of the logic of common ground that $\Box_{t'}\Diamond_{t'}p \rightarrow \Diamond_{t'}p$. From this and the fact that $\Gamma_{t'}$ is logically closed, it follows that $\Diamond_{t'}p \in \Gamma_{t'}$. Under the assumption that $\Gamma_{t'}$ is always consistent (assuming that some tacit adjustment happens when this would otherwise not be the case), it follows that the $\Box_{t'}\neg p \notin \Gamma_{t'}$, and thus $\bar{p}$ is not common ground at t', and thus that p is compatible with the common ground at t'.

Suppose second that this claim is rejected. Then $\Box_{t'}\neg\Diamond_{t'}p$ is in $\Gamma_{t'}$; equivalently, $\Box_{t'}\Box_{t'}\neg p \in \Gamma_{t'}$. As proved in §3, it is a theorem in the logic of common ground entails that $\Box_{t'}\Box_{t'}p \rightarrow \Box_{t'}p$. From this and the fact that $\Gamma_{t'}$ is logically closed, it follows that $\Box_{t'}\neg p \in \Gamma_{t'}$, and thus that $\bar{p}$ is common ground at t', and p is not compatible with the common ground at t'.

Since ⌜Must p⌝ will be regimented as $\Box_{t'}p$, these considerations also show that an assertion of ⌜Must p⌝ is a proposal which, if accepted, ensures that p is in the common ground; and, if rejected, ensures that $\bar{p}$ is compatible with the common ground.

B Semantics of a fragment

With g_c a variable assignment initialized by the context c and κ a local context, I propose the following semantics (see Mandelkern 2019 for extended discussion):

(1) $[\![\text{Might}_i\ \text{p}]\!]^{c,\kappa,w}$

a. defined only if $\forall w' \in \kappa : g_c(i)(w') \subseteq \kappa$.

b. Where defined, true iff $\exists w' \in g_c(i)(w) : [\![\text{p}]\!]^{c,\kappa,w'} = 1$.

In unembedded environments, κ is initialized as the context set. For unembedded modals, then, ⌜Might p⌝ will be true just in case p is compatible with a subset of the context set selected by the context. I assume as a pragmatic default this subset will be the worlds compatible with the prospective common ground, yielding the prospective contextualist assertoric content. (Modal claims which are interpreted as quantifying over a set of worlds outside of the global context, as in 'exocentric' readings, will be interpreted as associated with a covert operator which has shifted their local context.)

When 'might' is embedded, κ may be shifted, explaining some of the peculiarities of how we interpret embedded modals. I propose that we follow the symmetric version of Schlenker (2009)'s theory of local contexts to spell this out. To get a sense of how this would work, consider these entries for connectives:

(2) $[\![\text{p and q}]\!]^{c,\kappa,w}$=1 iff $[\![\text{p}]\!]^{c,\kappa\cap q,w}$=$[\![\text{q}]\!]^{c,\kappa\cap p,w}$=1.

(3) $[\![\text{Not p}]\!]^{c,\kappa,w}$=1 iff $[\![\text{p}]\!]^{c,\kappa,w}$=1.

(4) $[\![\text{If p, then q}]\!]^{c,\kappa,w}$=$[\![\text{q}]\!]^{c,\kappa\cap p,w}$.[53]

Note that these entries suffice to guarantee that sentences like ⌜p and might not p⌝, and ⌜Might p and not p⌝, are contradictions (false at every world where defined). This explains their infelicity when embedded under

[53] This assumes that the consequents of conditionals always contain modals whose interpretation is sensitive to their local context. Note that this need not be an epistemic modal, though.

'Suppose' and 'If' (whatever semantics we adopt for this operators), answering a challenge raised by Yalcin 2007 for the contextualist treatment of modals.

References

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

Anscombe, G. (1963). *Intention*. Cornell University Press.

Austin, J. (1962). *How To Do Things With Words*. Oxford University Press.

Austin, J. L. (1979). Performative utterances. In *Philosophical Papers*. Oxford University Press.

Bach, K. (1975). Performatives are statements too. *Philosophical Studies*, 28:229–36.

Bach, K. (1994). Conversational impliciture. *Mind and Language*, 9:124–162.

Bach, K. and Harnish, R. (1979). *Linguistic Communication and Speech Acts*. MIT Press.

Beaver, D. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications: Stanford, CA.

Beddor, B. and Egan, A. (2018). Might do better: Flexible relativism and the qud. *Semantics & Pragmatics*.

C.E Alchourrón and P. Gärdenfors and D. Makinson (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530.

Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2009). Attention! *Might* in inquisitive semantics. In Cormany, E., Ito, S., and Lutz, D., editors, *Semantics and Linguistic Theory (SALT)*, volume 19, pages 91–108.

Cohen, L. J. (1964). Do illocutionary forces exist? *The Philosophical Quarterly*, 14:118–137.

Condoravdi, C. and Lauer, S. (2011). Performative verbs and performative acts. In Reich, I. et al., editors, *Sinn und Bedeutung*, volume 15, pages 149–164, Saarbrücken, Germany. Saarland University Press.

Davidson, D. (1979). Moods and performances. In Margalit, A., editor, *Meaning and Use*, number 3 in Sythese Language Library. Springer, Dordrecht.

Dowell, J. (2011). A flexibly contextualist account of epistemic modals. *Philosophers' Imprint*, 11(14):1–25.

Eckardt, R. (2009). Integrated speech acts. Manuscript, University of Göttingen.

Eckardt, R. (2012). *Hereby* explained: an event-based account of performative utterances. *Linguistics and Philosophy*, 35(1):21–55.

Eckardt, R. and Schwager, M. (2009). How to do things with words. ESSLLI 2009, Bordeaux.

Egan, A. (2007). Epistemic modals, relativism, and assertion. *Philosophical Studies*, 113(1):1–22.

Egan, A. (2011). Relativism about epistemic modals. In Hales, S., editor, *Blackwell Companion to Relativism*, pages 219–241. Wiley-Blackwell.

Egan, A., Hawthorne, J., and Weatherson, B. (2005). Epistemic modals in context. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy: Knowledge, Meaning and Truth*, chapter 6, pages 131–169. Oxford University Press.

von Fintel, K. (2008). What is presupposition accommodation, again? *Philosophical Perspectives*, 22(1):137–170.

von Fintel, K. and Gillies, A. (2007). An opinionated guide to epistemic modality. In Gendler, T. and Hawthorne, J., editors, *Oxford Studies in Epistemology 2*, pages 32–62. Oxford.

von Fintel, K. and Gillies, A. (2008). CIA leaks. *Philosophical Review*, 117(1):77–98.

von Fintel, K. and Gillies, A. (2010). Must...stay...strong! *Natural Language Semantics*, 18(4):351–383.

von Fintel, K. and Gillies, A. (2011). 'Might' made right. In Egan, A. and Weatherson, B., editors, *Epistemic Modality*. Oxford University Press.

Groenendijk, J. and Stokhof, M. (1975). Modality and conversational information. *Theoretical Linguistics*, 2(1-3):61–112.

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

Hacking, I. (1967). Possibility. *Philosophical Review*, 76(2):143–168.

Karttunen, L. (1969). Pronouns and variables. In Binnick, R. et al., editors, *Papers from the Fifth Regional Meeting of the Chicago Linguistic Society*, pages 108–116, Chicago. Chicago Linguistic Society.

Karttunen, L. (1972). Possible and must. In Kimball, J., editor, *Syntax and Semantics*, volume 1, pages 1–20. Academic Press, New York.

Karttunen, L. (1974a). Presupposition and linguistic context. *Theoretical Linguistics*, 1(1-3):181–194.

Karttunen, L. (1974b). Presupposition and linguistic context. *Theoretical Linguistics*, 1(1-3):181–194.

Khoo, J. (2015). Modal disagreements. *Inquiry*, 58(5):511–534.

Khoo, J. and Knobe, J. (2018). Moral disagreement and moral semantics. *Noûs*, 52(1):109–143.

Knobe, J. and Yalcin, S. (2014). Epistemic modals and context: Experimental data. *Semantics and Pragmatics*, 7(10):1–21.

Kölbel, M. (2009). The evidence for relativism. *Synthese*, 166:375–95.

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. (2012). Conditionals. In *Modals and Conditionals*, pages 86–108. Oxford University Press.

Lackey, J. (2007). Norms of assertion. *Noûs*, 41(4):594–626.

Lasersohn, P. (2009). Relative truth, speaker commitment, and control of implicit arguments. *Synthese*, 166(2):359–374.

Lassiter, D. (2016). *Must*, knowledge, and (in)directness. *Natural Language Semantics*, 24(2):117–163.

Lewis, D. (1979). A problem about permission. In Saarinen, E. et al., editors, *Essays in Honor of Jaakko Hintikka*. D. Reidel, Dordrecht.

MacFarlane, J. (2011). Epistemic modals are assessment sensitive. In Egan, A. and Weatherson, B., editors, *Epistemic Modality*. Oxford University Press.

MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford University Press, Oxford.

Mandelkern, M. (2017). A solution to Karttunen's problem. In Truswell, R., Cummins, C., Heycock, C., Rabern, B., and Rohde, H., editors, *Sinn und Bedeutung 21*, pages 827–844.

Mandelkern, M. (2019). Bounded modality. *The Philosophical Review*, 181(1).

Mandelkern, M. (To appear). What 'must' adds. *Linguistics and Philosophy*.

Moore, G. (1962). *Commonplace Book*. Routledge.

Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81.

Ninan, D. (2010). Semantics and the objects of assertion. *Linguistics and Philosophy*, 33(5):355–380.

Price, H. (1983). Does 'probably' modify sense? *Australasian Journal of Philosophy*, 61(4):396–408.

Roberts, C. (1998/2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6):1–69.

Rothschild, D. (2011). Expressing credences. In *Proceedings of the Aristotelian Society*, volume 112, pages 99–114.

Rothschild, D. (2013). Do indicative conditionals express propositions? *Noûs*, 47(1):49–68.

Schlenker, P. (2009). Local contexts. *Semantics and Pragmatics*, 2(3):1–78.

Searle, J. (1968). Austin on locutionary and illocutionary acts. *The Philosophical Review*, 77:405–424.

Searle, J. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press.

Silk, A. (2015). What normative terms mean and why it matters for ethical theory. In Timmons, M., editor, *Oxford Studies in Normative Ethics*, volume 5, pages 296–325. Oxford University Press.

Silk, A. (2018). Weak and strong necessity modals. In Dunaway, B. and Plunkett, D., editors, *Meaning,*

Decision, and Norms: Themes from the Work of Allan Gibbard.
Stalnaker, R. (1970). Pragmatics. *Synthese*, 22(1-2):272–289.
Stalnaker, R. (1974). Pragmatic presuppositions. In Munitz, M. K. and Unger, P., editors, *Semantics and Philosophy*, pages 197–213. New York University Press, New York.
Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and semantics*, volume 9, pages 315–322. New York: Academic Press.
Stalnaker, R. (1993). A note on non-monotonic modal logic. *Artifical Intelligence*, 64(2):183–96.
Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language and Information*, 7(1):3–19.
Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25(5-6):701–721.
Stalnaker, R. (2014). *Context*. Oxford University Press.
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, editor, *Sinn und Bedeutung*, volume 11, pages 583–97.
Stojanovic, I. (2007). Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy*, 30:691–706.
Stojnić, U. (2016). One's *Modus Ponens*: Modality, coherence, and logic. *Philosophy and Phenomenological Research*, 93(3):1–48.
Strawson, P. F. (1964). Intention and convention in speech acts. *The Philosophical Review*, 73:439–460.
Swanson, E. (2011). How not to theorize about the language of subjective uncertainty. In Egan, A. and Weatherson, B., editors, *Epistemic Modality*. Oxford University Press.
Swanson, E. (2015). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, 45(121):121–146.
Szabó, Z. G. (2016). Prospective interpretation. *Philosophical Studies*.
Teller, P. (1972). Epistemic possibility. *Philosophia*, 2(4):303–320.
Truckenbrodt, H. (2009). Performatives and agreement. Manuscript, ZAS Berlin.
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.
Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
Wittgenstein, L. (2001/1953). *Philosophical Investigations*. Wiley-Blackwell, 3rd edition.
Yablo, S. (2014). *Aboutness*. Princeton University Press.
Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.
Yalcin, S. (2011). Nonfactualism about epistemic modality. In Egan, A. and Weatherson, B., editors, *Epistemic Modality*, pages 295–332. Oxford University Press.
Yalcin, S. (2012). Bayesian expressivism. *Proceedings of the Aristotelian Society*, CXII(2):123–160.
Yalcin, S. (2015). Epistemic modality *de re*. *Ergo*, 2(19):475–527.