

SHIFTY CONTEXTUALISM ABOUT EPISTEMICS

CALEB PERL

Shandong University

According to a highly natural, orthodox view, epistemic modals like *might* and *must* are contextually variable, allowing us to express different propositions in different contexts of utterance. This view (contextualism about epistemic modals) is the orthodox one because the only other ways of making sense of how epistemic expressions are sensitive to information (views like relativism, expressivism, and dynamicism) carry such unorthodox commitments. Yet it has faced more than its share of challenges. In this paper, I will argue that two important challenges for contextualism about epistemic modals receive the very same solution: one problem about disagreement, and one problem about the reasonableness of our epistemic beliefs. The first of these challenges is very familiar, and the second less so, but equally important.

IMAGINE Ignorant Ian and Well-Informed Wendy looking for some keys.

Ian: The keys might be in the drawer.

Wendy: What you said is not true—I've already checked.

There is a highly natural, orthodox view about epistemic modals like *might* and *must*: that the modals are contextually variable, allowing us to express different propositions in different contexts of utterance.¹ But it has been difficult for the orthodox view to capture all the facts about our use of epistemic modals.

There are two central facts that have been especially difficult to capture. The first central fact is that it can be appropriate for Wendy to disagree with Ian, to say that what he said is not true, or to say that he was wrong, or that he made a mistake. (Imagine that she knows that the keys aren't in the drawer.) The second central fact is that Wendy can acknowledge that Ian is reasonable even while insisting that he's wrong. After all, Ian might not know whether the keys are in

Contact: Caleb Perl <chperl@gmail.com>

1. Angelika Kratzer (1977; 1981; 2012) has been particularly influential in developing this approach.

the drawer. If he doesn't, it's appropriate and reasonable for him to think that the keys might be in the drawer.

It is hard for contextualists to explain both features of the example. They can explain why the two agents are reasonable in making their utterances. The two agents are simply describing different bodies of information. Wendy is describing one that represents the contents of the drawer, and Ian is describing one that doesn't. But this suggestion seems to get disagreement wrong. If the two parties are only describing different bodies of information, they wouldn't be disagreeing. If you said *I'm hungry*, and I said *I'm not hungry*, we wouldn't be disagreeing. We're just talking about different people. Something similar for the contextualist: the two parties would just be talking about different bodies of information. It's then hard for the contextualist to explain why Wendy can take herself to disagree with Ian's reasonable belief.²

I will argue that these problems arise because we haven't thought about the relationship between modals and attitude reports in a systematic enough way. This diagnosis shouldn't be surprising: disagreement about what might be comes from our attitudes about what might be. I will defend this diagnosis by developing a new contextualist account of epistemic expressions and showing how this new account builds from one systematic picture of the relationship between attitudes and modals. This new account crisply explains how Wendy and Ian can both be reasonable while also disagreeing with each other. But my goal isn't just to show how contextualists can solve these two problems. Many philosophers have the impression that contextualists have to make strange and unmotivated claims to solve these problems. I want to undermine this impression. When we think systematically enough about what contextualists should say about attitude to modals, we see that contextualists do have strikingly principled answers to these questions. I will close with a further illustration of this point. I show how

2. There has been an explosion of work on this challenge—articulations of important theoretical approaches and arguments for and against. There isn't always perfect clarity about the relations between these different approaches, as Mark Schroeder (2015c) describes in detail. But broadly expressivist approaches (or approaches that draw heavily on expressivist resources) include at least Ernest Adams (1975), Joshua Knobe and Seth Yalcin (2014), Ben Lennertz (2013; 2014), Jonathan Bennett (2003), Dorothy Edgington (1995), Martin Montminy (2012), Sarah Moss (2013; 2015), Benjamin Schneider (2010), Eric Swanson (2006), and Seth Yalcin (2007; 2011). Broadly relativist approaches include at least Andy Egan (2007), Andy Egan, John Hawthorne, and Brian Weatherson (2005), John MacFarlane (2007; 2011; 2014), and Tamina Stephenson (2007). Broadly contextualist approaches include at least Pranav Anand and Valentine Hacquard (2013), Kent Bach (2011), Gunnar Björnsson and Alexander Almér (2010), Gunnar Björnsson and Stephen Finlay (2010), David Braun (2012), Herman Cappelen and John Hawthorne (2009), Ezra Cook (2013), Janice Dowell (2011; 2013), Angelika Kratzer (1977; 1981; 2012), Kai von Fintel and Anthony Gillies (2008; 2011), John Hawthorne (2007), Torfinn Huvenes (2015), Justin Khoo (2015), and Crispin Wright (2007). Broadly dynamic views include at least Frank Veltman (1996) and Malte Willer (2013; 2014).

this new contextualist account crisply explains another important phenomenon: the “epistemic contradictions” that Seth Yalcin (2007) has highlighted.

1. The New Theory

But I will build to this general argument gradually. I will start only by introducing my new contextualist theory, and showing how it works. I’ll wait until Section 2 to show how it follows from a systematic picture of attitudes about what might be.

The new theory starts from the basic contextualist idea. Epistemic use of modals express propositions about some further argument, a *modal base*, which is a body of information. But my new contextualist theory relies on two fairly novel resources. One resource is a new property which I will call *epistemic correctness*: the property of containing all truths. A modal base is epistemically correct if it contains all truths. And one modal base is more epistemically correct than another if the first contains more truths. The second resource is a distinctive propositional attitude. Suppose you know that only one of p_1 , p_2 , and p_3 is true. You think think p_1 is more likely: p_1 is .4 likely, p_2 .3 likely, and p_3 .3 likely. Suppose further that the best course of action depends on which of these three propositions is true. In deciding on the best course of action, you can reasonably treat p_1 as true. In doing that, you’re disposed to use it as a premise in further reasoning. (You’ve got to act as though one is true, and it’s the best candidate.) When you do that, I will say that you’re *accepting* p_1 . Importantly, though, acceptance is a quite weak attitude. It is less demanding, for example, than *belief*. After all, p_1 is only .4 likely—so it’s not plausible that you believe it.

I propose that might-judgments involve *accepting* a modal base as *epistemically correct*. For example, asserting an epistemic is a hybrid state, consisting of two attitudes. It both consists in asserting something about what’s compatible or incompatible with a particular modal base, and also consists in accepting that modal base as epistemically correct. Here’s an example.

In asserting that the keys might be in the drawer, I’m

- asserting that some modal base m_1 is compatible with the keys being in the drawer, and
- accepting that m_1 is epistemically correct.

Other attitudes are also a hybrid of two attitudes. Belief and knowledge, for example, both are.

If you *believe* that the keys might be in the drawer, you

- *believe* that \mathbf{m}_1 is compatible with the keys being in the drawer, and
- accept that \mathbf{m}_1 is epistemically correct.

If you *know* that the keys might be in the drawer, you

- *know* that \mathbf{m}_1 is compatible with the keys being in the drawer, and
- accept that \mathbf{m}_1 is epistemically correct.

1.1. How the Hybrid Theory Explains Both Disagreement and Reasonableness

The hybrid theory explains both of the central desiderata that we want to explain. It explains how two agents can *disagree* while both being *reasonable*. Say that Ignorant Ian and Well-Informed Wendy are looking for some keys.

Ian: The keys might be in the drawer.

Wendy: No, they can't be there. (I've already checked.)

Modal bases are, let's say, sets of propositions. Ian is talking about \mathbf{m}_1 , and Wendy is talking about \mathbf{m}_2 .

$\mathbf{m}_1 = \{\text{the keys are not in the garage, the keys are not in the kitchen}\}$

$\mathbf{m}_2 = \{\text{the keys are not in the garage, the keys are not in the kitchen, the keys are not in the drawer}\}$

The hybrid theory then makes these two predictions.

Ian asserted that \mathbf{m}_1 is compatible with the keys being in the drawer, while accepting that \mathbf{m}_1 is epistemically correct.

Wendy asserted that \mathbf{m}_2 rules out the keys being in the drawer, while accepting that \mathbf{m}_2 is epistemically correct.

Then Wendy and Ian disagree *in what they accept*. They are both accepting propositions that can't both be true. The proposition that \mathbf{m}_1 is correct and the proposition that \mathbf{m}_2 is correct can't both be true, since it can't be that both contain *all* truths. That's the simple account of *disagreement*. How can the two agents nonetheless be reasonable? Reasonable belief is also a hybrid state, consisting of acceptance of some modal base as epistemically correct, plus reasonable belief about that modal base. It crucially does *not* require reasonable belief about what's epistemically correct. So Ian is reasonable because he has reasonable beliefs about what is compatible with his own modal base, the one that leaves open whether the keys are in the drawer.

Now the hybrid theorist is *expanding* the kind of disagreements that the contextualist can explain. She is suggesting that our two agents can be disagreeing about what's epistemically correct. Importantly, though, she still allows for the sort of disagreement that more traditional contextualists can explain. Suppose that two mathematicians are disagreeing about what follows from a set of axioms. One might assert that p must be true, while the other asserts that p might be false. In this case, they would be disagreeing about what's compatible with some modal base: the set of axioms. Moreover, we can evaluate the reasonableness of the different beliefs about what's compatible with the axioms. We might deny that a mathematician is reasonable in believing that p must be true if we don't think that the mathematician has a proof of p . Traditional contextualists have an elegant explanation of this kind of case, and the hybrid theory retains that elegant explanation.

The hybrid account of disagreement generalizes straightforwardly to all cases of disagreement, including eavesdroppers. Imagine that Ian is talking with Wendy. Eavesdropping Evelyn isn't part of the conversation, but she knows where they are.

Ian: The keys might be in the drawer.

...

Eavesdropping Evelyn [to a friend]: No, they can't be there. (I've already checked.)

Many contextualists struggle to explain this exchange, but the hybrid theorist explains it crisply. The hybrid theorist takes disagreement to be disagreement about what's epistemically correct. But she thinks of epistemic correctness in a striking and distinctive way. It's a comparative notion: some modal base is more correct than some group of bases. Which ones? *Any modal base whatsoever*. Let's illustrate with Ian. He's thinking about some modal base m_1 . In accepting that m_1 is epistemically correct, he's taking it to be more epistemically correct than any other modal base **whatsoever**. Now m_1 is incomplete. It doesn't say where they keys are. Imagine that Evelyn does know where the keys are: say, that they're in the drawer. Then she accepts a modal base m_2 that includes that proposition, as epistemically correct. That's why they disagree. Evelyn accepts that m_2 is more correct than m_1 , and Ian accepts that m_1 is more correct than m_2 .

This proposal works only because it takes Ian to be doing something very odd. He's accepting a proposition that he's in a position to know is false. Ian knows that either the keys are in the drawer or they're not. He also knows that if the keys are in the drawer, m_1 is not epistemically correct, because m_1 plus the proposition about the keys is *more* correct. And he knows that some other modal base is more correct if they're not in the drawer. So Ian is in a position to know that it's false that m_1 is epistemically correct.

In general, though, it can be perfectly appropriate to accept propositions that you can know to be false. Here are two examples. First, you can accept that Hamiltonian mechanics is true and relativist mechanics is false to simplify your calculations in an engineering context. In doing that, you're accepting a proposition that you know to be false. Second, some religious people will insist that God loves his children, while also insisting that God isn't male. Since *his* presupposes that the salient individual is male, they're accepting a proposition that they don't believe. In this case, they're accepting it because they regard it as the least unsatisfactory option. Other possessive pronouns (in English) presuppose that the individual is female, or that the group has several members.

The central feature of my account of epistemics is that epistemics pattern with the use of Hamiltonian mechanics, or gendered language for God. In both cases, your practical interests can make it perfectly appropriate to accept propositions you know to be false. The practical interest in the epistemic case is to reason from the largest amount of information that you have, without relying on anything that your information leaves open. The latter goal is why ordinary agents only accept an incomplete modal base as epistemically correct.

1.2. *Contrasting the Hybrid Theorist with Other Contextualists*

The hybrid theorist strictly improves on the approaches that other contextualists have developed. To see why, imagine that Wendy told us (1).

(1) Ian is heading towards that drawer, because he reasonably believes that the keys might be there. But I know that they can't be there; I've already checked. So it would be a mistake to follow Ian in believing that the keys might be there.

On the hybrid suggestion, uses of (1) communicate these three propositions.

- Ian is heading towards that drawer, because he (reasonably believes that m_1 is compatible with the keys being there and accepts that m_1 is epistemically correct).
- But I, Wendy, know that m_2 is incompatible with the keys being there and accept that m_2 is epistemically correct.
- So Ian is wrong in (believing that m_1 is compatible with the keys being there and accepting that m_1 is epistemically correct).

Different modal bases are incorporated into the complements of Ian's knowledge reports than are incorporated into the complements of Wendy's knowledge reports. Why? Because the two agents take different modal bases to be epistemically correct.

This suggestion contrasts with an important contextualist suggestion from Janice Dowell (2011; 2013). She holds that, when one person assertively utters an epistemic and her interlocutor assertively utters its negation, the two parties are disagreeing because they're expressing propositions about the *very same* modal base. Dowell gives a powerful defense of this proposal, suggesting that speaker intentions determine the modal base incorporated into the propositions expressed, and that speakers can intend to talk about the same modal base even if they don't know exactly what modal base they're talking about. We can simplify the comparison with Dowell by pretending that speakers intend to talk about whatever modal base is *epistemically correct* in the present context. Dowell thinks that they're talking about whatever base is in fact epistemically correct, while the hybrid theory takes them to be talking about different particular modal bases which they accept to be epistemically correct.

Does Dowell's view differ substantively from the hybrid theory? Yes. The two make different predictions about the reasonableness of judgments about what might be. The hybrid theorist predicts that it's really easy to have reasonable beliefs about what might be. Reasonable belief depends only on reasonable belief about what's compatible with a set of propositions that you're already thinking about. Dowell does not make it as easy to have reasonable beliefs about what might be. Since Ian and Wendy disagree, the two attitude reports should express propositions about the very same modal base. But it's hard to see how Ian could have reasonable beliefs about that modal base, unless he is sensitive to the information that Wendy takes to bear on the question. And it's unclear how he could be sensitive to that information. He isn't aware of what it is!

We can appreciate another distinctive feature of the hybrid picture by comparing it with a proposal by Kai von Fintel and Anthony Gillies (2011). Properties like epistemic correctness play no role in their proposal. They instead propose that utterances of epistemic modals communicate a *cloud* of propositions, propositions about several different bodies of information. One of those propositions describes only the speaker's body of information; another describes the body of information of the entire group. They suggest that it's appropriate to assert an epistemic modal when you accept one of those propositions—typically, the one about your body of information. But your assertion puts the other propositions in play. So other speakers can react to your utterance by reacting to some of the other propositions put in play. This proposal explains why it's easy for a speaker to appropriately assert an epistemic. It's appropriate as long as her own information is compatible with the prejacent. But it also explains why her hearers can object. They're objecting to the other propositions put in play.

Their proposal doesn't generalize to the attitude reports under discussion.

(1) Ian is heading towards that drawer, because he reasonably believes that the keys might be there. But I know that they can't be there; I've

already checked. So it would be a mistake to follow Ian in believing that the keys might be there.

What proposition does Ian believe? It seems like he only believes the proposition about his own body of information. (Ian may have no beliefs about any other salient bodies of information; he may suspend judgment on them.) But then it's unclear what Ian is mistaken in believing. He is not mistaken in believing the proposition about his own body of information. That proposition is true. But that proposition may be the only relevant proposition he believes. Since \lceil Ian is wrong in believing $S \rceil$ entails that Ian believes S , it's hard to see how von Fintel and Gillies can explain this case.³ The hybrid theory does much better.⁴

The new contextualist account makes it much easier to see how someone can have reasonable but mistaken beliefs about what might be. However, there is a danger that the account makes it *too* easy to have mistaken beliefs about what might be. Virtually all 'might' beliefs involve accepting an incomplete modal base as epistemically correct. (Ordinary agents aren't omniscient, and mostly know that they aren't.) But we saw in the last section that no incomplete modal base could be epistemically correct. So the account seems to predict that *all* ordinary 'might' beliefs are mistaken. And this prediction would be very costly. For one thing, I don't regard my own 'might' beliefs as mistaken, whatever this account might predict.⁵

Fortunately, though, the proffered account does not make this prediction. Distinguish two pictures of a mistake in acceptance. On a *belief-centric* picture, you can appropriately say \lceil it would be a mistake to accept $p \rceil$ if you yourself *believe* $\neg p$. On an *acceptance-centric* picture, you can appropriately say \lceil it would be a mistake to accept $p \rceil$ if you yourself *accept* $\neg p$. The acceptance-centric picture turns out to disarm the problem in the previous paragraph, while also being significantly more plausible than the belief-centric alternative.

I'll develop the difference between the belief-centric picture and the acceptance-centric picture with an example with gendered religious language. Suppose that Bob and Carl both use gendered pronouns for God, because they regard those pronouns as the least inadequate option. But they believe that God is neither male nor female. Now since *his* presupposes that the salient individual is male, an utterance of (2) communicates the proposition (+).

3. Ben (Lennertz 2013: Chapter 2, Section 5) has emphasized this kind of problem.

4. It's also important to compare the hybrid theory with other similar contextualist approaches. Several contextualists supplement their approach an expressivist pragmatics: examples include Stephen Finlay and Gunnar Björnsson (2010), Ben Lennertz (2013), and Martin Montminy (2012). The hybrid approach is importantly different from those approaches, but it's more useful to contrast the two approaches late in the paper, in Section 4.2.

5. I'm grateful to an anonymous *Ergo* referee for posing this objection in a particularly forceful way.

(2) It would be a mistake to follow Carl in believing that God infuses grace in his children.

(+) It would be a mistake to follow Carl in (believing that God infuses grace in God's children and accepting that God is male)

The belief-centric picture predicts that Bob can appropriately say (2) even if Bob agrees that God infuses grace in God's children. After all, Bob doesn't *believe* that God is male. (He believes that God is neither male nor female.) Since he doesn't believe the proposition Carl accepts, the belief-centric account predicts that (2) is appropriate. And that prediction is incorrect. Bob is *ex hypothesi* perfectly happy to use gendered pronouns for God—he doesn't disagree with Carl about his use of pronouns. In order for (2) to be appropriate, Bob and Carl have to disagree about the infusion of grace. The acceptance-centric picture, by contrast, predicts that (2) is *inappropriate*. That picture grounds mistakes in acceptance in disagreement in acceptance. Since Bob and Carl both accept that God is male, Bob can appropriately say (2) only if he disagrees with Carl about whether God infuses grace in God's children.

And our initial problem about ordinary 'might' beliefs does not arise on the acceptance-centric conception. I grant that we should *believe* that no incomplete modal base is epistemically correct. (The last page gave a compelling argument for that point.) But that belief doesn't require us to *accept* that no incomplete modal base is epistemically correct. On the contrary, we have very good reason to accept incomplete modal bases as epistemically correct. And in diagnosing mistakes in acceptance, I reason from the propositions that I myself accept. Consider someone who is less informed than me. The modal base that they accept to be correct includes fewer propositions than mine does. So I can reasonably take them to accept something false, since, given the propositions that I myself accept, one of the propositions that they accept is false. Reasoning from what I accept also explains why I won't regard myself as mistaken about what I accept. Given the propositions that I myself accept, all the propositions that I accept are true.⁶

6. This observation explains why it's Moore-paradoxical to assertively utter *the keys might be there but I am wrong in so believing*. In asserting the first conjunct, you're asserting that some modal base \mathbf{o}_1 is compatible with the keys being there and accepting that \mathbf{o}_1 is epistemically correct. And in asserting the second conjunct, I'm asserting that I'm wrong in (believing that \mathbf{o}_1 is compatible with the keys being there and accepting that \mathbf{o}_1 is epistemically correct). When I'm evaluating whether it's a mistake to accept that \mathbf{o}_1 is epistemically correct, I reason from the propositions I accept, which includes the proposition that \mathbf{o}_1 is epistemically correct. Since that proposition trivially entails itself, it's not a mistake to accept that \mathbf{o}_1 is epistemically correct. But then the only way for the second conjunct to be assertable is for me to not believe that \mathbf{o}_1 is compatible with the keys being there. So the assertion has Moore-paradoxical assertability-conditions.

Now there are lots of fruitful and interesting questions to explore about the acceptance-centric conception, and we'll take up some of them as the paper goes on. The important point for now is that that conception has the right structure for a Goldilocks solution to the contextualist's problem: not too few mistaken beliefs about what might be, and not too many, either.

2. A Systematic Picture of Attitudes towards Epistemics

The broader goal of this paper is to connect the contextualist's problems about disagreement to more general questions about attitudes about what might be. The last section showed that it's *possible* for contextualists to explain epistemic disagreement, while also explaining reasonable epistemic belief. The explanation draws on two distinctive tools: the property *being-epistemically-correct*, and the attitude of *acceptance*. But it's reasonable to worry that these tools are both *ad hoc*—that is, unmotivated additions to the basic contextualist idea. This section introduces a systematic account of attitudes about what might be. It shows how those two tools are essential parts of that systematic account. That is, it shows how thinking more systematically about attitudes about what might be illuminates the questions in the last section. After doing this, I'll explain how the account introduced in Section 1 bears a range of other exciting fruit, include an explanation of the "epistemic contradictions" that Seth Yalcin (2007) has highlighted.

2.1. A Central but Underappreciated Question

My systematic account answers a central but underappreciated question about the epistemic use of modals. This question *is* genuinely underappreciated. So I will begin by arguing for the importance of this question, and arguing for it in some depth.

This central question is about *attitudes* towards epistemics. It's an instance of a more general question about attitudes towards modals, like (3).

(3) Jill believes that she must keep this promise.

(3) can be used to talk about the demands of morality, the demands of prudence, or the demands of some particular set of laws. It's standard to explain this point by supposing that *she must keep this promise* has one more argument place than its surface syntax suggests. The moral use of (3) communicates a proposition with that argument place saturated with something that represents the demands of morality, the prudential use with something that represents the demands of prudence, and so on.

My central question arises *after* we explain what all the uses of (3) have in

common. It asks what *distinguishes* the different uses of (3). What distinguishes the proposition the moral use communicates from the proposition the legal use communicates? A NAIVE THOUGHT is that they're distinguished because the propositions believed have different truth-conditions.

NAIVE THOUGHT: The moral use of (3) communicates that Jill believes a proposition with different truth-conditions than the proposition that the legal use of (3) takes her to believe.

And this NAIVE THOUGHT is pretty natural. Jill's moral belief is different from her belief about the laws, because the two propositions have different modal profiles. The moral proposition is true because of what the moral facts are, and the legal proposition is true because of what the legal facts are. And the moral and legal facts don't always line up.

Unfortunately, though, the NAIVE THOUGHT doesn't work in full generality. It is possible for the laws of some society to coincide exactly with the laws of morality, across all of modal space, even while someone gets confused, and thinks that the laws don't coincide. Imagine that some set of laws does coincide with morality in just that way. It's still possible for Jill to think that she's *morally* required to keep this promise, without thinking that she's *legally* required to. In that case, the moral use of (3) asserts something true, and the legal use asserts something false. But the NAIVE THOUGHT can't explain this fact. In this case, the moral proposition is true at a world w iff the legal proposition is true at w . There isn't any difference in modal profile for the NAIVE THOUGHT to latch onto to explain how the moral but not the legal use of (3) asserts something true.

Let's make this point more concrete. We might model the demands of morality with orderings on worlds: w_1 is ranked higher than w_2 , which is tied with w_3 , which are both above... And I'm imagining a case where the very same ordering that models the demands of morality also models the laws of some society. (The society designed its laws so that they correspond with morality throughout all of modal space.) Call that ordering \mathbf{o}_1 . Given the NAIVE THOUGHT, the moral and legal uses of (3) both communicate that Jill believes that \mathbf{o}_1 requires her keeping this promise. And that won't work, as it stands. The moral use asserts something true and the legal use asserts something false.

As a result, we need to draw on richer resources than the NAIVE THOUGHT uses. We need to explain how Jane can be thinking of \mathbf{o}_1 in different ways, as the demands of morality, or the demands of positive law. I will call the different ways we think about the orderings *modes of presentation*. You might think of these modes of presentation as properties that we predicate of orderings. Or you might think of them as the concepts that we use to cognize the orderings. Or you might think of them as grist for Stalnakerian diagonalization.⁷ Then we can say that

7. In setting aside the NAIVE THOUGHT, I am not taking a stand on the nature of

the moral use of (3) asserts that Jane believes that \mathbf{o}_1 (thought of as the demands of morality) requires her keeping this promise, which can be true even if she doesn't think that \mathbf{o}_1 (thought of as the demands of the law) requires her keeping this promise.

Let's return to epistemic modals. I've just given an argument that it's necessary to posit modes of presentation to distinguish the moral use of modals from the legal use. Once we're familiar with this kind of argument, we're in a position to appreciate how to generalize it for the epistemic use. The argument requires two different uses of modals: in the initial case, moral and legal uses. So I need a use of modals that pairs with the epistemic use, to power my argument. I will use what I'll call the *circumstantial* use of modals. We can say *the tree can grow to be ten feet tall* to indicate something about its nature, that it's compatible with its nature that it grows to be that tall. Given this use, we can construct the same sort of case constructed earlier. It's possible for someone's information to coincide exactly with the facts that are relevant for the circumstantial use. Say that John finds himself in that situation. There is still a difference between epistemic belief ascriptions and circumstantial belief ascriptions. (4a) can be appropriate even if (4b) isn't.

(4a) John believes that this tree might grow to be ten feet tall.

(4b) John believes that this tree can grow to be ten feet tall.

John is aware of the information that he has. He knows that that information is compatible with the tree growing to be ten feet tall. (4a) is then perfectly appropriate. But he doesn't realize that that information coincides with the facts about the nature of the tree. In fact, he suspends judgment about the nature of the tree. Then (4b) isn't appropriate.

The best way to explain this contrast is to posit a mode of presentation for the circumstantial use, and a mode of presentation for the epistemic use. (4a) and (4b) both communicate propositions about the same body of information—the same modal base, to use a phrase from Kratzer.⁸ But (4b) is *inappropriate* even though (4a) is appropriate, because appropriate use of (4b) requires John to think of that body of information *as* being about the nature of the tree. The explanation is exactly like the explanation of other instances of Frege's puzzle.

propositions—whether they're individuated more finely than sets of worlds. I'm just pointing out that you need to draw on whatever tools you use to capture hyper-intensional differences in your account of modals.

8. This assumption is very important for this style of argument. And I haven't defended it here. A full defense of it would need to show that there are only two unpronounced argument-places: one for something like a modal base, and one for something like an ordering or an ordering source. But I take Angelika Kratzer (1977; 1981; 2012) to have given powerful considerations in favor of that claim, and for those considerations to have been borne out in the subsequent literature.

John thinks that Hesperus is bright can be appropriately assertable even if *John thinks that Phosphorus is bright* isn't. Why? The use of *Phosphorus* suggests a different mode of presentation than the use *Hesperus* does. So we should posit a mode of presentation to distinguish the epistemic use of modals.⁹

Appreciating this point puts us in a better position to appreciate the hybrid theory from Section 1. That theory was using a particular account of the mode of presentation that distinguishes the epistemic use. It appealed to a modal base being *epistemically correct*—that is, as containing all truths. It uses that property as the mode of presentation for the epistemic use. (4a) is appropriate because the modal base that John accepts to be epistemically correct is compatible with the tree growing that tall. Now remember that the overarching goal of this section is to show that the hybrid theory is highly principled, and not *ad hoc*. I'll pursue that goal by showing that the hybrid theory is a simple and natural answer to what I will call the CENTRAL QUESTION.

CENTRAL QUESTION: What mode of presentation is associated with the epistemic use of modals, and how is it associated with that use?

Features of the hybrid account that initially look *ad hoc* do not look the same way when we see them as answering this question.

The appeal to epistemic correctness is a natural idea about the mode of presentation that distinguishes the epistemic use. And we should be willing to entertain this idea, once we see that we do need some such mode. (The conclusion will briefly consider possible modes of presentation that distinguish the epistemic use.)

2.2. *Motivating the Rest of the Hybrid Theory*

The rest of the hybrid theory follows from a very simple conjecture about our CENTRAL QUESTION. It holds that the mode of presentation of the epistemic use is the property *being-epistemically-correct*, and that that property is part of a *presupposition* of a use of an epistemic.

What are presuppositions? I won't say much here. I'll just understand the presuppositions of a sentence as information that the sentence (i) is somehow associated with, but (ii) is interpreted as being part of the background assumptions in the conversation. It's easier to appreciate this notion from particular

9. Other accounts of the difference between (4a) and (4b) might tempt you. You might note the syntactic differences that Valentine Hacquard (2006) notes between epistemic and root modality, and take those syntactic differences to make some kind of semantic difference. I'm happy to allow for this point. Even granting it, though, why need some account of why (4a) can be appropriate even while (4b) isn't. And I'm claiming that the best account of that difference appeals to modes of presentation, which may be interestingly informed by syntactic considerations.

examples.

- (i) It wasn't Tom who stole the tarts.
- (ii) John was also in NY last night.
- (iii) John has stopped beating his dog.

Each of these sentences are somehow associated with two propositions, with one proposition understood as foregrounded and at-issue, and the other as backgrounded and not-at-issue. (i) is associated with the backgrounded proposition that someone did steal the tarts. You wouldn't use (i) unless you thought that someone stole the tarts, but the main point of using (i) is to communicate that Tom didn't steal the tarts. The other examples are similar—(ii) is associated with the backgrounded proposition that someone else was in NY last night, and (iii) is associated with the backgrounded proposition that John used to beat his dog.¹⁰

The hybrid theorist conjectures that the property *being-epistemically-correct* is part of the presupposed, not-at-issue content associated with bare epistemic utterances.

- (5) The keys might be in the drawer.

AT-ISSUE: m_1 is compatible with the keys being in the drawer

NOT-AT-ISSUE: m_1 is epistemically correct

This conjecture is a simple empirical hypothesis. It's definitely not the only viable hypothesis about epistemic modals under attitude reports, or the only way to answer our CENTRAL QUESTION, not even when we appeal to epistemic correctness. But if this empirical hypothesis is right, the hybrid account of epistemic modals from Section 1 follows immediately.

The hybrid theory follows immediately because of a highly general fact about presuppositions under attitude reports: that acceptance is the attitude appropriate for presuppositions of the complement. I will spend some time belaboring this general fact, because it's so important for evaluating the hybrid theory. Consider some uncontroversial examples of presupposition triggers, like a cleft sentence like *it wasn't Tom who took out the trash*. If you assertively utter that sentence, you presuppose that someone took out the trash, while asserting that Tom didn't take out the trash. In general, an utterance with the form \lceil it wasn't F who G-ed \rceil presupposes that someone G-ed. An interesting pattern emerges when cleft sentences are embedded in attitude ascriptions.

- (6a) John **hopes** that it wasn't Tom who took out the trash.
- (6b) John **fears** that it wasn't Tom who took out the trash.
- (6c) John **conjectures** that it wasn't Tom who took out the trash.

10. The characterization of (ii) isn't quite right; see Saul Kripke (2009) for a more careful account.

(6a) can be true even if John doesn't hope that someone took out the trash. (Maybe he knows that someone took out the trash, and regrets that the trash was taken out.) The other examples are the same. (6b) can be true even if John doesn't fear that someone took out the trash; and (6c) can be true even if John doesn't conjecture that someone took out the trash. In each case, the sentences are appropriate only if John has some attitude towards the proposition that someone took out the trash. In general, $\lceil A \text{ V-es that } S \rceil$ can be true even if A does not V the presuppositions of S's complement.¹¹

We should note another important fact about these examples. It's normally odd to use (6a), (6b), or (6c) unless John is in some sense *aware* that someone took out the trash. The best explanation of this oddity is that attitude reports usually communicate that the matrix subject *accepts* the presupposition of the complement. That is, the matrix subject has to be treating the presupposition of the complement as true for some reason. (Robert Stalnaker 2002 has influentially emphasized this point.) Now you might wonder if I've failed to notice a stronger generalization: that that matrix subject needs to *believe* the presupposition of the complement. This alternative is not liberal enough. Suppose that we're trying to figure out who would have taken out the trash, if someone did take out the trash. You can temporarily assume that someone did take it out, to simplify the conversation that follows. Once you've made that assumption, (6d) is perfectly appropriate.

(6d) You **know** that it wasn't Tom who took out the trash.

In this case, though, you don't *believe* that someone took out the trash. You're just accepting it as a way to figure out who would have taken out the trash,

11. Now you might wonder whether these examples illustrate the point I take them to. In general, it's possible to hope in some content without hoping in everything that the content entails. You know that Bob is coming to the party, and think him a bore—you'd rather that he not be there. But Mary would redeem Bob's presence. (Mary's riffs on Bob's personality are *really* funny.) *You hope that Bob and Mary come* can then be true, even though *you hope that Bob comes* isn't. Maybe (6a) is appropriate for the same reason. (6a)'s complement really does express a conjunctive content. It's just that you don't need to hope in both conjuncts for (6a) to be appropriate.

Maybe that's all these examples show—that you can hope in a content without hoping in everything the content entails. Clefts express conjunctive propositions: that someone took out the trash, and that Tom didn't. So there isn't anything special about these clefts; they just illustrate the general point that you can hope in a content without hoping in all its entailments. But that diagnosis is implausible, because it overgenerates. Suppose that John hopes, fears, or conjectures that someone took out the trash. The present idea predicts that (6a)–(6c) should be appropriate then. After all, he's hoping, fearing, or conjecturing one of the two conjuncts, and that's imagined to be enough to make the attitude reports appropriate. But it's very hard to understand the attitude reports as appropriate for that reason. The present idea can't be sustained.

if someone did. (6d) is appropriate even in that case. So we should follow Stalnaker in being more liberal about the attitude that the matrix subject has towards the presupposition of the complement. The matrix subject need not believe the presupposition of the complement in order for (6d) to be appropriate.¹²

To summarize:

- you can *accept* a proposition without believing it, or knowing it,
- the matrix subject of an attitude report $\lceil A \text{ V-es that } S \rceil$ need not also V presuppositions of the complement,
- but many attitude reports *do* communicate that the matrix subject accepts the presuppositions of the complement.

I take these generalizations to be constraints on a full theory of presupposition; any such theory will vindicate these three generalizations. It is unfortunately difficult to explain why these generalizations hold.¹³ But the important point for my purposes is *that* these generalizations are true, not why they are true.

The goal of this section was to motivate the hybrid theory which associates attitude reports like (7) with the following pair of attitudes.

(7) Jane thinks that the keys might be in the kitchen.

- Jane thinks that \mathbf{m}_1 is compatible with the keys being in the kitchen, and
- Jane accepts that \mathbf{m}_1 is epistemically correct.

Two features of the hybrid view were particularly apt to look *ad hoc*: the appeal to epistemic correctness, and the appeal to acceptance. We've now seen the simple conjecture that motivates these two features. Epistemic correctness is introduced as the mode of presentation that distinguishes beliefs about what epistemically might be, from other kinds of modal beliefs. And the hybrid theorist is someone who conjectures that that mode of presentation is a *presupposition* of utterances with epistemic modals. The hybrid theory is not *ad hoc*. Instead, it is one the simplest and most conservative ways of extending a basically Kratzerian understanding of modals to answer a question that everyone needs to answer.

12. There is broad agreement that attitude reports communicate something about the matrix subject's attitude towards presuppositions; Paul Elbourne (2005), Bart Geurts (1998), Irene Heim (1992), and Lauri Karttunen (1974) all agree on this basic point. They don't consider the choice between the attitude being *belief* and its being *acceptance*. So I take the present Stalnakerian consideration to be compelling enough to motivate the conclusion that it's acceptance and not belief that matters.

13. For more empirical details about these generalizations and steps towards explaining them, see Bart Geurts (1998), Irene Heim (1992), Caleb Perl (2017b), Judith Tonhauser, David Beaver, Craige Roberts, and Mandy Simons (2013)

2.3. Consequences of this Explanation of the Hybrid Theory

This presuppositional explanation of the hybrid theory has an important advantage. To appreciate this advantage, let's recall one reason the hybrid theory initially looked implausible. That theory takes ordinary talk about what might be to involve acceptance of a proposition that the speakers can know to be false. Asserting that the keys might be in the drawer means asserting that your own body of information \mathbf{m}_1 is compatible with the keys being in the drawer, while also *accepting* that \mathbf{m}_1 is epistemically correct. In accepting that \mathbf{m}_1 is epistemically correct, you're accepting that it's *more* correct than any modal base whatsoever. Now you're not omniscient. So you're in a position to know that the proposition that you accept is false. Why? Let \mathbf{m}_1 be the modal base that you accept to be epistemically correct. Since you're not omniscient, \mathbf{m}_1 is incomplete. So there is some way of extending it to be more complete: there's some p where adding either p or $\neg p$ would make it more complete. If p is true, \mathbf{m}_1 plus p is more epistemically correct. And if $\neg p$ is true, \mathbf{m}_1 plus $\neg p$ is more epistemically correct. \mathbf{m}_1 can't be epistemically correct; in order for it to be, it would have to be more correct than every other modal base.

This feature of the hybrid account cleanly explains eavesdroppers. It's always reasonable for a better-informed eavesdropper to disagree with you, because they reject the proposition you accept: that \mathbf{m}_1 is epistemically correct. But this feature of the hybrid account might look implausible. Ordinary discourse about what might be just doesn't involve this sort of commitment! I appealed to two analogies to make the idea more plausible. One was about religious language: religious people who use a gendered pronoun to talk about God, even though they don't believe that God is male. I suggested that they're accepting that God is male, for conversational purposes. Now I take pronouns like *he* and *him* to trigger the presupposition that the salient individual is male.

$\llbracket \text{he} \rrbracket = \lambda f_{\langle et \rangle}$

AT-ISSUE: x is f

NOT-AT-ISSUE: x is male

An utterance of *God loves his children* presupposes that God is male.¹⁴ It's nonetheless intelligible for a religious person to make that utterance, even if they don't believe this presupposition. They only need to *accept* it for their conversational purposes.

The hybrid theorist takes the epistemic use of modals to involve accepting a proposition you don't believe. And the analogy with gendered pronouns is quite

14. See Robin Barwise and Cooper (1981), Philippe Schlenker (2003), and Yasutada Sudo (2012) for a defense of the claim.

tight. *The two phenomena have exactly the same linguistic explanation: they are both due to the fact that you only need to accept presuppositions of your utterance.* Though this picture of the epistemic use is initially surprising, it ends up having a lot going for it. The epistemic use involves temporarily treating a body of information as the only truths. That's why any eavesdropper will disagree with you. They don't treat that body of information as the only truths. But you yourself have great *practical* reasons for treating that body of information as the only truths. You have practical reason for not relying on any more complete body of information: any such body of information will contain a proposition that you're not sure about.

My presuppositional explanation of the hybrid theory makes the theory significantly more plausible. It makes the theory much more simple and elegant. And it shows that natural language already has the resources that the hybrid theorist uses in explaining the data: the analogy with gendered pronouns is exact. But the presuppositional explanation gives rise to a further question. I've emphasized that the hybrid theory can explain *disagreement* about what might be. And given the presuppositional explanation of the hybrid theory, disagreement about what might be is usually disagreement in presupposition. We're disagreeing about which modal base is epistemically correct. But you might wonder if disagreement in presupposition licenses the full range of disagreement markers that are in fact licensed. After all, it's very odd to use certain locutions if you want to reject the presuppositions of an utterance.

(8) It wasn't Bill who stole the tarts.

(9) ?? [No/that's wrong], because the tarts weren't stolen.

No and *that's wrong* seem appropriate only as objections to what (8) asserts (that Bill didn't steal the tarts), not as objections to the presupposition. Epistemics pattern quite differently. *No* and *that's wrong* are both perfectly appropriate.

(8') The keys might be in the drawer.

(9') [No/that's wrong.] They can't be there.

Someone who assertively utters (9') disagrees with someone who assertively uttered (8'). And that disagreement is possible even if the two agents agree about what's compatible with the salient bodies of information. The hybrid theorist takes them to be disagreeing in presupposition—the second speaker is rejecting the presupposition that the first speaker has made. But it doesn't seem like presuppositions can bear the weight that the hybrid theorist needs. You can't say *no* or *that's false* if you reject the presupposition.

This worry confuses a feature of some presupposition triggers with a feature of all presupposition triggers. We should distinguish two subclasses of presupposition triggers: those whose appropriateness depends on their already being

satisfied in the initial context, and those whose appropriateness doesn't. It-clefts fall in the first subclass. It's inappropriate to say (8) totally out of the blue, if we don't already accept that someone stole the tarts.¹⁵ But other presupposition triggers fall into the other class, and can be used even if the initial context doesn't satisfy the presupposition. For example, you can use *start* to introduce new information: I can say that John started smoking even if my hearers are ignorant about whether he smoked in the past.¹⁶ Crucially, though, you can use our locutions to object to the presupposition that *start* triggers.

(8'') John started dancing at noon.

(9'') [No he didn't/that's wrong]—he's been dancing since ten.

This difference illustrates a general pattern. Locutions like *that's wrong* can target the presupposition trigger iff the presupposition trigger can be felicitously used to introduce new information. It's not yet clear why the two phenomena go together—why you can use *that's wrong* to reject presuppositions if the presupposition trigger can be used to introduce new information. But the generalization does seem robust.¹⁷

Now the hybrid theorist's presupposition about epistemic correctness can definitely be used for new information: the hearer may not realize that some new modal base is epistemically correct until the speaker says her bit. It's thus unsurprising that you can use locutions like *that's wrong* to reject the presupposition.

2.4. General Problems about Agreement

Let's close by considering a problem about *agreement*. Suppose that I know only that the keys aren't in the kitchen or garage, and you know only that they aren't in the dining room. Then it's true that

(10) You and I agree that the keys might be in the bedroom.

As a first pass, the hybrid theory seems to take (10) to communicate that

15. Saul Kripke (2009) and Judith Tonhauser et al. (2013) both helpfully discuss this point.

16. Explanations of this point are highly fraught—some theorists explain it by supposing that it's easier to *accommodate* the presupposition of *start* than the presupposition of a cleft; Robert Stalnaker (2002) is an influential example. Other theorists do not think that that is the best strategy: skeptics include Mandy Simons, Judith Tonhauser, David Beaver, and Craig Roberts (2011), Kai von Stechow (2008), and Christopher Gauker (2008). Importantly, though, the parties all agree about the data that matter here—that it's easier to use *start* to introduce new information.

17. You can verify this point for the triggers that Judith Tonhauser et al. (2013) call anaphoric triggers (of which clefts are one example): you can't use them to introduce new information, but neither can you target the presuppositions with locutions like *that's wrong*.

(10f) You and I agree that \mathbf{m}_x is compatible with the keys being in the bedroom, and accept that \mathbf{m}_x is epistemically correct.

But there is no modal base \mathbf{m}_x such that you and I both accept that \mathbf{m}_x is epistemically correct. If \mathbf{m}_x contains the proposition that the keys aren't in the dining room, I won't accept that \mathbf{m}_x is correct. (I don't have a high enough credence in that proposition.) Similarly if \mathbf{m}_x contains the proposition that the keys aren't in the garage. Then you won't accept \mathbf{m}_x as correct. (You don't have a high enough credence in that proposition.) The hybrid theory doesn't seem capable of explaining why (10) is true.

This problem parallels a central problem about anaphora in attitude reports. Solutions to that central problem are also solutions to this problem. Suppose that I think that Miley is the only kitten who killed a mouse, and I hope that Miley didn't torture it first. You think that *Tabby* is the only kitten who killed that mouse, and hope that Tabby didn't torture it first.

(11) You and I agree that a kitten killed that mouse, and we both hope that it didn't torture the mouse first.

This example involves linked attitudes about the same object—agreement plus hope. It's natural to think that the two attitudes are linked because they are both singular attitudes about the same object.

(11f) $[\exists x: \text{kitten}(x)]$ (You and I agree that x killed that mouse, and we both hope that x didn't torture the mouse first)

But this suggestion is mistaken for exactly the same reason that the initial hybrid gloss on (10) is mistaken. No assignment of values to variables verifies the constraints in (11f). You don't think that Tabby did the killing, and I don't think that Miley did it.

We need to draw on some other formal tool to explain how (11) is true. One option is to appeal to the sort of *unanchored discourse referents* that Irene Heim (1982) and Hans Kamp (1981) have influentially pioneered. Functionally, those discourse referents are *denoting concepts* that play a similar role to Fregean individual concepts. They allow individuals to think and talk about individuals without requiring the individuals to be able to think singularly about those individuals. Where i is a term for an unanchored discourse referent $\lceil A V_1$ -es that i is F and $A V_2$ -es that i is $G \rceil$ is true iff, in every maximal consistent extension of the propositions that $A V_1$ -es and the propositions that $A V_2$ -es, there is some individual that is both F and G . (11) is then true. Let F abbreviate "kill the mouse" and G abbreviate "torture the mouse". I think that Miley F -ed and I hope that Miley didn't G . Every relevant maximal consistent extension of my attitudes includes the two embedded propositions. At every such extension, then, there's an

object that F-ed and didn't G: Miley. You think that Tabby F-ed and you hope that Tabby didn't G. So at every maximal extension of your attitudes, there's an object that F-ed and didn't G: Tabby. Appealing to unanchored discourse crisply explains how (11) is true.

Unanchored discourse referents can crisply explain why (10) is true.

(10) You and I agree that the keys might be in the bedroom.

Where *i* is a term for an unanchored discourse referent:

(Official LF) You and I agree that *i* is compatible with the keys being in the bedroom, and accept that *i* is epistemically correct.

Then (10) is true. Every consistent extension of my beliefs includes that \mathbf{m}_m is compatible with the keys being in the bedroom and \mathbf{m}_m is epistemically correct. So every consistent extension verifies the constraints on the discourse referent—(10) is true of me. And every consistent extension of *your* beliefs includes that \mathbf{m}_y is compatible with the keys being in the bedroom and \mathbf{m}_y is epistemically correct. So every consistent extension verifies the constraints on the discourse referent—(10) is true of you too.

This explanation preserves the two central advantages of the hybrid account. It gets disagreement about what might be exactly right. If my modal base includes the negation of the prejacent and your modal base doesn't, we disagree. And in that case, (10) would be false. (There is no maximal extension of my attitudes where the modal base that is epistemically correct is compatible with the keys being in the bedroom.) At the same time, this explanation also preserves the central epistemic virtues of the hybrid theory. The rationality of attitudes towards unanchored discourse referents can bubble up from the rationality of the corresponding singular thoughts. Suppose that I'm rational in believing that Miley killed the mouse, and you're rational in believing that Tabby did it. Then (11') is true.

(11') You and I are rational in agreeing that a kitten killed that mouse, and also in hoping that it didn't torture the mouse first.

In other words, rational attitudes about particular objects (Miley and Tabby) ground rational attitudes about discourse referents. The appeal to unanchored discourse referents is fully compatible with the central epistemic virtues that the hybrid theorist claims.

Now there are many other ways of theorizing about intensional anaphora in examples like (11).¹⁸ But we should expect any adequate theory of intensional

18. Hawthorne and Manley (2012), for example, would favor a quite different approach.

anaphora to solve the problem for the hybrid theorist. Given the hybrid theory, the present problem *just is* a problem about intensional anaphora.¹⁹

19. For more detail, see (Perl 2017a: Chapters 9–11). This account also generalizes to other related problems. Can I say (a)?

(a) You're mistaken in believing that the keys might be in the bedroom.

It seems like (a) should be true, on the hybrid theory.

(LF1) You're mistaken in (believing that \mathbf{m}_y is compatible with the keys being in the bedroom and accepting that \mathbf{m}_y is epistemically correct).

It seems like I should be willing to assert (a), if the hybrid theory is right. By my lights, \mathbf{m}_y is not epistemically correct. But note that (a) might have another interpretation, given what we've said in this section. It can communicate something about unanchored discourse referents:

(LF2) You're mistaken in (believing that i is compatible with the keys being in the bedroom and accepting that i is epistemically correct).

And LF2 is *false*. Let F = compatible with the keys being in the bedroom, and G = epistemically correct. Every consistent extension of what I believe and accept includes the proposition that my own modal base \mathbf{m}_m is F and the proposition that \mathbf{m}_m is G. So every relevant maximal extension verifies the constraints on the discourse referent in LF2. As a result, it's not a mistake to believe that i is F and accept that i is G. In other words, LF2 is false, for reasons very much like the reasons why (10) is true.

The hybrid theorist makes the right prediction about (a) if she associates it with LF2. What should she say about the LF1 interpretation of the sentence? Should she say that (a) is ambiguous between a reading where the sentence is false (that is, LF2) and a reading where it is true (that is, LF1)? That's not very attractive. We tend not to hear an interpretation of (a) where it's true. She should rather hold that general constraints on interpretation cause us not to hear LF1. Consider a similar case. Suppose that you believe that the mouse-killing kitten is *not* a mouse-torturer. Then (b) is *false*.

(b) A kitten killed the mouse. You mistakenly think that it tortured the mouse first.

And (b) is false because you don't even have the relevant belief. In our present discourse-referent framework, we would say that you don't have the belief about the unanchored discourse referent that is the mouse-killing kitten that it is a torturer. But note that there is actually a reading of (b) where that supposition is true but the reading of (b) is true. Imagine I think that Miley killed and didn't torture the mouse. You think that Miley didn't kill the mouse but did torture it. You also think that Tabby killed the mouse but didn't torture it first. You do have the mistaken singular belief about Miley, that he tortured the mouse—and Miley is the individual that I regard as the mouse-killing kitten. And pronouns like *it* can be used to attribute belief in singular propositions. But somehow we don't hear that reading of (b), especially not given our initial supposition. That is, there are general constraints on interpretation that cause us to prefer the discourse-referent interpretation of (b). The hybrid theorist will hold that those same constraints are what explain why we prefer the LF2 reading of (a).

3. Contrast with Expressivism, Relativism, and Expressivist Contextualism

There has already been an enormous amount of literature on the epistemic use of modals.²⁰ One of the reasons that there is such a large literature is that the epistemic use raises important foundational questions in the philosophy of language. Some philosophers argue for heterodox approaches like *relativism* or *expressivism* as the best explanation of the epistemic use. They argue that more traditional approaches cannot explain it. My hybrid theory is a fully traditional approach. If it's right, the epistemic use of modals does *not* give any aid or comfort to heterodox views. I will spend quite a bit of time on this point, because it's central to the broader philosophical significance of the questions discussed in this paper.

A traditional conception of semantics takes its goal to be to articulate the assertability conditions of a sentence by articulating what it would be for it to be true. And it has a quite conservative understanding of truth. Sentences express propositions, and propositions are either monadically true, or true at worlds and/or times. There are two importantly different ways to reject this traditional approach. *Expressivists* articulate the assertability conditions of a sentence by articulating the mental state that makes it appropriate. That is one way to reject traditional views: to deny the starting assumption that the fundamental goal is to articulate truth-conditions. It is also possible to retain that starting assumption while departing from traditional semantic theories, by accepting a *relativist* conception of truth. The relativist allows that a proposition can be true when assessed from one point and false when assessed from a different point.

These two heterodox approaches are very interesting. They promise to illuminate foundational questions about truth and representation. And they have very interesting applications in other domains. Expressivism, for example, has very interesting *metaethical* applications. It transforms difficult questions about how we think about moral properties into more tractable questions about how we think about our own preferences. But these applications are interesting only if there is some reason to take these heterodox proposals seriously in the first place. And there is some reason to worry. Expressivists have the Frege-Geach problem, of giving a constructive account of embedded sentences.²¹ And there are important foundational challenges to the relativist's conception of truth.²²

Epistemic expressions are an important *license for optimism* on behalf of the expressivist and relativist. Both heterodox approaches have powerful and ele-

20. See the references in Footnote 2.

21. Important discussions of these problems include those by Peter Geach (1965), Bob Hale (1993), Mark Schroeder (2008; 2015b), John Searle (1962), and Nicolas Unwin (1999; 2001). Schroeder (2015) discusses these issues with particular attention to the epistemic case.

22. John Hawthorne and Herman Cappelen (2009) and Scott Soames (2011) both articulate some challenges.

gant explanations of how different agents can be reasonable in their epistemic beliefs but also be reasonable in disagreeing with others. The relativist predicts that it's easy for both agents to have reasonable beliefs about what might be: they each assess propositions about what might be given their own information, and it's comparatively easy to check what their own information says about the keys. But they also can intelligibly disagree: Wendy assesses Ian's proposition given her own information, and his proposition is false given her information. The expressivist has an elegant explanation of her own. Ian's reasonableness in saying that the keys might be in the drawer depends only on the mental state he's expressing in making that assertion. And Wendy and Ian disagree if the mental states they express aren't cotenable. And as emphasized earlier, it's initially hard to explain these facts in a traditional contextualist framework. This difficulty is some evidence that natural language does not work in the way the traditionalist supposes, but instead works in the way that the relativist or expressivist supposes. So it's a license for optimism that the problems for the heterodox can be solved.

The hybrid theory from Section 1 undercuts this license for optimism. It captures all the facts about disagreement and reasonableness that the heterodox could have reasonably taken to be a license for optimism. Now there are still other points that the heterodox can appeal to as licenses for optimism. (Maybe. The hybrid theory may also undercut those licenses for optimism as well, as noted later.) But the points about disagreement and reasonableness are among the cleanest and crispest licenses for optimism, so it's significant that the hybrid theorist captures them in a wholly traditional way.

3.1. Expressivist Contextualism

It's also illuminating to contrast the hybrid view with mixed views that draw on both contextualist and expressivist resources. They also take general difficulties about epistemic expressions as a license for optimism for the legitimacy of what they're doing. Work by Stephen Finlay and Gunnar Björnsson (2010) illustrate this general strategy.²³ They start by noting our practical reasons for using the most complete body of information available. The fuller body is less likely to lead us astray. So when Wendy says that Ian is wrong to think that the keys might be in the kitchen, she is recommending against relying on the information that Ian has. This account of disagreement is different than any account from previous sections. It's not disagreement about what's compatible with a modal base. But neither is it disagreement about which modal base has some further property, like being epistemically correct. Both sorts of disagreements are disagreements about which propositions are true. By contrast, Finlay and Björnsson are thinking

23. Ben Lennertz (2013) and Martin Montminy (2012) have made similar suggestions.

of the disagreement as practical disagreement: disagreement about what to do, rather than what to believe or what to accept.

This kind of proposal is viable only given an answer to the general problems that dog expressivists, including the Frege-Geach problem. Suppose I know that the cat isn't hungry. I can appropriately assert that

(12) Ian wrongly believes that the cat might be hungry.

In asserting (12), I need not be asserting that Ian is wrong about what is compatible with his information. (He and I can agree about that point, even as I insist on (12).) I can use (12) to assert that he's wrong in having the whole complex of attitudes that the theory associates with bare utterances of *the cat might be hungry*. This is the basic point that the expressivist contextualist needs to explain.

And the expressivist contextualist cannot explain this point, unless she has an answer to the Frege-Geach problem. To illustrate the problem, let's switch to a more complicated case. Suppose that some axioms of set theory T are provably compatible with some proposition p. Susan the set theorist knows about the proof. To express that T is compatible with p, she assertively utters *p might be true*. Now she and I agree that other mathematical considerations rule out p. (She's just using *p might be true* to express a fact about set-theory, in a context where it's clear that she's talking about set theory.) Doug the dogmatist disagrees with us. He does not acknowledge the other mathematical considerations that rule out p. He says *p might be true*, because he also knows that p is provably compatible with T. But he insists that T are the only acknowledged facts that bear on p. I'm willing to criticize Doug in a way I'm not willing to criticize Susan. Contrast this pair.

(13) Susan wrongly believes that p might be true.

(14) Doug wrongly believes that p might be true.

I'm willing to assert (14) but I'm not willing to assert (13). I'm asserting something different about Susan than I am about Doug. The challenge is to explain what that difference is.

It is easy to answer this challenge given a solution to the Frege-Geach problem. Expressivist contextualists associate an utterance U with a proposition P and some kind of commitment A.²⁴ U is appropriately assertable for me iff I have commitment A and I satisfy the norm of assertion for P. (*p might be true* is associated, say, with the proposition that \mathbf{m}_T is compatible with p and a preference for relying on \mathbf{m}_T .) Given the knowledge norm for assertion, say, it's appropriate to assert that sentence if you know that proposition and have that preference. Then

24. I'll use *association* as a fudge word for the mix of semantic and pragmatic processes that explain how P and A are relevant for our evaluation of uses of S.

we can give a crisp account of why (14) asserts more. (14) asserts that Doug has a preference for relying only on m_T , and (13) asserts nothing of the sort.

Utterances of (13) assert that (*) Susan believes P and it's wrong that P.
 Utterances of (14) assert that (**) Doug believes that P and A and it's wrong that P and A.

I'll accept (**), because I think it's wrong to prefer relying only on m_T . And I'll reject (*), because Susan and I both believe P.

Crucially, though, this answer works only given a solution to the Frege-Geach problem. In particular, the second conjunct of (**) is interpretable only given a solution to the Frege-Geach problem. It requires us to interpret commitments like A under a negation operator like *it's wrong that*. At this point, then, I've given one illustration of what it would take to answer my challenge. In fact, though, I think that this is the only way for an expressivist contextualist to answer this challenge. My challenge is to articulate what proposition is asserted by utterance of (13) and (14). I don't care how that happens. (I'll allow my opponent as much semantic blindness as she wants.) I just want to know what the difference is between what's asserted by utterances of (13) and utterances of (14).

3.2. *The Severity of this Problem*

The only way I see to answer this challenge without answering the Frege-Geach problem is to appeal to metalinguistic reasoning.²⁵ But we should trust that she can do that only when we understand what exact propositions she takes uses of (13) and (14) to assert. Let's start with one natural thought.

Utterances of (13) assert that (+) every context where Susan would assertively utter *it might be that p* is a context where one of the attitudes communicated is wrong.

Utterances of (14) assert that (++) every context where Doug *is talking about the modal base that he prefers* and where he would assertively utter *it might be that p* is a context where one of the attitudes communicated is wrong.

These propositions do capture the intuitive differences. I don't believe (+). After all, one of the contexts where Susan would assertively utter *it might be that p* is a context where she's not talking about the modal base she prefers. She's just talking about the axioms of set theory. And she acknowledges mathematical truths that go beyond those axioms. Moreover, I *do* believe (++) . Any context

25. Stephen Finlay (2014) seems to have defended something like this view, but has since abandoned it in favor of another approach.

where Doug is talking about the modal base that he prefers is a context where his preference is wrong. (There are other truths beside those axioms!) And we can evaluate these propositions without a solution to the Frege-Geach problem. Evaluating them just requires evaluating the unembedded preferences.

But this kind of proposal cannot work. It forces the wrong predictions about context sensitive vocabulary.

(15) Susan wrongly believes that I might be in Germany now.

Utterances of (15) assert that (+++) every context where Susan would assertively utter *I might be in Germany now* is a context where one of the attitudes communicated is wrong.

(+++) is true iff Susan has incorrect beliefs about *herself*. Her beliefs about me are totally irrelevant. Those truth-conditions are definitely wrong. (15) is about Susan's beliefs about *me*. It's not just that (15) has a reading that is about me: it's that (15)'s only readings are about me. None of them are about Susan. This attempt at a concrete, testable account of the proposition asserted by any of (12) through (14) definitely fails. And it's hard to see an attempt that does better.

The general problem is that allowing for metalinguistic readings of the complement predicts *too many* readings of sentences like (15). It's really hard to fix a theory that overgenerates. (For one thing, you can't appeal to pragmatic processes that add in additional content. The problem is that there's already too much content floating around.) I conclude that whatever is asserted by (12) through (14) can't be a metalinguistic proposition. But the kind of metalinguistic proposal is the only apparent way to explain the difference between (13) and (14), without needing an answer to the Frege-Geach problem. It's reasonable for us to conclude that the expressivist contextualist does need an answer to the Frege-Geach problem.

Now it is possible to draw on other heterodox resources to explain disagreement. I've been focusing on contextualists who draw on *expressivist* resources to explain disagreement. But it's also possible for a contextualist to draw on *relativist* resources to explain disagreement. It's helpful to see the recent suggestion by Stephen Finlay (2016) in this vein. On his proposal, (14) expresses the false proposition (++++):

(14) Doug wrongly believes that p might be true.

(++++) Doug wrongly believes that p is compatible with the propositions in \mathbf{m}_T .

(This proposition is false, because p is in fact compatible with the propositions in \mathbf{m}_T .) Finlay suggests that ordinary speakers are willing to assert (14), because

they think that there is some proposition that Doug is disposed to assert, and that is false relative to our information. It's helpful to see that proposition as the same sort of proposition that a relativist like John MacFarlane (2014) takes bare epistemic utterances to assert: a proposition that can be assessed as true relative to one point and false relative to another. One important difference between Finlay and MacFarlane is that MacFarlane thinks that these propositions exist and are genuinely assessment sensitive.

Finlay does not. He thinks that the commitment to these propositions is an error in ordinary epistemic discourse. Finlay's proposal raises a number of issues. But one important issue is that he is assuming that ordinary speakers are committed to assessing some propositions as true relative to their own body of information. And some critics of relativism doubt that it's even intelligible to do that. (Scott Soames (2011) is one example.) This criticism applies as much to Finlay as to MacFarlane. Those critics are challenging the intelligibility of the relativist explanation, not its correctness. And Finlay needs the explanation to be intelligible, even if he doesn't need it to be correct.

3.3. *How the Hybrid Theorist Gets this Data Right*

I've noted that expressivists and relativists both take the epistemic use of modals to be an important *license for optimism* on their behalf. They think that some sort of heterodox explanation has to be on the right track, because no traditional explanation will do. Expressivist contextualists might take themselves to have the same license for optimism. In order to explain disagreement and reasonable belief, you have to draw on some sort of heterodox resource. There might be interesting technical questions about how best to do that, but those technical questions shouldn't shake our confidence that it's possible.

The hybrid theory from Section 1 undercuts this license for optimism. So it makes the technical questions much more pressing. It eliminates the license for optimism that natural language works in a way that guarantees that those technical questions have a solution. The hybrid theory cleanly explains the data that expressivist contextualists don't explain. Why? She takes what the agent believes and what he accepts to both embed under *wrongly*.

(14) Doug wrongly believes that p might be true.

Utterances of (14) assert that [Doug [wrongly [believes that m_1 is compatible with p and accepts that m_1 is epistemically correct]]].

The relevant uses of (13) assert less.

(13) Susan wrongly believes that p might be true.

Utterances of (13) assert that [Susan [wrongly [believes that m_1 is compatible with p]]].

So we'll reject (13), in the story I gave, because Susan is right about whether m_1 is compatible with p . Doug is right about that point, too. But he's wrong in a further way. He accepts m_1 as epistemically correct. Now I haven't explained why (13) and (14) have such different readings. But I have answered the challenge that I've issued to the expressivist contextualist. I've explained how uses of (13) can assert something different from what uses of (14) would assert. This answer is available precisely because the hybrid theorist is so conservative. She doesn't draw on any heterodox resources (expressivist or relativist), so she doesn't need to solve the problems that dog those approaches. But I just don't see how expressivists or expressivist contextualists can clear even this low bar, absent a solution to the Frege-Geach problem.

And this lesson is important for other philosophical questions, too. Finlay and Björnsson are interested in epistemic expressions in part because they accept a broadly relativist conception of *moral* belief. Finlay, for example, defends an *end-relational* semantics for moral sentences, where they express propositions about our *ends*.²⁶ He allows that different agents may have different ends even when they're talking about what they morally ought to do. When I say *you ought to keep this promise* and you disagree, we may well be talking about different ends. And we may be making true claims about our different ends. But we'll continue to disagree morally even when we agree about what would promote various ends. So he needs some additional account of moral disagreement that goes beyond the propositions that are communicated. He fills that need with a broadly expressivist account that parallels his treatment of epistemics.

In fact, Finlay and Björnsson are explicit in taking their account of epistemic expressions as a license for optimism in defense of their relativism. They write: "since contextualism invokes the very same sort of pragmatic considerations (in particular an appeal to privileged contexts) to account for both sorts of context-insensitive assessments, no significant extra cost is added to defend standard relativity" (Björnsson & Finlay 2010: 36). The hybrid theorist undercuts this license for optimism. You don't need expressivist resources to defend the basic contextualist idea. There is no guarantee that the Frege-Geach problem can be solved. So *a fortiori* there's no guarantee that Finlay and Björnsson's metaethical relativism is defensible.

4. Conclusion

In fact, the comparison with Finlay and Björnsson gives us a more complete picture of the hybrid account this paper has introduced. Finlay and Björnsson implicitly adopt the same account of the epistemic use of modals as of the moral

26. Finlay (2014)

use. In both cases, the propositions communicated aren't enough to explain all cases of disagreement, so they also draw on expressivist tools. But this uniform treatment of the epistemic use and the moral use is somewhat unusual. It's more common for contextualists to treat the moral uses of modals quite differently than they treat the epistemic uses. To explain the moral use, they introduce a mode of presentation *being-morally-correct*, over and above the ordering or ordering source. This mode of presentation explains moral disagreement. Moral disagreement is disagreement about what is morally correct. But extant contextualists tend *not* to use the same strategy for epistemics. They don't rely on an additional property over and above the modal base. The whole situation is very odd. Contextualism is appealing as a way of unifying all the different uses with a common semantic core. But extant contextualists tend to use that common semantic core in a very different way in the epistemic case than in the moral case.

The hybrid theorist is distinctive, because she uses the same strategy in both cases. In that way, her approach is like Finlay and Bjornsson's. Unlike them, though, she is tacitly assuming an altogether traditional picture of moral disagreement. It's always disagreement about what is morally correct. The hybrid theorist generalizes this traditional picture *from* the moral case *to* the epistemic case. Epistemic disagreement is also disagreement about what's epistemically correct. There is, however, a good reason why most contextualists haven't extended the traditional strategy from the moral case to the epistemic case. It's hard to identify a property like *being-morally-correct* for the epistemic use that captures two central desiderata: first, that it covers all cases of disagreement, and, second, that it's something that ordinary agents can have reasonable beliefs about. Some philosophers have tried. Dan Lopez de Sa (2008), for example, has also defended contextualism by positing a new presupposition: the presupposition that the modal base represents what's common information to parties to the conversation. Alex Silk (2017) has suggested in a somewhat similar vein that a use of an epistemic modal presupposes that the salient modal base is *endorsed* in the context.²⁷ Both of these suggestions are closer to satisfying the second desideratum, that it be something that ordinary agents can have reasonable beliefs about.

Unfortunately, though, it's hard to develop these ideas to satisfy the first requirement, that it cover all cases of disagreement. As John MacFarlane (2007) emphasizes, Lopez de Sa's proposal struggles to explain *eavesdroppers*. Ignorant Ian asserts that the keys might be in the drawer. Someone much better informed hears what Ian said, but has decisive evidence that the keys aren't there. Our Eavesdropping Evelyn could intelligibly react by saying that Ian is wrong, that

27. And he intends this claim to dissolve arguments from disagreement, as it clearer in other work, like Silk (2016) and Silk (2017b).

the keys can't be there. It's very hard for Lopez de Sa to explain what they're disagreeing about. They're talking about different modal bases, so they aren't disagreeing about what's compatible with any one modal base. Moreover, Evelyn knows that they're not party to the same conversation. So it's not plausible that she is disagreeing with Ian about what's common information to parties in the conversation.²⁸

The hybrid theory, by contrast, does have an adequate account of eavesdroppers. But it has that adequate account because it abandons the second desideratum, that the property that covers all cases of disagreement is something that ordinary agents can have reasonable beliefs about. Eavesdroppers are disagreeing about what is epistemically correct. It's always intelligible for a better-informed eavesdropper to do that, given the hybrid theory. Someone who asserts an epistemic is accepting a proposition that they can know to be false: the proposition that their own incomplete modal base is epistemically correct.²⁹

The hybrid theory is like a three legged stool. Take away any leg, and the whole thing falls apart. First, disagreement about what might be is disagreement about what is more epistemically correct than anything whatsoever. Take away that leg, and you end up with problems in explaining all cases of disagreement. Second, *reasonable belief* about what might be is not reasonable belief about what is epistemically correct. Take away that leg, and you end up with problems in explaining how easy it is to have reasonable belief about what might be. Third, modals *presuppose* that the salient modal base is epistemically correct. Take away that third leg, and it's hard to see how the resulting view is coherent. (If you can disagree with a claim about what might be because the modal base isn't correct, why doesn't reasonable belief about what might be require reasonable belief about epistemic correctness?) Together, though, the three legs support a powerful explanation of the full range of data. You don't need to know the presupposition of the complement in order to know the complement. So the speaker's knowledge about what might be doesn't depend on her knowledge about what's epistemically correct.

One payoff of this paper is to start to show that epistemic expressions are no evidence for heterodox semantics, whether relativist or expressivist. The hy-

28. And it's very unclear if Silk's account can capture the full range of disagreement that my account captures. It would take a full paper to figure out if it can. By contrast, there isn't any question that my account does capture the full range of disagreement.

29. I introduced epistemic correctness as the mode of presentation to distinguish attitudes about epistemic modal bases from attitudes about circumstantial modal bases. But this way of motivating the mode of presentation does not lead directly to epistemic correctness. There are lots of different ways of thinking about the relevant mode of presentation. Eavesdroppers are relevant to this choice: it would be good to identify a mode of presentation that they could disagree about. That's why I've opted for epistemic correctness, but other choices might be possible.

brid theorist gives a fully traditional account that captures two central facts that the heterodox doubt can be explained in that way. Now there is a further range of complicated data about epistemics that I haven't had space to explore here. For example, there are complicated issues about when you can appropriately *retract*.³⁰ The hybrid theorist also has powerful new resources to bring to bear on that data. What you accept as epistemically correct over time can change. The fact that you accept one modal base as epistemically correct now doesn't also require you to do the same later. This fact opens up exciting new possibilities for theorizing about epistemics, especially because some presuppositional constructions exhibit some of the same puzzling behavior as epistemics do.

One illustration of those exciting possibilities concerns a puzzle from Seth Yalcin. He has noted that (16) and (17) both seem defective.

(16) # Suppose it is raining and it might not be raining.

(17) # If it is raining and it might not be raining, then... (Yalcin 2007: 985)

And he has argued that traditional contextualists can't explain this data. But my hybrid theory can. It associates the utterances with the following logical forms:

(16lf)

AT-ISSUE: suppose it is raining and m_1 is compatible with it not raining

NOT-AT-ISSUE: m_1 is epistemically correct

(17lf)

AT-ISSUE: if it is raining and m_1 is compatible with it not raining, ...

NOT-AT-ISSUE: m_1 is epistemically correct

The content supposed and the content in the antecedent are each inconsistent with the not-at-issue content. If it is raining and m_1 is compatible with it not raining, then there is another modal base that is more epistemically correct than m_1 : m_1 plus the proposition that it is raining. In general, this sort of inconsistency leads to significant infelicity.

(18a) # Suppose that it's not raining and John knows that it's raining.

AT-ISSUE: suppose that it's not raining and John knows that it's raining

NOT-AT-ISSUE: it's not raining

(18b) # If John knows that it's raining and it's not raining, ...

30. An important range of complicated cases are discussed in Andy Egan, John Hawthorne, and Brian Weatherson (2005), Andy Egan (2007), and Anthony Gillies and Kai von Fintel (2008).

AT-ISSUE: if it's not raining and John knows that it's raining, ...

NOT-AT-ISSUE: it's not raining

In both of these cases, the content supposed and the content in the antecedent are each inconsistent with the not-at-issue content. In fact, presuppositions in general behave like this example illustrates.³¹ For my purposes, it doesn't matter why they behave like this—only that they do.³² Since presuppositions in general behave like this, the hybrid theory explains Yalcin's data, without any further assumptions.³³

The hybrid theory also explains some further facts about Yalcin's data. Cian Dorr and John Hawthorne have noted that there are special cases where constructions like Yalcin's are assertable.

(19) If there is no spider in the closet and I am having a fatal heart attack because there might be, then I am about to die an absurd death. (Dorr & Hawthorne 2013: 877)

Similar constructions are possible with presupposition triggers, like the Krugmanesque (20).

(20) If Obama hasn't greatly increased government spending, and the budget is under discussion because everyone knows that Obama *has* greatly increased government spending, the discussion of the budget will be a mess.

31. Here are three more examples:

(18c) # Suppose it never has rained and it just stopped raining.

(18d) # Suppose that no one stole the tarts and it was John who stole the tarts.

(18e) # Suppose that no one stole the tarts and the tart thief was English.

32. You might think that this example has nothing at all to do with presupposition. The at-issue content is itself contradictory—so it's no surprise that we hear them as inconsistent! In making this suggestion, you are assuming that presuppositions are *entailments* of their triggers. That is, you're assuming that if an utterance of a sentence *S* presupposes *p*, then *p* is one of *S*'s entailments. Someone who thinks this needs to explain why attitudes towards presupposition triggers don't require the same attitude towards the presupposition. (You can hope that John knows that his mother died, without hoping that his mother died.) Once you explain this point, I will adopt your explanation to my hybrid theory. I will say that 'Might *p*' entails that the salient modal base is epistemically correct, and draw on your account of why attributions of reasonable belief don't require reasonable belief about that entailment.

33. Alex Silk (2017) has suggested a similar account of Yalcin's data, an account which also draws on a presupposition. But the distinctive contribution of this paper is its account of reasonable belief and disagreement, not its account of Yalcin's data. I'm mentioning Yalcin's data only to illustrate how views like mine can capture it.

It is possible (if difficult) for a presupposition to be somehow cancelled, in the way that (20) illustrates.³⁴ Given this point, the hybrid theory *predicts* that there would be examples like (19), where the presupposition about epistemic correctness is somehow cancelled. And if that presupposition is somehow cancelled, (19) is perfectly consistent. Moreover, it's plausible that the presupposition would be cancelled in this context. A use of (19) doesn't seem to presuppose that the salient modal base is epistemically correct! But the hybrid theory also predicts that this sort of example is possible only in special cases, because presuppositions can be canceled only in special contexts. The hybrid theory has exactly the right structure to capture Yalcin's phenomenon.³⁵

The general success of the hybrid theory suggests that it's well worth exploring. It is simple and highly principled, and powerfully explains a range of otherwise recalcitrant data.

Acknowledgements

For extremely helpful comments and conversations about this paper, I'm grateful to two referees at *Ergo*, and to Steve Finlay, John Hawthorne, Nathan Howard, Ben Lennertz, John MacFarlane, Barry Schein, Scott Soames, Ralph Wedgwood, and Jon Wright. I'm especially grateful to Mark Schroeder for invaluable discussion of the ideas in this paper and of the broader project.

References

- Abusch, Dorit (2010). Presupposition Triggering from Alternatives. *Journal of Semantics*, 27(1), 37–80.
- Adams, Ernest (1975). *The Logic of Conditionals*. D. Reidel.
- Anand Pranav and Valentine Hacquard (2013). Epistemics and Attitudes. *Semantics & Pragmatics*, 6(8), 1–59.
- Bach, Kent (2011). Perspectives on Possibilities: Contextualism, Relativism, or What? In Andy Egan and Brian Weatherson (Eds.), *Epistemic Modality* (19–59). Oxford University Press.

34. Robert Stalnaker (1974) influentially unified a range of otherwise difficult data by supposing that this is possible, and Scott Soames (2009) notes important further points about this phenomenon. Dorit Abusch (2010) has an especially helpful discussion of these points.

35. Silk (2017b: Section 3.4) discussed this point in much more detail. There are some structural similarities between this presupposition explanation and the one from Dorr and Hawthorne (2013). The explanations are still quite different, as Silk helpfully brings out. Their explanation appeals to a conversational implicature, not to a presupposition. And my explanation follows from a simple and highly conservative conjecture about the epistemic use of modals.

- Barwise, Jon and Robin Cooper (1981). Generalized Quantifiers and Natural Language. *Linguistics and Philosophy*, 4(2), 159–219.
- Bennett, Jonathan (2003). *A Philosophical Guide to Conditionals*. Oxford University Press.
- Björnsson, Gunnar and Stephen Finlay (2010). Metaethical Contextualism Defended. *Ethics*, 121(1), 7–36.
- Björnsson Gunnar and Alexander Almér (2010). The Pragmatics of Insensitive Assessments: Understanding the Relativity of Assessments of Judgments of Personal Taste, Epistemic Modals, and More. *The Baltic International Yearbook of Cognition, Logic, and Communication*, 6, 1–45.
- Braun, David (2012). An Invariantist Theory of ‘Might’ Might be Right. *Linguistics and Philosophy*, 35(6), 461–489.
- Cappelen, Herman and John Hawthorne (2009). *Relativism and Monadic Truth*. Oxford University Press.
- Cook, Ezra (2013). Epistemic Modals and Common Ground. *Inquiry*, 56(2–3), 179–209.
- Dorr, Cian and John Hawthorne (2013). Embedding Epistemic Modals. *Mind*, 122(488), 867–913.
- Dowell, Janice (2011). A Flexible Contextualist Account of Epistemic Modals. *Philosopher’s Imprint*, 11(14), 1–25.
- Dowell, Janice (2013). Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity. *Inquiry*, 56(2–3), 149–178.
- Edgington, Dorothy (1995). On Conditionals. *Mind*, 104(414), 235–329.
- Egan, Andy (2007). Epistemic Modals, Relativism, and Assertion. *Philosophical Studies*, 133(1), 1–22.
- Egan, Andy, John Hawthorne, and Brian Weatherson (2005). Epistemic Modals in Context. In Gerhard Preyer and Georg Peter (Eds.), *Contextualism in Philosophy* (131–168). Oxford University Press.
- Elbourne, Paul (2005). *Situations and Individuals*. The MIT Press.
- Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Finlay, Stephen (2016). Disagreement Lost and Found. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 12, 187–205). Oxford University Press.
- Gauker, Christopher (2008). Against accommodation: Heim, van der Sandt and the presupposition projection problem. *Philosophical Perspectives*, 22(1), 129–163.
- Geach, Peter (1965). Assertion. *Philosophical Review*, 74(4), 449–465.
- Geurts, Bart (1998). Presuppositions and Anaphors in Attitude Contexts. *Linguistics and Philosophy*, 21(6), 545. ISSN 0165-0157.
- Hacquard, Valentine (2006). *Aspects of Modality* (Doctoral dissertation). Massachusetts Institute of Technology.

- Hale, Bob (1993). Can There Be a Logic of Attitudes? In Crispin Wright and John Haldane (Eds.), *Reality, Representation and Projection* (337–363). Oxford University Press.
- Hawthorne, John (2007). Eavesdroppers and Epistemic Modals. *Philosophical Issues*, 17, 92–101.
- Hawthorne, John and David Manley (2012). *The Reference Book*. Oxford University Press.
- Heim, Irene (1982). *The Semantics of Definite and Indefinite Noun Phrases* (Doctoral dissertation). University of Massachusetts.
- Heim, Irene (1992). Presupposition Projection and the Semantics of Attitude Verbs. *Journal of Semantics*, 9(3), 183–221.
- Huvenes, Torfinn (2015). Epistemic Modals and Credal Disagreement. *Philosophical Studies*, 172(4), 987–1011.
- Kamp, Hans (1981). A Theory of Truth and Semantic Representation. In Groenendijk Janssen and Martin Stokhof (Eds.), *Formal Methods in the Study of Language* (277–322). Mathematical Centre.
- Karttunen, Lauri (1974). Presupposition and Linguistic Context. *Theoretical Linguistics*, 1(1–3), 181–194.
- Khoo, Justin (2015). Modal Disagreements. *Inquiry*, 58(5), 511–534.
- Knobe, Joshua and Seth Yalcin (2014). Epistemic Modals and Context: Experimental Data. *Semantics & Pragmatics*, 7, 1–21.
- Kratzer, Angelika (1977). What ‘Must’ and ‘Can’ Must and Can Mean. *Linguistics and Philosophy*, 1(3), 337–355.
- Kratzer, Angelika (1981). The Notional Category of Modality. In Hans J. Eikmeyer and Hannes Rieser (Eds.), *Words, Worlds, and Contexts* (38–74). De Gruyter.
- Kratzer, Angelika (2012). *Modals and Conditionals. New and Revised Perspectives*. Oxford University Press.
- Kripke, Saul (2009). Presupposition and Anaphora: Remarks on the Formulation of the Projection Problem. *Linguistic Inquiry*, 40(3), 367–386.
- Lennertz, Ben (2013). *Reasoning with Uncertainty and Epistemic Modals* (Doctoral dissertation). University of Southern California.
- Lennertz, Ben (2014). Simple Contextualism about Epistemic Modals is Incorrect. *Thought: A Journal of Philosophy*, 3(4), 252–262.
- López de Sa, Dan (2008). Presuppositions of Commonality: An Indexical Relativist Account of Disagreement. In García-Carpintero and Max Kolbel (Eds.), *Relative Truth* (297–310). Oxford University Press.
- MacFarlane, John (2007). Relativism and Disagreement. *Philosophical Studies*, 132(1), 17–31.
- MacFarlane, John (2011). Epistemic Modals are Assessment Sensitive. In Andy Egan and Brian Weatherson (Eds.), *Epistemic Modality* (144–178). Oxford Uni-

- versity Press.
- MacFarlane, John (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
- Montminy, Martin (2012). Epistemic Modals and Indirect Weak Suggestives. *Dialectica*, 66(4), 583–606.
- Moss, Sarah (2013). Epistemology Formalized. *Philosophical Review*, 122(1), 1–43.
- Moss, Sarah (2015). On the Semantics and Pragmatics of Epistemic Vocabulary. *Semantics and Pragmatics*, 8, 1–81.
- Perl, Caleb (2017a). *Positivist Realism* (Doctoral dissertation). University of Southern California.
- Perl, Caleb (2017b). Presuppositions, Attitudes, and Why They Matter. Manuscript in preparation.
- Schlenker, Philippe (2003). A Plea for Monsters. *Linguistics and Philosophy*, 26(1), 29–120.
- Schneider, Benjamin (2010). Expressivism Concerning Epistemic Modals. *Philosophical Quarterly*, 60(240), 601–615.
- Schroeder, Mark (2008). *Being For: Evaluating the Semantic Program of Expressivism*. Oxford University Press.
- Schroeder, Mark (2015a). Attitudes and Epistemics. In *Expressing Our Attitudes* (225–256). Oxford University Press.
- Schroeder, Mark (2015b). *Expressing Our Attitudes: Explanation and Expression in Ethics, Volume 2*. Oxford University Press.
- Schroeder, Mark (2015c). Is Semantics Formal? In *Expressing our Attitudes* (209–224). Oxford University Press.
- Searle, John (1962). Meaning and Speech Acts. *Philosophical Review*, 71(4), 423–432.
- Silk, Alex (2016). *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. Oxford University Press.
- Silk, Alex (2017a). How to Embed an Epistemic Modal: Attitude Problems and Other Defects of Character. *Philosophical Studies*, 174(7), 1773–1799.
- Silk, Alex (2017b). Normative Language in Context. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Volume 12*. Oxford University Press.
- Simons, Mandy, Tonhauser, Judith, Beaver, David, and Roberts, Craige (2011). What Projects and Why. *Proceedings of Semantics and Linguistic Theory (SALT)*, 22, 309–327.
- Soames, Scott (2009). Presupposition. In *Philosophical Essays, Volume 1: Natural Language: What It Means and How We Use It* (73–130). Princeton University Press.
- Soames, Scott (2011). True At. *Analysis*, 7(1), 124–133.
- Stalnaker, Robert (1974). Pragmatic Presuppositions. In Milton K. Munitz and Peter K. Unger (Eds.), *Semantics and Philosophy* (197–214). New York University

- Press.
- Stalnaker, Robert (2002). Common Ground. *Linguistics and Philosophy*, 25(5–6), 701–721.
- Stephenson, Tamina (2007). Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30(4), 487–525.
- Sudo, Yasutada (2012). *On the Semantics of Phi Features on Pronouns* (Doctoral dissertation). Massachusetts Institute of Technology.
- Swanson, Eric (2006). *Interactions with Context* (Doctoral dissertation). Massachusetts Institute of Technology.
- Tonhauser, Judith, David Beaver, Craig Roberts, and Mandy Simons (2013). Towards a Taxonomy of Projective Content. *Language*, 89(1), 66–109.
- Unwin, Nicholas (1999). Quasi-Realism, Negation and the Frege-Geach Problem. *The Philosophical Quarterly*, 49(196), 337–352.
- Unwin, Nicholas (2001). Norms and Negation: A Problem for Gibbard’s Logic. *The Philosophical Quarterly*, 51(202), 60–75.
- Veltman, Frank (1996). Defaults in Update Semantics. *Journal of Philosophical Logic*, 25(3), 221–261.
- von Fintel, Kai (2008). What Is Presupposition Accommodation Again? *Philosophical Perspectives*, 22(1), 137–170.
- von Fintel, Kai and Anthony Gillies (2008). CIA Leaks. *Philosophical Review*, 117(1), 77–98.
- von Fintel, Kai and Anthony Gillies (2011). Might Made Right. In Andy Egan and Brian Weatherson (Eds.), *Epistemic Modality* (108–130). Oxford University Press.
- Willer, Malte (2013). Dynamics of Epistemic Modality. *Philosophical Review*, 122(1), 45–92.
- Willer, Malte (2014). Dynamic Thoughts on Ifs and Oughts. *Philosophers’ Imprint*, 14(28), 1–30.
- Wright, Crispin (2007). New Age Relativism and Epistemic Possibility: The Question of Evidence. *Philosophical Issues*, 17, 262–283.
- Yalcin, Seth (2007). Epistemic Modals. *Mind*, 116(464), 983–1026.
- Yalcin, Seth (2011). Nonfactualism about Epistemic Modality. In Andy Egan and Brian Weatherson (Eds.), *Epistemic Modality* (295–332). Oxford University Press.