

Epistemic Modals and Epistemic Modality

Brian Weatherson Andy Egan

2011

Abstract: This chapter introduces the main themes of the volume, summarizes the chapters in it, and looks at the various arguments that have been raised for semantic relativism over the past decade. It concludes that two of these arguments seem to be resistant to the anti-relativist replies that have appeared in response to this work on relativism. One of these is an argument from agreement. It is argued that contextualist theories about various puzzling locutions have a hard time explaining why it is so easy for people who would happily utter the same words to describe themselves as agreeing, if those words were really context-sensitive. Another is an argument concerning attitude ascriptions. It seems there are quite different restrictions on what values the (allegedly) context-sensitive expressions can take inside and outside of attitude ascriptions. Since this isn't how context-sensitive terms usually behave, this phenomena tells against contextualism, and in favour of relativism.

1 Epistemic Possibility and Other Types of Possibility

There is a lot that we don't know. That means that there are a lot of possibilities that are, epistemically speaking, open. For instance, we don't know whether it rained in Seattle yesterday. So, for us at least, there is an epistemic possibility where it rained in Seattle yesterday, and one where it did not. It's tempting to give a very simple analysis of epistemic possibility:

- A possibility is an epistemic possibility if we do not know that it does not obtain.

But this is problematic for a few reasons. One issue, one that we'll come back to, concerns the first two words. The analysis appears to quantify over possibilities. But what are they? As we said, that will become a large issue pretty soon, so let's set it aside for now. A more immediate problem is that it isn't clear what it is to have de re attitudes towards possibilities, such that we know a particular possibility does or doesn't obtain. Let's try rephrasing our analysis so that it avoids this complication.

- A possibility is an epistemic possibility if for every p such that p is true in that possibility, we do not know that p is false.

If we identify *possibilities* with metaphysical possibilities, this seems to rule out too much. Let p be any contingent claim whose truth value we don't know. We do know, since it follows from the meaning of *actually*, that p iff actually p is true. But that biconditional isn't true in any world where p 's truth value differs from its actual truth value. So the only epistemic possibilities are ones where p 's truth value is the same as it actually is. But p was arbitrary in this argument, so the only epistemic possibilities are ones where *every* proposition has the same truth value as it actually does. This seems to leave us with too few epistemic possibilities!

A natural solution is to drop the equation of *possibilities* here with metaphysical possibilities. We've motivated this by using a proposition that is easy to know to be true, though it isn't true in many metaphysical possibilities. There are many problems from the other direction; that is, there are many cases where we want to say that there is a certain kind of epistemic possibility, even though there is no matching metaphysical possibility. We'll go through five such examples.

First, there are necessary *a posteriori* claims that arise from the nature of natural kinds. The standard example here is *Water is atomic*. That couldn't be true; necessarily, anything atomic is not water. But until relatively recently, it was an epistemic possibility.

Second, there are claims arising from true, and hence metaphysically necessary, identity and non-identity statements. A simple example here is *Hesperus is not Phosphorus*. This could not be true; by necessity, these celestial bodies are identical. But it was an epistemic possibility.

Third, there are claims about location. It isn't quite clear what proposition one expresses by saying *It's five o'clock*, but, plausibly, the speaker is saying of a particular time that that very time is five o'clock. It's plausible that if that's true, it's true as a matter of necessity. (Could this very time have occurred earlier or later? It doesn't seem like it could have.) So a false claim about what time it is will be necessarily false. But often there will be a lot of epistemic possibilities concerning what time it is.

Temporal location raises further matters beyond the necessary *a posteriori*. We want there to be epistemic possibilities in which it is four o'clock, five o'clock and so on. But it isn't altogether clear whether claims like that can be true in metaphysical possibilities. If we identify a metaphysical possibility with a possible world, then it isn't clear what would make it the case that it is four o'clock in a possible world. (What time is it in this possible world?) This might suggest there are different *kinds* of facts at a metaphysical possibility as at an epistemic possibility.

Fourth, there are issues about mathematics. Actually, there are two kinds of puzzle cases here. One concerns propositions that are logical consequences of our mathematical beliefs, but which we haven't figured out yet. Twenty years ago, it certainly seemed

to be an epistemic possibility that the equation $a^n + b^n = c^n$ had positive integer solutions with $n > 2$. Now we know that there are no such solutions. Moreover, if mathematics is necessarily true, then there isn't even a metaphysical possibility in which there are such solutions. So we shouldn't think that there was some metaphysical possibility that twenty years ago we hadn't ruled out. Rather, we were just unsure what metaphysical possibilities there are.

Finally, there are issues about logic. (Some views on the nature of logic and mathematics will deny that our fourth and fifth categories are different.) Getting the metaphysics of taste right is hard. One option that we think at least can't be ruled out is that intuitionist logic is the correct logic of taste talk. That is, when it comes to taste, we don't even know that it's true that everything either is or is not tasty. But that doesn't mean we're committed to the existence of a possibility where it isn't true that everything is tasty or not tasty; if such a state isn't actual, it probably isn't possible. The liar paradox is even harder than the metaphysics of taste. Anything should be on the table, even the dialethist view that the Liar is both true and false. That is, the Liar might be true and false. In saying that, we certainly don't mean to commit to the existence of some possibility where the Liar is true and false. We're pretty sure (but not quite certain!) that no such possibility exists.

The last two cases might be dealt with by being more careful about what an epistemic possibility is. There are quite simple cases in which we want to resist the identification of epistemic possibilities with what we don't *know* to be the case. For discussion of several such cases, see Hacking (1967), Teller (1972) and DeRose (1991). If we could very easily come to know that p does not obtain, perhaps because p is clearly ruled out by things we do know, then intuitively it isn't the case that p is epistemically possible. If we know that if q then not p , and we know q , then p is not possible, even if we haven't put conditional and antecedent together to conclude that p is false. So we need to put some constraints on the epistemically possible beyond what we know to be false. Perhaps those constraints will go so far as to rule out anything inconsistent with what we know. In that case, it wasn't possible all along that Fermat's Last Theorem was false. And, assuming the non-classical approaches to taste and alethic paradoxes are incorrect, those approaches aren't even possibly correct. We're not endorsing this position, just noting that it is a way to rescue the idea that all epistemic possibilities are metaphysical possibilities.

The papers in this volume that most directly address these issues are by Frank Jackson, David Chalmers and Robert Stalnaker. Jackson argues against the view that accounting for epistemic possibilities requires us to think that there is a kind of possibility, conceptual possibility, that is broader than metaphysical possibility. He briefly reviews the reasons some people have had for taking this position, including those we've

just discussed, and some of the reasons he rejected it in “From Metaphysics to Ethics”. But he adds some new arguments as well against this position, what he calls the ‘two space’ view of possibility. One argument says that if there is a possibility of any kind where water is not H_2O , then being water and being H_2O must be different properties by Leibniz’s Law. But then we have an implausible necessary connection between distinct properties. Another argument turns on the difficulty of identifying the water in these supposed conceptual possibilities that are not metaphysically possible.

David Chalmers discusses what kind of thing epistemic possibilities, or as he calls them, ‘scenarios’, might be. He discusses the strengths, weaknesses and intricacies of two proposals: what he calls the ‘metaphysical’ and ‘epistemic’ constructions. The metaphysical construction is fairly familiar: it takes epistemic possibilities to be centered possible worlds. The epistemic construction takes epistemic possibilities to be maximally possible sentences of a specially constructed language. The metaphysical construction requires several assumptions before it matches up with the intuitive notion of epistemic possibility, while the epistemic construction requires a primitive notion of epistemic possibility. But both constructions seem to illuminate the elusive notion of an epistemic possibility. Chalmers ends with a discussion of several applications of his constructions in semantics, in formal epistemology and in moral psychology.

Another place where one finds an important role for a distinctively epistemic (or at least doxastic) sort of possibility is in theorizing about indicative conditionals. In Robert Stalnaker’s contribution, he examines two types of accounts of indicative conditionals, which differ in where they locate the conditionality. One view analyzes assertions of indicative conditionals as a special sort of *conditional assertion*, and another analyzes them as an ordinary assertion of a special sort of *conditional proposition*. Stalnaker argues that the two views are not so different as we might initially have thought.

2 Three Approaches to Epistemic Modals

Even when we settle the issue of what epistemic possibilities are, we are left with many issues about how to talk about them. Speakers will often say that something is (epistemically) possible, or that it might be true. (It’s plausible that claims that p must be true, or that p is probable, are closely related to these, but we’ll stick to claims about (epistemic) possibility at least for this introduction.) It’s plausible to think that a proposition isn’t possible or impossible simpliciter, it’s rather that it is possible or impossible relative to some person, some group, some evidence or some information. Yet statements of epistemic possibility in plain English do not make any explicit reference to such a person, group, evidence set or information state. One of the key issues confronting a semanticist attempting to theorise about epistemic modals is what to do about this lack of

a reference. We'll look at three quite different approaches for dealing with this lack: contextualist, relativist and expressivist.

2.1 Contextualism

Consider a particular utterance, call it u , made by speaker S , of the form *a might be F*, where the might here is intuitively understood as being epistemic in character. To a first approximation, the sentence is saying a 's being F is consistent with, or not ruled out by, a certain body of knowledge. But whose body of knowledge? Not God's, presumably, for then *a might be F* would be true iff *a is F* is true, and that's implausible. The contextualist answer is that the relevant body of knowledge is supplied by context.

When discussing the ways in which context fills in content, some writers will start with the pronoun I as an example. And to some extent it's a useful example. The sentence *I am a fool* doesn't have truth-conditional content outside of a context of utterance. But any utterance of that sentence does express something truth conditional. Which truth conditional sentence it expresses is dependent on facts about the context of its use. In fact, it is dependent on just one fact, namely who utters it. So when Andy utters *I am a fool* he expresses the proposition that Andy is a fool. And when Brian utters *I am a fool* he expresses the proposition that Brian is a fool.

So far I is a useful example of a context-sensitive expression. But in many ways it is an unusual example of context-sensitivity, and focussing too much on it can lead to an overly simplistic view of how context-sensitive terms work. In particular, I has three properties that are unusual for a context-sensitive expression.

- Its content in a context is computable from the context by a simple algorithm - namely the content is the speaker.
- Its content does not depend on any properties of the intended audience of the utterance.
- It behaves exactly the same way in embedded and unembedded contexts.

Some terms have none of these properties. Consider, for example, *we*.

There isn't any obvious algorithm for computing the content of a particular use of *we*. The content may depend on the intentions of the speaker. It may depend on which people have been talked about. In sentences of the form *We are F*, different values of F might constrain what values can be rationally assigned to *we*. And when that is so, the interpretation of *we* will (usually) be constrained to those groups.

Perhaps most notably, it depends a lot on the audience. If S is talking to H , and says *We should grab some lunch*, the content is that S and H should grab some lunch. And that's the content because H is the intended audience of the utterance. Intended audiences can change quickly. If Andy says *We will finish the paper this afternoon, then*

we will go for a walk, talking to Brian when he utters the first conjunct, and Fido when he utters the second, the content is that Andy and Brian will finish the paper this afternoon, then Andy and Fido will go for a walk.

That *we* has neither of the first two properties is uncontroversial. What is perhaps a little more controversial is that it does not have the third either. When *we* is in an unembedded context it (usually) behaves like a free (plural) variable. Under certain embeddings, it can behave like a bound variable. Barbara Partee and Phillippe Schlenker offer the following examples.

(5.9) John often comes over for Sunday brunch. Whenever someone else comes over too, *we* (all) end up playing trios. (Partee 1989)

(5.10) Each of my colleagues is so difficult that at some point or other *we've* had an argument. (Schlenker 2003)

In neither case does *we* contribute a group consisting of the speaker plus some salient individuals. Indeed, in neither case does it so much as contribute a group, since it is (or at least behaves like) a bound variable. There's nothing in the contextualist story about *we* that prevents this.

It's worthwhile reviewing these facts about *we*, because on the most plausible contextualist stories about *might*, it too has these three properties. The contextualist theory we have in mind says that the content of *u* is *For all that group X could know using methods M, a is F*. The group *X* will usually consist of the speaker and some salient others, perhaps including the intended audience. The salient methods might include little more than easy deduction from what is currently known, or may include some wider kinds of investigation. (See DeRose (1991) for arguments that the relevant methods include more than deduction, and that they are contextually variable.)

Now it isn't part of the contextualist theory that there is an easy method for determining who is in *X*, or what methods are in *M*. So in that respect *might* is like *we*. But, just as the group denoted by *we* typically includes the intended audience of the utterance, the group *X* will typically include the intended audience of *u*. And the methods *M* will typically include any method that can be easily carried out. This can be used to explain some phenomena about disagreement. So if Andy says, to Brian, *a might be F*, and Brian knows that *a* is not *F* (or can easily deduce this from what he knows), Brian can disagree with what Andy says. That is, he can disagree with the proposition that it is consistent with what members of the conversation know that *a* is *F*. And, the contextualist says, that's just what Andy did say. If Brian presents Andy with his grounds for disagreement, Andy might well retract what he said. Since arguments about disagreeing with utterances like *u* have been prominent in the literature, it is worth noting that the contextualist theory can explain at least some facts about disagreement.

Nor is it part of the contextualist theory that *might* behaves exactly the same way in embedded and unembedded contexts. Indeed, like *we*, *might* can behave like a bound variable. On the most natural reading of *Every pedestrian fears that they might be being watched*, there is no single group X such that every pedestrian fears that for all X (could easily) know, that pedestrian is being watched. Rather, every pedestrian fears that for all they themselves know, they are being watched. The naturalness of this reading is no embarrassment to the contextualist theory, since it is a common place that terms that usually get their values from context can also, in the right setting, behave like bound variables.

Indeed, thinking about these parallels between context-sensitive expressions and epistemic modals seems to provide some support for contextualism. In his contribution to the volume, Jonathan Schaffer argues that various features of the way epistemic modals behave in larger sentences support the idea that an evaluator place must be realised in the syntax. For instance, consider the natural interpretation of “Anytime you are going for a walk, if it might rain, you should bring an umbrella.” We interpret that as saying that whenever you go for a walk, you should bring an umbrella if your evidence at that time is consistent with rain. Schaffer interprets that as evidence that there is hidden syntactic structure in epistemic modals, and argues that the contextualist offers the best account of how the hidden structure gets its semantic values.

So the contextualist has a lot of explanatory resources, and a lot of flexibility in their theory, which are both clear virtues. But there are some limits to the flexibility. There are some things that the contextualist, at least as we’re using the term ‘contextualist’ is committed to. In particular, the contextualist is committed to the content of a particular speech act (or at least of a particular assertion) is absolute, not assessor-relative. And they’re committed to the truth value of those contents being the same relative to any assessor. Let’s give those two commitments names.

- (C) The semantic content of an assertion is the same relative to any assessors.
- (T) The truth value of the semantic content of an assertion is the same relative to any assessors.

The first of these rules out the possibility that the semantic content of an assertion differs with respect to different groups. The second rules out the possibility that semantic contents have assessor relative truth values. Modern relativists have proposed theories that dispense with these dogmas, and we’ll investigate those in the next section, after going over some of the motivations for relativism.

2.2 Relativism

In many fields, relativism is motivated by instances of “faultless disagreement”, and epistemic modals are not left out of this trend. Here is the kind of case that we used in Egan, Hawthorne, and Weatherson (2005) to motivate relativism.

Consider the following kind of case. Holmes and Watson are using a primitive bug to listen in on Moriarty’s discussions with his underlings as he struggles to avoid Holmes’s plan to trap him. Moriarty says to his assistant

(24) Holmes might have gone to Paris to search for me.

Holmes and Watson are sitting in Baker Street listening to this. Watson, rather inexplicably, says “That’s right” on hearing Moriarty uttering (24). Holmes is quite perplexed. Surely Watson knows that he is sitting right here, in Baker Street, which is definitely not in Paris.

Here we have Watson somewhat surprisingly agreeing with Moriarty. In some sense, it seems wrong for him to have done so. He should have disagreed. Well, imagine that he did, by saying “That’s not right”. The quick argument for relativism is that the contextualist cannot make sense of this. Whatever group’s knowledge Moriarty intended to be talking about when he spoke, it presumably didn’t include Holmes and Watson; it just included him and his intended audience, i.e. the underlings. And it’s true that for all they know, Holmes is in Paris. So the content of Moriarty’s utterance is true. But it seems that Watson can properly disagree with it (and can’t properly agree with it). That, we thought, was a problem.

There are three kinds of response to this argument on behalf of the contextualist that we think look promising. All of these responses are discussed in Fintel and Gillies (2008). We might look harder at the denotation of the ‘that’ in Watson’s reply, we might think again about what the relevant group is, and we might look at other cases where the contextualist story is more promising, as a way of motivating the first two responses. Let’s look at these in turn.

Above we said that Watson disagreed with Moriarty by saying “That’s not right”. But that’s potentially reading too much into the data. What seems correct is that Watson can say “That’s not right”. But that’s only to disagree with Moriarty if the ‘that’ denotes what Moriarty said. And that might not be true. It’s possible that it picks out, say, the embedded proposition that Holmes has gone to Paris. And it’s fine for Watson to disagree with that.

Even if Watson is disagreeing with the semantic content of Moriarty’s utterance, it might be that he’s doing so properly, because what Moriarty said is false. That might

be the case because it might be that, in virtue of hearing the utterance, Watson became part of the relevant group X . Typically speaker intentions, particularly singular speaker intentions, are not the final word in determining the content of a context-sensitive expression. If Brian points over his shoulder, thinking a nice glass of shiraz is behind him, and says *That is tasty*, while in fact what he is pointing at is a vile confection of Vegemite infused chardonnay, he's said something false. The simplest thing to say about a case like this is that Brian intended the denotation of 'That' to be the thing he was pointing at, whatever it is. Similarly, Moriarty might have intended the relevant group X to be whoever heard the utterance at that time, even if he didn't know Watson was in that group. (Or it might be that, whatever Moriarty's intentions, the semantic rules and conventions for 'might' in English determine that the relevant group includes everybody who heard the utterance at the time.)

This second response would seem somewhat ad hoc were it not for a class of examples von Fintel and Gillies describe concerning assessors from radically different contexts. Typically the anti-contextualist commentary on cases like these suggest that any hearer who knows that a is not F can disagree with u . But that doesn't seem to be in general true.

Or consider the case of Detective Parker. He has been going over some old transcripts from Al Capone's court case in the 1920s—Capone is being asked about where some money is in relation to a particular safe:

(20)

1. Capone: The loot might be in the safe.
2. Parker: ??Al was wrong/What Al said is false. The safe was cracked by Geraldo in the 80s and there was nothing inside.
(2008, 86)

The knowledge of at least some hearers, such as Detective Parker, does not seem to count for assessing the correctness of Capone's speech. A contextualist might suggest that's because contemporaneous hearers are in the relevant group, and later reviewers are not.

So there are definitely some contextualism-friendly lines of response available to the argument for relativism from disagreement. But interestingly, some of these contextualist responses do not work as well as a response to a similar argument from agreement. Imagine that Andy, after doing some reading on the publicly available evidence, correctly concludes that it doesn't rule out Prince Albert Victor. He doesn't think this is very likely, but thinks it is possible. Andy hears someone on TV talking about the Ripper who says "Prince Albert Victor might have been Jack the Ripper", and Andy

says “That’s true”. Intuitively Andy is right to agree with the TV presenter, but this is a little hard to explain on the contextualist theory.

Note that here we can’t say that Andy is agreeing because he is agreeing with the embedded proposition, namely that Prince Albert Victor was the Ripper. That’s because he doesn’t agree with that; he thinks it is an open but unlikely possibility.

Nor does it particularly matter that Andy, as one of the people watching the TV show, is part of the relevant group X . All that would show is that if Andy knew Prince Albert Victor wasn’t the Ripper, the presenter’s assertion is false. But unless Andy is the group X , the fact that Andy’s knowledge, or even what is available to Andy, does not rule out the Prince does not mean Andy should agree with the statement. For all Andy knows, someone else watching, perhaps even someone else the presenter intends to include in her audience, has evidence exculpating the Prince. If that’s right, then he does not know that the proposition the contextualist says the speaker asserted is true. But yet he seems justified in agreeing with the presenter. This seems like a real problem for contextualism.

A quite different objection to contextualism comes from metasemantic considerations. The most casual reflection on the intuitive content of utterances like μ suggests there is staggeringly little rhyme or reason to which group X or method M might be relevant. The argument here isn’t that the contextualist’s semantic proposal is mistaken in some way. Rather, the argument is that the accompanying metasemantic theory, i.e. the theory of how semantic values get fixed, is intolerably complicated. Slightly more formally, we can argue as follows.

1. If contextualism is true, the metasemantic theory of how a particular use of ‘might’ gets its semantic value is hideously complicated.
2. Metasemantic theories about how context-sensitive terms get their values on particular occasions are never hideously complicated.
3. So, contextualism is false.

The problem with this argument, as Michael Glanzberg (2007) has argued, is that premise 2 seems to be false. There are examples of uncontroversially context-sensitive terms, like ‘that’, for which the accompanying metasemantic theory is, by any standard, hideously complicated. So the prospects of getting to relativism from metasemantic complexity are not, we think, promising.

But there is a different metasemantic motivation for relativism that we think is a little more promising. Compare the difference between (1) and (2).

- (1) Those guys are in trouble, but they don’t know that they are. (2)
 ??Those guys are in trouble, but they might not be.

Something has gone wrong in (2). This suggests that (2) can't be used to express (1). That is, there's no good interpretation of (2) where those guys are the group X . This is a little surprising, since we've made the guys pretty salient. Cases like this have motivated what we called the Speaker Inclusion Constraint (hereafter SIC) in "Epistemic Modals in Context". That is, in unembedded uses of 'might' the group X always includes the speaker. Now the explanation of the problem with (2) is that for the speaker to assert the first clause, she must know that the guys are in trouble, but if that's the case, and she's in group X , then the second clause is false.

Now a generalisation like this doesn't look like it should be grounded in the meaning (in some sense of 'meaning') of 'might'. For comparison, it seems to be part of the meaning of 'we' that it is a first-person plural pronoun. It isn't just a metasemantic generalisation that the speaker is always one of the group denoted by 'we'. By analogy, it is part of the meaning of 'might' that the speaker is always part of the group X .

Further, when the meaning of a context-sensitive expression constrains its value, those constraints still hold when the term is used as a bound variable. For instance, it is part of the meaning of 'she' that it denotes a female individual. If Smith is male, then the semantic content of *She is happy* can't be that Smith is happy. Similarly, when 'she' is behaving like a bound variable, the only values it can take are female individuals. So we can't use *Every student fears she will fail the test* to quantify over some students some of whom are male. And there's no interpretation of *Every class hopes we will win* where it means that every class hopes that that class will win. Even when under a quantifier and an attitude ascribing verb, 'we' must still pick out a group that includes the speaker. The natural generalisation is that constraints on context supplied by meaning do not get overridden by other parts of the sentence.

The problem for contextualists about 'might' is that it doesn't behave as you'd expect given these generalisations. In particular, the SIC doesn't hold when 'might' is in certain embeddings. So there is a reading of *Every student fears they might have failed* where it means that every student fears that, for all they know, they failed. The knowledge of the speaker isn't relevant here. Indeed, even if the speaker knows that many students did not fail, this sentence can be properly uttered. This suggests the following argument.

1. If contextualism is true, then the explanation of the SIC is that it is part of the meaning of 'might' that the relevant group X includes the speaker.
2. If it is part of the meaning of 'might' that the relevant group X includes the speaker, then this must be true for all uses of 'might', included embedded uses.
3. When 'might' is used inside the scope of an attitude ascription, the relevant group need not include the speaker.
4. So, contextualism is not true.

Premise 1 would be false if the metasemantics was allowed to be systematic enough to explain why the SIC holds even though it is not part of the meaning. Premise 2 would be false if we allowed ‘might’ to have a systematically different meaning inside and outside the scope of attitude ascriptions. And premise 3 would be false if any attitude ascriptions that are made are, contrary to intuition, tacitly about the speaker’s knowledge. Since none of these seems particularly plausible, there does seem to be a problem for contextualism here.

In their contribution to this volume, Kai von Fintel and Thony Gillies reject one of the presuppositions of the argument we’ve just presented. Classical contextualism, what they call ‘the canon’, says that context picks out a particular group, and an utterance of ‘It might be that p ’ is true iff that group’s information is consistent with p . That’s what we’ve taken as the stalking horse in this section, and von Fintel and Gillies are certainly right that it is the canonical version of contextualism. Von Fintel and Gillies agree that the broad outline of this contextualist story is correct. But they deny that context picks out a determinate group, or a determinate body of information. Rather, uttering an epistemic modal will ‘put into play’ a number of propositions of the form ‘For all group G knows, p ’. This ambiguity, or perhaps better indeterminacy, is crucial they argue to the pragmatic role that epistemic modals play. And once we are sensitive to it, they claim, we see that contextualism has more explanatory resources than we’d previously assumed, and so the motivation for relativism fades away.

In summary, there are four motivations for relativism that have been floated in the literature. These are:

- Intuitions about disagreement;
- Intuitions about agreement;
- Arguments from metasemantic complexity; and
- Arguments from semantic change in attitude ascriptions,

As noted, the third argument doesn’t seem very compelling, and it is a fairly open question whether the first works. But the second and fourth do look like good enough arguments to motivate alternatives.

2.3 Two Kinds of Relativism

We said above that contextualism is characterised by two theses, repeated here for convenience.

- (C) The semantic content of an assertion is the same relative to any assessors.
- (T) The truth value of the semantic content of an assertion is the same relative to any assessors.

So there are two ways to not be a relativist, deny (C) and deny (T). One might deny both, but we'll leave that option out of our survey.

What we call *content relativism* denies (C). The picture is that contextualists were right to posit a variable X in the structure of an epistemic modal claim. But the contextualists were wrong to think that X gets its value from the context of utterance. Rather, the value of X is fixed in part by the context of assessment. In the simplest (plausible) theory, X is the speaker and the assessor. So if Smith asserts that Jones might be happy, the content of that assertion, relative to Andy, is that for all Smith and Andy know, Jones is happy, while relative to Brian its content is that for all Smith and Brian know, Jones is happy.

The primary motivation for content relativism is that it keeps quite a bit of the contextualist picture, while allowing enough flexibility to explain the phenomena that troubled contextualism. So for the content relativist, contents are exactly the same kinds of propositions as the contextualist thinks they are. So we don't need to tell a new kind of story about what it is for a content to be true, to be accepted, etc. Further, because we keep the variable X , we can explain the 'bound variable' readings of epistemic modals discussed in the first section.

A worry about content relativism is that the 'metasemantic' argument against contextualism might equally well tell against it. The worry there was that the constraints on X seemed to depend, in an unhappy way, on where in the sentence it appeared. The content relativist has a move available here. She can say that as a rule, whenever there's a variable like X attached to a term, and that term is in an attitude ascription, then the variable is bound to the subject of the ascription. This might be an interesting generalisation. For instance, if she is a content relativist about both epistemic modals and predicates of personal taste, she has a single explanation for why both types of terms behave differently inside and outside attitude ascriptions.

There are two interesting 'near cousins' of content relativism. One is a kind of *content pluralism*. We might hold (a) that an assertion's content is not relative to an assessor, but (b) some assertions have many contents. So if S says *a might be F*, and this is assessed by many hearers, S asserts *For all s and b know, a is F*, for each H who hears and assesses the speech. Now when a hearer b_1 does this, she'll probably focus on one particular content of S 's assertion, namely that *For all s and b_1 know, a is F*. But the content pluralist accepts (while the content relativist denies) that even relative to b_1 , S 's assertion also had the content *For all s and b_2 know, a is F*, where b_2 is a distinct assessor.

Another near cousin is the view, defended in this volume by Kent Bach, that the semantic content of an epistemic modal is typically not a complete proposition. In the case just described, it might be that the semantic content of what S says is *For all _____*

knows, *a* is *F*, and that's not a proposition. Now a given hearer, *H*, might take *S* to have communicated to them that *For all s and b know, a is F*, but that's not because that's the semantic content of what *S* says. It's not the absolute content (a la contextualism), the content relative to *H* (a la content relativism) or one of the contents (a la content pluralism).

It's a very big question how we should discriminate between these theories. Some readers may even worry that there is no substantive differences between the theories, they are in some sense saying the same thing in different words. One big task for future research is to clearly state the competing theories in the vicinity of here, and find arguments that discriminate between them.

A quite different kind of relativism denies (T). This view says that the content itself of an assertion can be true for some assessors, and false for others. Such a view is not unknown in recent philosophy. In the 1970s and 1980s (and to a lesser extent in subsequent years) there was a debate between temporalists and eternalists about propositions. The temporalists thought that a tensed proposition, i.e. the content of a tensed assertion, could be true at one time and false at another. The eternalists denied this, either taking truth to be invariant across times, or in some cases denying that it even made sense to talk about truth being relative to something, e.g. a time.

Contemporary forms of truth relativism generalise the temporalist picture. The temporalists thought that propositions are true or false relative to a world-time pair. Modern relativists think that propositions are true or false relative to a world-assessor pair, or what loosely following Quine (1969) we might call a centered world. (Quine used this to pick out any world-place-time triple, but since most times and places don't have assessors at them, world-assessor pairs, or even world-assessor-time triples, are more restricted.) For example, as a first pass at a truth-relativism about predicates of personal taste, one might propose that the proposition expressed by a typical utterance of 'beer is tasty' will be true at any centered world where the person at the center of the world likes the taste of beer.

The truth relativist has an easy explanation of the data that motivated the rejection of contextualism. Recall two puzzles for the contextualist about terms like 'tasty': that it is so easy to agree with claims about what's tasty, and that reports of the form *X* thinks that beer is tasty are always about *X*'s attitude towards beer, not about *X*'s beliefs about how the speaker finds beer.

On the first puzzle, note that if to agree with an assertion is to agree with its propositional content, and that content is true at the center of your world iff you find beer tasty, then to agree with an assertion that beer is tasty, you don't have to launch an inquiry into the sincerity of the speaker, you just have to check whether you like beer. If you're in a world full in insincere speakers, and abundant beer, that's relatively easy.

On the second puzzle, if propositional attitude ascriptions report the subject's attitude towards a proposition, and if a proposition is a set of centered worlds, then the subject's attitude towards 'Beer is tasty' should be given by their attitude towards whether that proposition is true in their centered world. That is, it should be given by their attitude towards beer. And that's just what we find.

The extension of all this to epistemic modals is more or less straightforward. The simplest truth relativist theory says that an utterance of the form *a might be F* is true iff, for all the assessor at the center of the world knows, *a* is *F*. As Richard Dietz (2008) has pointed out, this won't do as it stands. If the speaker knows *a* is not *F*, then their utterance seems like it should be false relative to everyone. (Conversely, a speaker who knows *a* is *F* speaks truly, relative to any assessor, when they say *a must be F*.) If we're convinced of this, the solution is a mild complication of the theory. The utterance is both somewhat context-sensitive, and somewhat relative. So *S*'s utterance of *a might be F* is true at a centered world iff for all *S* plus the person at the center of the world know, *a* is *F*. We might want to add more complications (is it knowledge that matters or available information, for example?) but that's one candidate truth relativist theory.

There are three worries we might have about truth relativism. One is a very big picture worry that the very notion of truth being relative is misguided. This is a theme of Herman Cappelen and John Hawthorne's *Relativism and Monadic Truth*. Another is that it overgenerates 'explanations'. We can't explain cases like the Capone/Parker example. And a third is that, by making propositions so different from what we thought they were, we'll have to redo a lot of philosophy of language that presupposed propositions have the same truth value for everyone. In particular, we'll have to rethink what an assertion is. (That challenge is addressed – in different ways – in recent work by John MacFarlane and by Andy Egan.)

The strongest defence of relativism in this volume comes from John MacFarlane. His work on tense (MacFarlane 2003), and on knowledge attributions (MacFarlane 2005b), and on the broader philosophical status of relativism and other rivals to classical contextualism (MacFarlane 2005a, 2009), have been immensely influential in the contemporary debates. Here he develops a relativistic semantics for epistemic modals, along the lines of the proposals he has offered elsewhere for tense and knowledge attributions. He argues that many phenomena, several of which we've discussed in this introduction, raise trouble for contextualism and promote relativism. These phenomena include third-party assessments, retraction and disagreement. He argues that only the relativist can explain the troublemaking phenomena.

2.4 Expressivism

So far we've looked at two of our three major approaches to epistemic modals. The contextualist says that which proposition is asserted by an epistemic modal depends crucially on the context of utterance. The relativist says that the contextualist is ignoring the importance of the context of assessment. The content relativist says that they are ignoring the way in which the context of assessment partially determines what is said. The truth relativist says that they are ignoring the way in which propositions uttered have different truth values at different contexts of assessment.

The expressivist thinks that there is a common assumption behind all of these theories, and it is a mistaken assumption. The assumption is that when we're in the business of putting forward epistemic modals, we're in the business of asserting things that have truth values. The expressivist rejects that assumption. They say that when we say *a might be F*, we're not asserting that we are uncertain about whether *a* is *F*, we're expressing that uncertainty directly. The contextualists and relativists think that in making these utterances, we're expressing a second-order belief, i.e. a belief about our own knowledge, or lack thereof. The expressivists think we're expressing a much simpler mental state: uncertainty.

One way to motivate expressivism is to start with the anti-contextualist arguments, and then argue that relativism is not an acceptable way out. So we might, for instance, start with the argument from agreement. The expressivist notes that there are many ways to agree with a statement. If Smith says 'Let's have Chinese for dinner', and Jones agrees, there need not be any proposition that Smith asserted that Jones is agreeing to. We're happy to call all sorts of meetings of minds agreements. So the agreement phenomena that the contextualist can't explain, the expressivist can explain. When Smith says 'Brown might be a spy', and Jones agrees, there isn't necessarily any proposition they both accept. Rather, their agreement consists in having a common mental state, namely uncertainty about whether Brown is a spy.

The expressivist may then run out any number of arguments against relativism. For instance, they might argue (against content relativism) that it is a requirement of a speech act being an assertion that it have a determinate content. And they might argue, perhaps motivated by theoretical considerations about the role of assertions in conversation, that contents which vary in truth value among hearers couldn't be contents of assertions. If true, that would rule out truth relativism. We're moved, perhaps by elimination as much as anything, to expressivism.

There are more direct arguments for expressivism as well. Isaac Levi (1996, 55) motivated a view on which epistemic modals don't have truth values by thinking about learning. Imagine someone previously thought that Brown might be a spy, perhaps on

quite good grounds, then they learn that he is not a spy. If that's all they learned, then it seems odd to say that there's something that they previously knew, that now they don't know. It seems learning shouldn't destroy knowledge. That's what happens in standard models for belief revision (which were one of Levi's primary concerns) and it is independently plausible. But if epistemic modals express propositions, and those are true or false, then there is a proposition that the person did know and now don't know, namely that Brown might be a spy.

There are clearly a few possible responses to this argument. For one thing, we could make the epistemic modal claims explicitly tensed. Both before and after the learning experience, the subject knew that Brown might, at t_1 , have been a spy, but didn't know that Brown might, at t_2 , have been a spy. (Indeed, they learned that that was false.) Or, and this is more in keeping with the spirit of this introduction, we might spell out the epistemic modal claim. Before and after the learning experience, the subject knew that it was consistent with everything the subject knew prior to the learning experience that Brown was a spy. So there's no information lost.

The problem with this move is that it seems to make epistemic modals overly complex. Intuitively, it is possible for a child to grasp a modal, and for the most natural interpretation of that modal to be epistemic, without the child having the capacity to form second order thoughts. (This point is one that Seth Yalcin uses in his argument for a kind of expressivism in this volume.) This question seems like it would be good to test empirically, though we don't know of any existing evidence that settles the question. Introspectively, it does seem that one can think that the cat might be in the garden without thinking about one's own epistemic or doxastic states as such. Those kinds of introspections might tell in favour of an approach which identifies epistemic modality with a distinct kind of relation to content, rather than a distinct kind of content.

Following important work by Allan Gibbard (1990), there is a natural way to formalise an expressivist theory of epistemic modality. Identify a 'context' with a set of propositions. Sentences, whether epistemic modals or simple sentences, are satisfied or unsatisfied relative to world-context pairs, where a world and a context make a pair iff every proposition in the context is true at that world. Then an epistemic modal, say *Brown might be a spy*, is satisfied by such a pair iff Brown is a spy is consistent with everything in the context. A simple sentence, like *White is a spy* is satisfied by such a pair iff White is a spy is true at the world. The pairing becomes useful when considering, say, conjunctions. A conjunction is satisfied iff both conjuncts are satisfied. So *White is a spy and Brown might be* is satisfied by a world-context pair iff White is a spy at the world, and Brown's being a spy is consistent with the context.

So far this looks a lot like relativism. A world-context pair is just like a centered world, with the context being what's known by the person at the center of the world. If we

apply the formalism to real-life cases, perhaps taking the contexts to be genuine contexts in the sense of Stalnaker (1978), the two formalisms might look very close indeed.

But there is, or at least we hope there is, a substantive philosophical difference between them. The expressivist has a restricted sense of what it is to make an assertion, and of what it is for an expression to be an expression of a truth. The expressivist most insistently does not identify satisfaction with truth. The only sentences that are true or false are sentences that are satisfied by a world context pair $\langle w, c_1 \rangle$ iff they are satisfied by every other pair starting with the same world. The expression of such a sentence, and perhaps only of such a sentence, constitutes an assertion. Otherwise it constitutes some other speech act.

And this is no mere difference in how to use the words ‘truth’, ‘assertion’ and so on. Nor is it even just a difference about truth and assertion and so on. It hopefully makes a difference to what predictions we make about the way epistemic modals embed, especially how they embed in propositional attitude ascriptions. We used that fact to argue against expressivism in “Epistemic Modals in Context”, since we thought there were in some cases more examples of successful embedding of epistemic modals, especially in conditionals, than the expressivist would predict. On the other hand Seth Yalcin uses facts about embedding to argue, in his paper in this volume, in favour of expressivism. He argues that on a non-expressivist view, we should be able to suppose that p is true but might not be true, and that can’t be supposed.

This argument is part of the argument by elimination that Yalcin against what he calls ‘descriptivism’ about epistemic modals in this contribution to the volume. He uses ‘descriptivism’ to pick out a broad category of theories about epistemic modals that includes both contextualism and relativism. He argues against all descriptivist views, and in favour of what he calls ‘expressivism’. He says that when someone utters an epistemic modal, they do not describe their own knowledge (or the knowledge of someone else), rather they express their own mental state. Some of Yalcin’s arguments for expressivism are related to arguments against contextualism; in particular he thinks like we do that there isn’t a viable form of contextualism. But he also thinks that there are problems for relativism, such as the difficulty in supposing Moore paradoxical propositions. He also notes that it is a puzzle for descriptivists to make sense of belief ascriptions involving epistemic modals. On a descriptivist model, a sentence like ‘ X believes that it might be that p ’ reports the existence of a second-order belief state. But Yalcin notes there are reasons to doubt that is right. He develops in detail an expressivist model that avoids what he takes to be shortcomings of descriptivist approaches.

The two papers we haven’t discussed so far, those by Eric Swanson and Stephen Yablo, are both related to this expressivist family of theories, though their positive proposals head off in distinctive directions.

Eric Swanson's contribution locates epistemic modals within a broader category, which he calls "the language of subjective uncertainty". He also emphasizes the diversity of epistemic modal locutions, and draws attention to the risks involved in focusing too closely on just a few examples. In the literature so far, 'might' and 'must' have tended to get the lion's share of the attention, while other sorts of epistemic modality – including the more explicitly quantitative sorts ('four to one against that', 'there's a 55% chance that', etc.) – have gone mostly unnoticed. Swanson argues that attending to other instances of the language of subjective uncertainty serves to undermine many of the standard proposals about epistemic 'might' and 'must', and motivates a probabilistic semantics.

Somewhat relatedly, Stephen Yablo develops a theory about epistemic modals where their primary function is not to state facts about the world, but to update the conversational score. Theories of this kind are quite familiar from the dynamic semantics tradition, but Yablo notes that the existing dynamic theories of epistemic modals are quite implausible. One of the challenges a dynamic approach to epistemic modals faces is to say how we should update a context (or a belief state) with *It might be that p* when the context previously was incompatible with *p*. Yablo adopts some suggestions from David Lewis's "A Problem about Permission" (Lewis 1979) to try and solve this puzzle.

References

- DeRose, Keith. 1991. "Epistemic Possibilities." *Philosophical Review* 100 (4): 581–605. doi: 10.2307/2185175.
- Dietz, Richard. 2008. "Epistemic Modals and Correct Disagreement." In *Relative Truth*, edited by Manuel Garcia-Carpintero and Max Kölbel, 239–64. Oxford: Oxford University Press.
- Egan, Andy, John Hawthorne, and Brian Weatherson. 2005. "Epistemic Modals in Context." In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, edited by Gerhard Preyer and Georg Peter, 131–70. Oxford: Oxford University Press.
- Fintel, Kai von, and Anthony S. Gillies. 2008. "CIA Leaks." *Philosophical Review* 117 (1): 77–98. doi: 10.1215/00318108-2007-025.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Glanzberg, Michael. 2007. "Context, Content and Relativism." *Philosophical Studies* 136 (1): 1–29. doi: 10.1007/s11098-007-9145-5.
- Hacking, Ian. 1967. "Possibility." *Philosophical Review* 76 (2): 343–68. doi: 10.2307/2183640.
- Levi, Isaac. 1996. *For the Sake of the Argument: ramsey Test Conditionals, Inductive*

- Inference and Nonmonotonic Reasoning*. Cambridge: Cambridge University Press.
- Lewis, David. 1979. "A Problem about Permission." In *Essays in Honour of Jaakko Hintikka on the Occasion of His Fiftieth Birthday on January 12, 1979*, edited by Esa Saarinen, Risto Hilpinen, Illka Niiniluoto, and Merrill Provenge, 163–75. Dordrecht: Reidel. Reprinted in his *Papers in Ethics and Social Philosophy*, Cambridge: Cambridge University Press, 2000, 20–33. References to reprint.
- MacFarlane, John. 2003. "Future Contingents and Relative Truth." *The Philosophical Quarterly* 53 (212): 321–36. doi: 10.1111/1467-9213.00315.
- . 2005a. "Making Sense of Relative Truth." *Proceedings of the Aristotelian Society* 105 (1): 321–39. doi: 10.1111/j.0066-7373.2004.00116.x.
- . 2005b. "The Assessment Sensitivity of Knowledge Attributions." *Oxford Studies in Epistemology* 1: 197–233.
- . 2009. "Nonindexical Contextualism." *Synthese* 166 (2): 231–50. doi: 10.1007/s11229-007-9286-2.
- Partee, Barbara. 1989. "Binding Implicit Variables in Quantified Contexts." In *Papers from the Twenty-Fifth Regional Meeting of the Chicago Linguistic Society*, edited by Caroline Wiltshire, Randolph Graczyk, and Bradley Music, 342–56. Chicago: Chicago Linguistic Society. Reprinted in her *Compositionality in Formal Semantics*, Oxford, 2004.
- Quine, W. V. O. 1969. "Propositional Objects." In *Ontological Relativity and Other Essays*, 139–60. New York: Columbia University Press.
- Schlenker, Philippe. 2003. "Indexicality, Logophoricity, and Plural Pronouns." In *Afroasiatic Grammar II: Selected Papers from the Fifth Conference on Afroasiatic Languages, Paris, 2000*, edited by Jacqueline Lecarme, 409–28. Amsterdam: John Benjamins.
- Stalnaker, Robert. 1978. "Assertion." *Syntax and Semantics* 9: 315–32.
- Teller, Paul. 1972. "Epistemic Possibility." *Philosophia* 2 (4): 303–20. doi: 10.1007/bf02381591.

Published in *Epistemic Modality*, 2011, pp. 1–18.