

Relativism

Brian Weatherson

Patrick Shirreff

2017

Relativism is the view that the truth of a sentence is relative both to a context of utterance and to a context of assessment. That the truth of a sentence is relative to a context of utterance is uncontroversial in contemporary semantics. This chapter focuses on three points: whether the version of contextualism is vulnerable to the disagreement and retraction arguments, and if so, whether these problems can be avoided by a more sophisticated contextualist theory. The points include: whether relativism really does avoid the four problems posed for the other theories; and whether there are other theories that also avoid the problems, without running into the problems facing relativism or problems of their own. The chapter concentrates on two families of views that have been called relativist: Relativism about propositional truth; and Relativism about utterance truth.

Relativism, in the sense we're interested in here, is the view that the truth of a sentence is relative both to a context of utterance and a context of assessment. That the truth of a sentence is relative to a context of utterance is uncontroversial in contemporary semantics. If the authors of this paper were to both utter the sentence "I'm Canadian", then one of us would say something true, and the other something false. And that's because the truth of the utterance "I'm Canadian" is sensitive to a feature of the context of utterance, namely the identity of the speaker. And that in turn is explained by the fact that the proposition expressed by that sentence is different in different contexts. Relativists deny that this simple story can explain all the ways in which context language.

A core motivation for relativism comes from looking at the problems with other views. So let's start by thinking about what might be happening when a speaker says "This is tasty". (Call this utterance, which we'll come back to a bit, U, and its speaker S.) The demonstrative 'this' is context-sensitive, but let's assume it is clear what is being referred to, and think about what contribution the predicate is making. There are two natural proposals that are simple, easy to fit into familiar frameworks, and most likely wrong.

First, the predicate might pick out an objective property of tastiness, and when S utters U, she ascribes this property to a thing. This objective view has a number of problems. First, there is a metaphysical problem; just what is this objective tastiness? Second, there is an epistemological problem. Typically, a speaker like S could be prepared to utter U after taking one bite of the thing. But for most plausible answers to the first question, it is unclear how they could know that the object had that property.

The two problems reinforce each other; The more plausible answers to the metaphysical question make the property sensitive to very abstruse conditions, such as how idealised observers would react to ingesting the substance. But that makes it even more unlikely that typical utterances of U satisfy the epistemological requirement. (The points we're making here are well known, but our presentation owes a lot to Lasersohn (2005).)

Second, the predicate might be context-sensitive. For concreteness, let's focus on a very simple contextualist theory of 'tasty'. An utterance of U is true, relative to a context c , iff the speaker in c likes the taste of the referent of 'this' in c . So "This is tasty" and "I like the taste of this" express something very close to the same proposition. (At least they express propositions that are necessarily materially equivalent, but the relation between the two is probably closer than that.) Again there are two problems, both of them to do with the different dynamics of "This is tasty" and "I like the taste of this". A hearer H, who does not like the taste of this substance, could more readily disagree with "This is tasty" than with "I like the taste of this". If H doesn't like its taste, he'll nevertheless concede that S's utterance of "I like the taste of this" is true, but won't concede that "This is tasty" is true. And a similar phenomenon occurs when S herself changes her mind. If later she dislikes the taste of the substance, she will be disposed to retract her utterance of "This is tasty", but not of "I like the taste of this". So both disagreement and retraction data pose problems for the contextualist. (On disagreement, see especially Tamina Stephenson (2007). On retraction, and most everything else we'll talk about in this paper, see John MacFarlane (2014).)

If the objectivist and contextualist solutions fail, the relativist suggests that they have a useful alternative. Say that truth is doubly relativised, both to contexts of utterance and to context of assessment. So the utterance, or perhaps just the proposition expressed by it, is true relative to a context of assessment c_a and context of utterance c_u iff the agent at c_a likes the taste of the denotation of 'this' in c_u . So now the truth of the utterance is relative both to the context of utterance and the context of assessment. This solves the metaphysical problem; talk about what a person likes is unproblematic. It solves at least a version of the epistemological problem; a speaker knows what they like, so can make utterances that are true in their context. It solves a version of the disagreement problem; if H doesn't like the taste of the stuff, he can truly say that what S said is not true, since it isn't true in his context. And it explains the retraction data, since once S changes her taste, what she said is no longer true relative to her new context, and hence should be retracted.

As everywhere else in philosophy, arguments and claims can be and are contested. In the quick case for relativism we've described so far, there are at least four points where one could readily disagree.

1. Whether objectivism is really vulnerable to the combination of the metaphysical and the epistemological arguments.
2. Whether this version of contextualism is vulnerable to the disagreement and retraction arguments, and if so, whether these problems can be avoided by a more sophisticated contextualist theory.

3. Whether relativism really does avoid the four problems posed for the other theories.
4. Whether there are other theories that also avoid the problems, without running into the problems facing relativism or problems of their own.

In this entry, we'll focus on the last three points, since they are more widely discussed in the literature than the first one. And indeed, the last three points are extremely actively debated in recent years. We won't try to take sides in those debates, though we will note that on no point is the final picture nearly as rosy for the relativist as this initial sketch may suggest.

We've so far focussed on one particular predicate, 'tasty', though the points we've been making generalise to most predicates of personal taste. But this is far from the only area where relativist theories have been proposed. There has been a lot of discussion of relativist theories of epistemic modals. The arguments here are fairly similar to the arguments about predicates of personal taste (though see the discussion below about syntax and control). And MacFarlane has argued for relativist treatments of future contingents, and knowledge ascriptions. We will not spend much time on those constructions explicitly, but the issues they raise are broadly similar to the issues raised by predicates of personal taste, and by epistemic modals.

We will largely bypass two quite different areas where relativist theories have been proposed. Mark Richard (2008) has argued for a relativist treatment of comparative adjectives. And Weatherson (2011) argued for relativist treatments of areas of certain areas discourse where common assumptions about the area are false, and a relativist treatment might be the most charitable fix. As interesting as we find these proposals, they haven't occasioned nearly as much discussion as the proposals discussed above, and so we'll set them aside.

What we will do is first clarify some of the many ways in which the term 'relativism' has been used in recent debates, then review some technical material about indices and contexts that is essential for understanding some relativist views, then look at the main line of recent debate concerning relativism, the one centered on issues about retraction and disagreement, and finally look at some syntactic evidence for relativism.

1 Varieties of Relativism

As we noted in the very first paragraph, it is uncontroversial that sentence truth is sensitive to the context of utterance. It is extremely contentious just how many such sentences are sensitive to the context of utterance. (See, for instance (DeRose 2009; Harman and Thomson 1996; Cappelen and Lepore 2005) for some of the disputes.) But that there is some sensitivity here is uncontroversial. The view that a sentence type's truth is sensitive to its context of utterance is sometimes called "indexical relativism" (Kölbel 2004; Einheuser 2008). For some purposes this is a useful name. In particular, the view that the content of moral predicates, and hence the truth value of ascriptions of moral predicates, is sensitive to the context of utterance does seem like a form of

moral relativism, as that term is standardly understood. It is, for instance, the view that Harman (1975) defends in a self-proclaimed defence of moral relativism. (That this was traditionally known as relativism is stressed in Cappelen and Huvnes (2014), who draw some interesting conclusions from this fact.) But for present purposes we want to clearly exclude those views, and focus on much more radical proposals that have been recently made.

There are two (overlapping) families of views that have been called relativist, and which we will be concentrating on for the bulk of this entry. They are:

- Relativism about propositional truth - Whether a proposition is true is not an absolute fact. Propositions can be true relative to some contexts of assessment, and false relative to others.
- Relativism about utterance truth - Whether an utterance is true is not an absolute fact. Utterances can be true relative to some contexts of assessment, and false relative to others. It isn't always clear whether a particular author, in defending relativism, is primarily defending the first or second of these claims. But it is natural to interpret most relativists as defending relativism about propositional truth. This is most clearly true for Kölbel (2002) and Egan (2007), but there are few relativists that it is hard to interpret as taking this to be the primary focus. The one big exception, however, is the most prominent defender of relativism in the contemporary literature, John MacFarlane (2014). He takes relativism about utterance truth to be the only genuine form of relativism, though he also endorses relativism about propositional truth. We are not going to engage with the dispute over what is *really* a relativist view. Instead we'll just describe a pair of views that are relativist in one of these senses but not the other, to explain how the two senses come apart.

First, let's think about how to make propositional truth, but not utterance truth, relative. Assume that Suzy actually swims, but that she might not have. Then, on a natural way of thinking about modality, the proposition *Suzy swims* will be true relative to the actual world, but false relative to any possible world in which she does not. So propositional truth is relative, with the relativity being to possible worlds.

Cappelen and Hawthorne (2009) argue that the reasoning of the previous paragraph misunderstands the nature of truth and modality. The proposition *Suzy swims* is, they say, simply true. It might have been false. But that doesn't mean it is merely true relative to some world or other; it just means that it might have had a different property than it actually has. Schaffer (2012) argues that the things we say and think, i.e., propositions, typically have worlds in their content, so what we express by the sentence "Suzy swims" is the proposition *Suzy swims at @*, and that proposition, like all propositions, has a non-relative truth value.

Even if all that is wrong, and the contents of sentences like "Suzy swims" are propositions that are true in some worlds, false in others, there is a reason not to call this 'relativism'. After all, other worlds don't exist, or at least don't exist in the way ours does. So it isn't true that the proposition expressed by "Suzy swims" is true relative to

something and false relative to something else; there isn't a something else for it to be false relative to.

Taking the last point on board, a serious relativism about propositional truth should say that there is a proposition p , and two relata a and b such that p is true relative to a , and false relative to b . One way to do that, as seen in for example Kaplan (1989), is to make truth relative to times, and believe that other times exist. (The contrast here is with the eternalist view of propositions defended by Evans (1985).) But the radical view that has been central to many recent versions of relativism has been to say that truth is relative to world, time, individual triples. So there is a proposition, for example, that is true relative to a triple $\langle w, t, a \rangle$ iff in w , the individual a disapproves of murder at t . Perhaps, on a relativist framework, this just is the proposition that murder is wrong. The picture of propositions here owes something to the discussion of *de se* beliefs in Lewis (1979), since Lewis thought the contents of such beliefs were sets of such triples. But Lewis put this forward as a theory about mental content; the move to extending it to linguistic content is more recent.

On this view, the contents of utterances will be propositions that can be true relative to one individual (in a world, at a time) and false relative to another. But the utterances themselves will naturally be thought to have absolute truth value. Think back to the view that propositional truth is just relative to worlds. Then an utterance of "Suzy swims" will, intuitively, be true iff Suzy swims in the world the utterance is made. When we think about a counterfactual utterance of "Suzy swims" in a world where she does not, we don't think the utterance is true just because its content is actually true. So although this is a form of relativism about propositional truth, it is not yet a form of relativism about utterance truth. MacFarlane, who takes utterance relativism to be central, calls this view 'non-indexical contextualism', with the name meant to highlight that it isn't, by his lights, genuinely relativist.

Now turn to relativism about utterance truth. A simple way to implement this is to say that the propositional content of an utterance is relative to an assessor. So consider a very simple, and surely false, view about 'you' in English. It holds that the content of an utterance of 'you', relative to a context of assessment, is the agent of that context. And it says that the utterance is true at a context of assessment iff the proposition it expresses relative to that context is (absolutely) true. So an utterance of "You swim" will express a different proposition relative to contexts with different agents. And since some of those propositions will be true, and some false, we have a version of relativism about utterance truth. But this view is consistent with the view that propositional truth is absolute, not relative even to worlds.

So the two versions of relativism are dissociable. But it is also possible to hold them simultaneously, as MacFarlane (2014) does. To set up MacFarlane's view, it is helpful to review the framework against which it was developed.

2 Index, Context and Content

The existence of indexical terms, like ‘I’, means that we can’t give a theory for truth conditions of sentences as such. The sentence-type *I am in Ann Arbor now* doesn’t have truth conditions; only occurrences of this sentence do. So a natural move is to build up a theory that assigns to each term a function from contexts to contents, and use that to provide a theory of which contexts a sentence is true in. It turns out that there are reasons to endorse more complications than this. In particular, sentences like (1) show the need for relativising truth conditions relative to both a context and an *index*.

(1) It might have been the case that my actual parents never met

That’s true. But think of how we might naturally provide truth conditions for it. First, we define the context-relative truth conditions for “My actual parents never met.”. The context will provide a world for ‘actually’ (i.e., the actual world), and an agent for ‘my’ (i.e., the speaker). Then, we say that the modal operator shifts the world of evaluation. Intuitively, *Might p* is true iff *p* is true in some possible world. So we need some way to see if the content of “My actual parents never met” is true in some other world. But that’s hard to do, since (given origin essentialism), there will be no context in which “My actual parents never met” can be truly uttered, so its content will be the function that maps every context into FALSE. And that’s the same content as, for example, “Two plus two is five”, and so it can’t be possibly true.

There are actually two related problems here. One is that while we need to look around the other worlds to see whether some thing is true, the nature of that thing is fixed by the actual context. So if the sentence is uttered by Sasha Obama in this world, the content is something like *It might have been the case that Barack Obama and Michelle Obama never met*. We need to fix the values of the contextually sensitive terms (‘my’ and ‘actual’) even when they appear inside operators. The second problem is that something as coarse-grained as a function is very hard to ‘shift’ in just one respect.

There is a well-known solution to this, developed primarily by Hans Kamp (1971), then put to important philosophical use by David Lewis (1980) and David Kaplan (1989). Say that our semantic theory will assign a function from both contexts and indices to terms. Ultimately, each sentence will get as its semantic value a function from contexts and indices to truth values. Lewis describes the distinction between contexts and indices thus,

A context is a location – time, place, and possible world – where a sentence is said. It has countless features, determined by the character of the location. An index is an n-tuple of features of context, but not necessarily features that go together in any possible context. Thus an index might consist of a speaker, a time before his birth, a worlds where he never lived at all, and so on. (Lewis 1980, 79)

indices are structured, and so they can be ‘improper’; it might consist of things that don’t go together. If the index contains just worlds and individuals, the index could

be $\langle w, \text{Sasha Obama} \rangle$, even if Sasha doesn't exist in w . Such an index will play a central role in making (1), as uttered by Sasha Obama, true. A context, on the other hand, will by its nature be proper; it will pick out a world, and a time in that world, and a person existing at that time, and so on.

Indices are, unlike contexts, 'shiftable'. There can be sentential operators that say the truth value of the whole sentence is given by looking at the truth value of the embedded sentence at some different index. Assume, for example, that there is a world parameter in the index, so the index is $\langle w, x \rangle$, where x contains everything else in the index. Then *Might p* will be true relative to context c and index $\langle w, x \rangle$ just in case for some w' , p is true relative to context c and index $\langle w', x \rangle$. Note that in this definition, we do *not* shift c ; so if c determines that the referent of 'I' is Sasha Obama, that stays even if we 'shift' the index to a world where Sasha does not exist.

The separation between context and index was developed to solve some technical problems, but it can do some philosophical work. The theory associates each sentence with a function from contexts and indices to truth values. Equivalently, it associates sentence-context pairs with functions from indices to truth values. So it is natural to say that the content of an utterance is (or is at least intimately connected to) the function associated with the pair consisting of the sentence uttered and the context it was uttered in. If the index is simply a world, then the contents will be functions from worlds to truth-values, as defended by classical contextualists such as Robert Stalnaker (1984). David Lewis (1980) objected that assigning contents at just this stage was arbitrary, why not associate contents with functions from context to truth values instead? François Recanati (2007) responds well to Lewis's arguments, although his defence requires making indices much more complex. In particular, he includes parameters for times and individuals in the index, leading him to support relativism about propositional truth.

With this background, we can more easily describe two distinctive features of MacFarlane's views. First, he makes indices very complicated indeed. The index might include, among other things, an information set (for interpreting epistemic and deontic modals), a standard of taste (for interpreting predicates of personal taste), relevant alternatives (for interpreting knowledge ascriptions), a time (for handling future contingents), and perhaps many more things. For simplicity, we'll call these things collectively a *perspective*, and say that indices are world-perspective pairs. (Though note that our terminology here requires perspectives to be structured entities, and for them to be potentially improper.) If the propositional content of an utterance is (intimately connected to) a function from indices to truth-values, then the one proposition will be true relative to one perspective and false relative to another.

So far we only have relativism about propositional truth, or what MacFarlane calls non-indexical contextualism. MacFarlane argues for a second revision to contextualist orthodoxy. Assume that we have a particular utterance U of a sentence S in context of utterance c_w , and that utterance is being assessed in context c_a . And assume that $w(c)$ is the world of context c , and $p(c)$ is the perspective of context c . Sentence S is associated with a function f_S from context-index pairs, i.e., context-world-perspective triples, to

truth values. The relativist about propositional truth, but not utterance truth, says that U is true iff $f_S(c_u, w(c_u), p(c_u)) = \text{TRUE}$. MacFarlane's relativist¹ says that U is true iff $f_S(c_u, w(c_u), p(c_a)) = \text{TRUE}$. Crucially, we assess the utterance itself, and not just its content, by the perspective of the assessor. MacFarlane argues that this move allows a better understanding of disagreement and retraction, which were central to the phenomena that motivated relativism. So let's turn to how well relativism explains the phenomena it was designed to explain.

3 Retraction and Disagreement

3.1 Retraction

Contextualists have a difficult time explaining why we retract earlier claims involving predicates of personal taste or epistemically modals assertions when our perspective has changed in the intervening time. Consider this example:

Kim (age 8): Lunchables are delicious.

Kim (age 26 being reminded of previous assertion): I take it back/I was wrong/what I said was false. Lunchables aren't delicious.

When Kim retracts her earlier assertion, it is natural to use any of the three forms of retraction listed here. It isn't natural to say either of the following things:

- Lunchables were delicious back then, but they aren't delicious any more.
- When I said that back then, I only meant that they were delicious to me back then.

There is a contrast with how clearly context sensitive terms like 'here' are used. It isn't natural to use any of the three forms of retraction we attribute to Kim below.

Kim (in a café): It is pleasant here.

Kim (in an oil refinery, reminded of previous assertion): I take it back/I was wrong/What I said was false. It isn't pleasant here.

On the other hand, it is natural to say things like:

- It was pleasant where we were, but it isn't pleasant where we are now.

¹Strictly speaking, to get MacFarlane's exact view which kinds of views are relativist we need to complicate things even more. MacFarlane explicitly leaves open the possibility that we have a non-indexical contextualist view about some terms, and a relativist view about others, and expressly says that such a view is a form of relativism. So what we should say is that perspectives are sub-divided into those features where the context of utterance is relevant for utterance truth, and those where the context of assessment is relevant to utterance truth. If the first set of features of p_u , and the second set is p_a , then the utterance U of sentence S in context c_u will be true, relative to c_a , iff $f_S(c_u, w(c_u), p_u(c_u), p_a(c_a)) = \text{TRUE}$. MacFarlane's own view seems to be that p_u will be empty, so the simplified view we've given in the text is a fair representation of how he thinks actual natural languages work. But there are possible languages where the complications noted here are relevant.

- When I said that back then, I only meant it was pleasant where we were.

Predicates of personal taste, like ‘delicious’, don’t behave like explicitly context-sensitive expressions like ‘pleasant here’. We see the same pattern with epistemic modals.

Hakeem: It might be snowing outside.

Fabian: No/that’s wrong/that’s false, it can’t be snowing. I just looked out the window and there were clear skies and the sun was out.

Hakeem: Really? Then I guess I was mistaken.

If *It might be that p* is true iff the speaker’s knowledge is compatible with *p*, as it does on a simple contextualist theory, then none of this conversation makes any sense. But, argue relativists, it is perfectly natural. A natural move here is to argue that the problem with simple contextualism isn’t the contextualism, but the simplicity. Exploring the moves that can be made here will take us too far afield; MacFarlane (2014) goes through a number of possible alternative contextualist theories and shows how examples like the ones involving Kim, Hakeem and Fabian can be generated to raise problems for each of them. Instead, let’s look at how the relativist handles the cases.

According to MacFarlane, relativists are committed to the following principle about retraction:

The speaker ought to retract the assertion if she has good grounds for thinking that its content is false (as assessed from the perspective she currently occupies).

Given that Kim and Hakeem are evaluating their earlier utterances from new perspectives, perspectives where the assertions are now judged to be false, they both should retract that earlier assertion because they now take them to be false. But arguably this principle overgenerates. Consider this example from Fintel and Gillies (2008).

Alex: The keys might be in the drawer.

Billy: (Looks in the drawer, agitated) They’re not. Why did you say that?

Alex: Look, I didn’t say that they were in the drawer. I said they might be there – and they might have been. Sheesh. (Fintel and Gillies 2008, 81)

The lesson they draw from the example is that retraction is somewhat voluntary. Despite what relativists claim, it is not always true that ‘might’ claims are retracted or rejected in light of new evidence. Instead, what seems to be true is that that solipsistic readings for the modals are virtually always available. They say that the relativist can’t easily explain this data. MacFarlane (2014, Ch. 10) responds by saying, in effect, that epistemic modal claims sometimes mean what contextualists say they mean, and this is compatible with relativism.

3.2 Disagreement and Agreement

The data about retraction are closely related to a phenomenon that classically was central to debates about relativism. A traditional argument for relativism is that it is necessary to explain ‘faultless disagreement’. If one person says Vegemite is tasty, and another says that it is not, the alleged datum is that they are disagreeing, but neither is wrong. Objective treatments of taste save the phenomenon of disagreement, but lose the faultlessness. Contextualist, or otherwise subjective, treatments of taste save the faultlessness, but lose the disagreement. Relativism was alleged to keep both. Faultless disagreement plays a big role in Kölbel (2002), but it has dropped out of the recent literature somewhat. But the focus on disagreement remains, with the central claim being that contextualists cannot explain why some conversations are genuinely disagreements. The following is how MacFarlane puts the point using an example from predicates of personal taste:

If the truth of my claim that a food is “tasty” depends on how it strikes me, while the truth of your claim that the same food is “not tasty” depends on how it strikes you, then our claims are compatible, and we do not disagree in making them. But it seems that we do disagree—even if we are aware that the source of our disagreement is our differing tastes. (MacFarlane 2014, 8)

This is no argument against any kind of objectivist treatment of predicates of personal taste. The issue is solely whether relativism or contextualism does a better job of explaining the facts about disagreement. We’ll look at a couple of reasons for thinking that the problem is less pressing for contextualists than it might first seem, then an interesting attempt to resolve the problems that remain for contextualism, then at some reasons for doubting that relativism helps solve the remaining problems.

3.3 Clarifying the Data

The puzzle for contextualism is supposed to be that there is a big difference between the felicity of Mark’s reply in the following two cases.

Sally: This chilli is tasty.

Mark: I disagree. It’s too hot.

Sally: I’m from Barcelona.

Mark: I disagree. #I’m from Oslo.

Assume Sally and Mark are both being sincere. So Sally does like the chilli and is from Barcelona, and Mark doesn’t like it and is from Oslo. In neither case can Mark sincerely repeat the words that Sally uttered. Indeed, he can sincerely utter the negations of the sentences Sally uttered. But in the first case, this seems to amount to a disagreement, and in the second case it does not. If the context-sensitivity of ‘tasty’ and ‘I’ is explained the same way, this is mysterious.

But note that there are other examples that do not seem to be amenable to a relativist treatment where Mark can express disagreement with Sally.

Sally: I like this chilli.

Mark: I disagree. It is too hot.

What seems to be going on there is that Mark is disagreeing with an attitude that Sally has, but not with any proposition she expressed. After all, Mark presumably agrees with the proposition that Sally likes the chilli. That's why he knows he is disagreeing with her. And that's the content of what he uttered. So some disagreements that are triggered by assertions are not with the content of what is asserted. This might be the germ of an idea for how contextualists can explain the data about disagreement, one nicely developed by Torfinn Huvenes (2012). He argues that in a lot of cases of disagreements involving taste, what we see is not a disagreement about any content, but a disagreement of attitudes.

Two parties disagree just in case there is something towards which they have conflicting attitudes. This sometimes means that there is a content that one party accepts and the other rejects, but that does not always have to be the case. Just as two parties may have conflicting beliefs, they may also have conflicting desires or preferences. (Huvenes 2012, 178)

If Huvenes is right, then we can't draw any conclusions for semantics from the facts about disagreement.

There is a further problem with using facts about disagreement to argue against contextualism. Consider the following dialogue.

Sally: Joe might be in Boston.

Mark: I disagree; he's definitely in China.

If *Joe might be in Boston* just means *For all I know, Joe is in Boston*, then it is prima facie unclear why Mark is disagreeing. After all, it is consistent with Joe's being in China that for all Sally knows, Joe is in Boston. But, say some contextualists, Mark need not disagree with the whole content of Sally's utterance. He may just be disagreeing with the *prejacent* of the modal, i.e., that Joe is in Boston.

Most of the terms of disagreement that have been used in examples motivating anti-contextualism can be seen to target something other than the entire proposition (Fintel and Gillies 2008, 83). This is a particular problem for arguments for relativism involving epistemic modals. It's possible to say "I disagree", "that's false", "no", or "you're mistaken" and disagree with the prejacent of someone's modal claim, not necessarily the whole claim.

One potential avenue for the relativist to get around this worry is to limit the range of disagreement markers that count as expressing the right kind of disagreement. For example, Tamina Stephenson (2007) limits the terms of disagreement in the examples

she uses to “no” and “nuh-uh”. And John MacFarlane (2014, 11) discusses disagreements that more explicitly target the entire asserted proposition such as “the proposition you expressed is false” and “what you asserted is false”. The issue with the two former disagreement markers is that they can explicitly target the embedded clause of an expression. Indeed, Stephenson herself provides a clear example of this phenomenon.

Mary: How’s the cake?

Sam: I think it’s tasty.

Sue: Nuh-uh, it isn’t tasty at all! (Stephenson 2007, 512)

The issue with the disagreement markers that MacFarlane uses is that they have become too technical to do the type of work they need to do. Relativism is meant to be an empirical thesis that relied on natural language data to back it up. We have moved well outside of the realm of natural language and the types of natural language intuitions about the acceptability of sentences that we can get when we move to “the proposition you expressed is false” and “what you asserted is false”.

Note that the two contextualist responses we’ve described here are rather complementary. The point Huvenes makes, that disagreements can involve attitudes other than belief, seems best served to defuse arguments from disagreements concerning predicates of personal taste. And the point that von Fintel and Gillies make, that disagreements can target prejacent, seems best served to defuse arguments from disagreements concerning epistemic modals. It is possible there are replies to this last point; Weatherston and Egan (2011) for example suggest that examples involving agreements are invulnerable to the response that von Fintel and Gillies make. But as it stands, the dominant trend in the literature seems to be in the direction of thinking the contextualist has the resources to answer these relativist arguments.

3.4 Presuppositions and Common Context

There is a more radical, and perhaps more concessive, response to the disagreement arguments available to the contextualist. Dan López de Sa (2008) develops a contextualist theory that explains the disagreement data by positing that for many context-sensitive terms, there is a presupposition that users of the term are in the same context. (Note that López de Sa calls his theory an “indexical relativist” theory, but it is a kind of contextualist theory in the way the terms are being used here.)

It isn’t true in general that there is a presupposition of commonality of context. If two speakers both say “I am happy”, they should be interpreted as putting forward different propositions. And that is because, in the relevant sense, they are in different contexts. But, perhaps, many terms are not like this. In particular, cases where there appears to be a problem with explaining the phenomena involving disagreement, perhaps this is not so. So consider the stock example López de Sa uses, a variant on one that may seem familiar by now.

Hannah: Homer Simpson is funny.

Sarah: I disagree. Homer is not funny.

If ‘funny’ denotes a different property when Hannah and Sarah use it, there is no disagreement about propositional content here. And the contextualist account of predicates of personal taste would predict that it could, and perhaps often will, denote a different property on different occasions of usage. But it seems that there is a disagreement here, and arguably even one about propositional content. The solution López de Sa offers is that in any conversation, there is a presupposition that we are applying the same aesthetic standards. The model he uses is a suggestion made by David Lewis (1989) in defending a contextualist treatment of claims about value.

Wouldn't you hear them saying 'value for me and my mates' or value for the likes of you'? Wouldn't you think they'd stop arguing after one speaker says X is a value and the other says it isn't? – Not necessarily. They might always presuppose, with more or less confidence (well-founded or otherwise), that whatever relativity there is won't matter in this conversation. (Lewis 1989, 84)

Here is how López de Sa develops the point.

According to the approach, ‘is funny’ triggers a presupposition of commonality to the effect that both Hannah and Sarah are similar with respect to humour. Thus, in any non-defective conversation where Hannah uttered ‘Homer is funny’ and Sarah replied ‘No, it is not,’ it would indeed be common ground that Hannah and Sarah are relevantly alike, and thus that they are contradicting each other. After all, provided they are alike, either both Hannah and Sarah are amused by Homer or they are not. (López de Sa 2008, 305)

This is an ingenious idea, but there are a few hurdles to be cleared before it could be declared a full solution to the problem. First, it needs a way to deal with eavesdroppers. If Hannah writes “Homer Simpson is funny” on a scrap of paper, and later Sarah chances upon that paper, she can still say that she disagrees. But it is very odd to think there is a presupposition of commonality of taste with anyone who chances upon one's writings. Second, it needs to account for cases where the presupposition is expressly cancelled. The Hannah/Sarah dialogue feels natural even in cases where it has been made explicit that Hannah and Sarah have completely different tastes, and they are displaying their differences for the amusement of their friends (MacFarlane 2014, 131–32). Third, it doesn't quite capture the idea that Hannah and Sarah are contradicting each other. When Sarah disagrees, it shows that either there is a proposition one accepts and the other rejects, or that the context is defective. If we thought the data was that there was a contradiction between what they say, López de Sa's approach can't explain that. And finally, it isn't obvious how to extend this theory to other terms for which relativism seems promising. It is one thing to say that conversations about humour presuppose a common standard for humour. It is much less plausible to say that conversations about what might be the case presuppose that the parties to the conversation know the same

things. None of these hurdles seem impossible to clear, but they do raise doubts about whether presuppositions can solve all the contextualists' problems with disagreement.

3.5 Relativism and Disagreement

If, after all that, we conclude that the contextualist still has a problem with disagreement, it's fair to ask whether the relativist does any better. Let's think again about López de Sa's example of Hannah and Sarah.

Hannah: Homer Simpson is funny.

Sarah: I disagree. Homer is not funny.

A simple relativist theory assigns truth conditions to Hannah's utterance relative to a context of utterance, her own, and an index that consists of a world-perspective pair. (Remember that we're using 'perspective' to cover everything other than a world that a relativist may want to put into an index, and that it will be a structured entity.) The content of Hannah's utterance will be set by her context. So it will be (or at least will determine) a function from indices, i.e., world-perspective pairs, to truth-values. Call the world Hannah and Sarah occupy w , and their perspectives p_H and p_S . Then the proposition that Homer Simpson is funny will be true relative to $\langle w, p_H \rangle$ and false relative to $\langle w, p_S \rangle$. So there is a proposition Hannah accepts and Sarah rejects, and so they disagree. Doesn't think mean that the relativist can explain the sense in which Hannah and Sarah disagree?

Not so fast. Consider another case involving One, who lives in w_1 where Mars has one moon, and Two, who lives in w_2 where Mars has two moons. They make the following utterances:

One: Mars has one moon.

Two: Mars has two moons.

Now most theorists would say that One and Two have expressed propositions that cannot be true together. (As noted earlier, Schaffer (2012) disagrees, though not in a way that helps relativism.) But they don't disagree. Among other things, they both think that the other speaks truly.

Hannah doesn't just think that the proposition that Homer Simpson is funny is true relative to $\langle w, p_H \rangle$; she thinks it is simply true. That is because she occupies (for want of a better word) the index $\langle w, p_H \rangle$. And one doesn't just think that the proposition that Mars has one moon is true relative to $\langle w_1, p_1 \rangle$, she thinks it is simply true. But in doing so, she need not be in disagreement with someone who occupies a different index, such as Two, and who thinks it is false. It isn't easy to read off the existence of disagreement from the endorsement of conflicting propositions when the parties occupy different indices. And that raises doubts about whether the relativist has really explained the disagreement. (The Mars example is from MacFarlane (2014, 128). Both Dreier (2009) and Francén (2010) raise doubts about whether the relativist can explain disagreement.)

Part of MacFarlane's response to this is to insist that the different elements of an index are treated differently. If we are relativists about propositional truth, but not utterance truth, it is natural to treat the two elements the same. A speaker, we'll then say, speaks truly iff the proposition they express is true relative to the context they occupy. But that isn't MacFarlane's view. He holds that an utterance is true, i.e., that a speaker speaks truly, relative to a context of assessment iff the proposition they express is true relative to the world they occupy, and the perspective of the context of assessment. This gives him a way to distinguish the Hannah/Sarah case from the One/Two case. There is still a lot more work to do to turn this into a full theory of disagreement, and chapter 5 of MacFarlane's book has a very careful study of the varieties of disagreement that are possible on a relativist theory, and how they can be used to explain the data. We're not going to attempt to evaluate the success of these responses, but rather conclude by noting that even if the contextualist has work to do to explain the phenomena involving disagreement, so does the relativist.

3.6 Problems with the Data

Most of the arguments in the literature have started with intuitions about disagreement. We don't think there is anything wrong with this in principle; indeed, it is what we've done so far. But when the intuitions get a little shaky, as they are in a few cases we've described so far, it is worth checking them more carefully, against a broader range of informants. And when that is done, it isn't clear that the data help the relativists as much as the relativists have claimed. Knobe and Yalcin (2014) provide evidence that the following claim, which we'll call (J), isn't as empirically justified as the relativists have made it out to be.

(J) Competent speaker/hearers tend to judge a present-tense bare epistemic possibility claim (BEP) true only if the prejacent is compatible with their information (whether or not they are the producer of that utterance); otherwise the BEP is judged false.

They argue that many relativists, in particular Egan and MacFarlane, are committed to this claim. Although Egan and MacFarlane differ on several points (Egan takes relativism about propositional truth to be primary, MacFarlane relativism about utterance truth), it does seem true that (J) is important to both of them. As Knobe and Yalcin put it,

Egan and MacFarlane are both clearly animated by the thought that "people tend to assess epistemic modal claims for truth in light of what they (the assessors) know, even if they realize that they know more than the speaker (or relevant group) did at the time of utterance" (MacFarlane 2011: 160; see also Egan 2007: 2-5, the section entitled "Motivation for relativism: eavesdroppers"). (Knobe and Yalcin 2014, 3-4)

This isn't what their data showed. The subjects were shown speakers whose evidence strongly, but falsely, suggested that Fat Tony was dead. They overwhelmingly said that

an utterance of “Fat Tony is dead” is false, but most said an utterance of “Fat Tony might be dead” was true. (Though it is worth noting that the responses displayed a considerably ambivalence; the answers weren’t in line with what either a contextualist or a relativist would straightforwardly predict.) The subjects did say that it would be correct for the speaker who said “Fat Tony might be dead” to retract that utterance once it was clear Fat Tony was alive. But this typically wasn’t because they thought the earlier utterance was false.

The point of this study was not to directly target intuitions about disagreement, but rather inter-contextual judgments and felicity of retraction. But the issues are closely related. If subjects who know Fat Tony is alive don’t judge that an utterance of “Fat Tony might be dead” is false, then either they don’t disagree with such an utterer, or the disagreement is, as Huvenes suggests, not the kind that motivates altering our semantics in the direction relativists suggest.

As Knobe and Yalcin are careful to note, even if relativists were completely wrong about inter-contextual evaluation of utterances, about disagreement, and about retraction, there are still other arguments for relativism. We’ll end with one other such argument, due primarily to Tamina Stephenson (2007).

4 Control and Syntax

There is a striking semantic/syntactic phenomenon that epistemic modals and predicates of personal taste share. It’s easiest to describe the phenomenon if we assume a contextualist semantics, though we’ll eventually use the puzzle to cast doubt on that semantics. So assume, for now, that *It must be that p* is true iff *p* is guaranteed to be true by the (possibly idealised) knowledge of *X*, where *X* is an individual or group supplied by context. (And perhaps the amount of idealisation is context-sensitive too.) And assume that *F is tasty* is true iff *F* tastes good to (the possibly idealised version of) *X*, where *X* is again supplied by context.

When ‘might’ or ‘tasty’ are not embedded, then *X* seems like it has to include the speaker, and perhaps not much more than that. Even making some other taster or knower salient does not suffice to change the value of *X*. For instance, if one utters “Joe is a great cook and connoisseur of fine food. The meals he prepares are always tasty,” the last sentence is not naturally construed as saying that *Joe* always likes the taste of what Joe cooks, but that the speaker and her hearer do (or will). Or if, to borrow an example from (Weatherson 2011), we say “Jones must know who the killer is,” the relevant *X* for interpreting the ‘must’ consists of the speaker and her hearers, not Jones. There is something strange about this; it normally isn’t that hard to make others relevant in a way that makes them the value of a contextually filled variable.

Or, perhaps, it isn’t normally that hard with one key exception. Some context-sensitive terms have, as part of their meaning, that their extension includes the speaker. Such terms include ‘I’ and ‘we’. Perhaps ‘must’ and ‘tasty’ are like ‘I’ and ‘we’ in this respect. Except, and here is Stephenson’s key insight, there is a big difference between ‘I’/‘we’ and ‘must’/‘tasty’. The former still pick out the speaker (and perhaps those

near her) under the scope of an attitude verb. The latter do not. Consider the natural interpretations of these sentences.

- Joe thinks my gumbo is tasty.
- Joe thinks we must have stolen the candy.

In each case, the personal pronouns ('my', 'we') have their customary denotations. It is the speaker's gumbo that is being praised, and the speaker and her friends who are being suspected of theft. But 'tasty' and 'must' do not behave like that. It isn't that Joe thinks the speaker likes the taste of her gumbo, but that Joe himself does. And it isn't that Joe thinks the speaker's evidence entails her guilt, but that Joe's evidence does.

Stephenson points out that we get even more dramatic results with more complex sentences. Consider these two sentences.

- Mary thinks that Sam thinks it must be raining.
- Mary thinks that Sam must think it is raining. (Stephenson 2007, 490)

The value of X for the first 'must' has to be Sam, and for the second 'must' it has to be Mary. In neither case is it the speaker, and in neither case is it even particularly optional how to interpret the 'must'. Compare *Mary thinks that Jane likes her house*, which is naturally read as three-way ambiguous. (The house might be Jane's, or Mary's, or a contextually supplied third party's.) The general point here is that the X values that the contextualist posits behave rather unlike other implicit or explicit context-sensitive terms.

The relativist explanation of this is that epistemic modals, and predicates of personal taste are, to use Stephenson's phrase, "inherently judge-dependent". That is, they are inherently dependent on some perspective for their truth. In the terms we've been using so far, we need a perspective in the index, and not just in the context, to explain the truth of these claims. When this is combined with a natural view that attitude predicates obligatorily shift the perspective (or judge) parameter, then it just naturally falls out that epistemic modals must take on the perspective of the immediate subjective of the attitude verb. Such a semantics is defended by Stephenson, and by Peter Lasnik (2005).

It falls out of this semantics that there cannot be "exocentric" uses of bare epistemic modals. These are uses of bare epistemic modals that take on the perspective of someone other than the speaker. (A similar point applies to predicates of personal taste.) This is a nice explanation of the fact that it is very hard to generate these exocentric uses. But it arguably overgenerates; there are cases where it seems we do get the exocentric reading. Here is one case from Egan, Hawthorne, and Weatherston (2005).

Ann is planning a surprise party for Bill. Unfortunately, Chris has discovered the surprise and told Bill all about it. Now Bill and Chris are having fun watching Ann try to set up the party without being discovered. Currently Ann is walking past Chris's apartment carrying a large supply of party hats. She sees a bus on which Bill frequently rides home,

so she jumps into some nearby bushes to avoid being spotted. Bill, watching from Chris's window, is quite amused, but Chris is puzzled and asks Bill why Ann is hiding in the bushes. Bill says, "I might be on that bus" (Egan, Hawthorne, and Weatherson 2005, 140)

The natural reading of what Bill says is that for all *Ann* knows, Bill is on the bus. And this is predicted to be impossible on the type of relativist view that gets us the correct results in the obligatory control cases. Stephenson suggests that there is an ellipsis in Bill's sentence. It should really be understood as:

- Ann is hiding in the bushes because I might be on that bus.

So in a sense, it isn't a bare epistemic modal; it is in a 'because'-clause. Moreover, suggests Stephenson, we should take 'because'-clauses to express something like a person's conscious reasoning or rationale. So in this case 'because' acts like an attitude verb and shift the perspective (or judge) of the epistemic modal. (A similar move, in response to a similar objection, is defended by John MacFarlane (2014, 272ff).)

The relativist needs two premises here for the defence to work. The first is that all these exocentric uses are either in the scope of an attitude verb, or are in an explanatory context. The second is that it is fair to treat 'because' as sufficiently like an attitude verb for these purposes. A contextualist could well object to either assumption. But even if they grant both assumptions, a contextualist may want to simply resist the whole line of reasoning. At most what these arguments about control show is that 'might', and 'tasty', behave rather differently to other context-sensitive terms. There is nothing inconsistent about just accepting that as a surprising fact. And given how radical a thesis relativism seems to many, accepting relativism as an explanation for these facts about control could well be an excessive reaction. The issues here have not been worked out in nearly as much detail as the issues concerning disagreement and retraction, and we are a long way from having a full accounting of the costs and benefits of the possible dialectical moves.

References

- Cappelen, Herman, and John Hawthorne. 2009. *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Cappelen, Herman, and Torfinn Tomesen Huvnes. 2014. "Relative Truth." In *The Oxford Handbook of Truth*, edited by Michael Glanzberg. Oxford: Oxford University Press.
- Cappelen, Herman, and Ernest Lepore. 2005. *Insensitive Semantics: A Defence of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- DeRose, Keith. 2009. *The Case for Contextualism: Knowledge, Skepticism and Context*. Oxford: Oxford.
- Dreier, James. 2009. "Relativism (and Expressivism) and the Problem of Disagreement." *Philosophical Perspectives* 23: 79–110. doi: 10.1111/j.1520-8583.2009.00162.x.

- Egan, Andy. 2007. "Epistemic Modals, Relativism and Assertion." *Philosophical Studies* 133 (1): 1–22. doi: 10.1007/s11098-006-9003-x.
- 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.
- Einheuser, Iris. 2008. "Three Forms of Truth-Relativism." In *Relativising Utterance Truth*, edited by Manuel Garcia-Carpintero and Max Kölbel, 187–203. Oxford: Oxford University Press.
- Evans, Gareth. 1985. "Does Tense Logic Rest on a Mistake?" In *Collected Papers*, 343–63. Oxford: Clarendon Press.
- Fintel, Kai von, and Anthony S. Gillies. 2008. "CIA Leaks." *Philosophical Review* 117 (1): 77–98. doi: 10.1215/00318108-2007-025.
- Francén, Ragnar. 2010. "No Deep Disagreement for New Relativists." *Philosophical Studies* 151 (1): 19–37. doi: 10.1007/s11098-009-9414-6.
- Harman, Gilbert. 1975. "Moral Relativism Defended." *Philosophical Review* 84 (1): 3–22. doi: 10.2307/2184078.
- Harman, Gilbert, and Judith Jarvis Thomson. 1996. *Moral Relativism and Moral Objectivity*. Cambridge, MA: Blackwell.
- Huvenes, Torfinn Thomesen. 2012. "Varieties of Disagreement and Predicates of Taste." *Australasian Journal of Philosophy* 90 (1): 167–81. doi: 10.1080/00048402.2010.550305.
- Kamp, Hans. 1971. "Formal Properties of 'Now'." *Theoria* 37: 227–74.
- Kaplan, David. 1989. "Demonstratives." In *Themes from Kaplan*, edited by Joseph Almog, John Perry, and Howard Wettstein, 481–563. Oxford: Oxford University Press.
- Knobe, Joshua, and Seth Yalcin. 2014. "Epistemic Modals and Context: Experimental Data." *Semantics and Pragmatics* 7 (10): 1–21. doi: 10.3765/sp.7.10.
- Kölbel, Max. 2002. *Truth Without Objectivity*. London: Routledge.
- . 2004. "Indexical Relativism Versus Genuine Relativism." *International Journal of Philosophical Studies* 12 (3): 297–313. doi: 10.1080/0967255042000243966.
- Laserson, Peter. 2005. "Context Dependence, Disagreement and Predicates of Personal Taste." *Linguistics and Philosophy* 28 (6): 643–86. doi: 10.1007/s10988-005-0596-x.
- Lewis, David. 1979. "Attitudes *de Dicto* and *de Se*." *Philosophical Review* 88 (4): 513–43. doi: 10.2307/2184646. Reprinted in his *Philosophical Papers*, Volume 1, Oxford: Oxford University Press, 1983, 133–156. References to reprint.
- . 1980. "Index, Context, and Content." In *Philosophy and Grammar*, edited by Stig Kanger and Sven Öhman, 79–100. Dordrecht: Reidel. Reprinted in his *Papers in Philosophical Logic*, Cambridge: Cambridge University Press, 1998, 21–44. References to reprint.
- . 1989. "Dispositional Theories of Value." *Aristotelian Society Supplementary Volume* 63 (1): 113–37. doi: 10.1093/aristoteliansupp/63.1.89. Reprinted in his *Papers in Ethics and Social Philosophy*, Cambridge: Cambridge University Press,

- 2000, 68-94. References to reprint.
- López de Sa, Dan. 2008. "Presuppositions of Commonality." In *Relativising Utterance Truth*, edited by Manuel Garcia-Carpintero and Max Kölbel, 297–310. Oxford University Press.
- MacFarlane, John. 2014. *Assessment Sensitivity*. Oxford: Oxford University Press.
- Recanati, François. 2007. *Perspectival Thought: A Plea for (Moderate) Relativism*. Oxford: Oxford University Press.
- Richard, Mark. 2008. *When Truth Gives Out*. Oxford: Oxford University Press.
- Schaffer, Jonathan. 2012. "Necessitarian Propositions." *Synthese* 189 (1): 119–62. doi: 10.1007/s11229-012-0097-8.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stephenson, Tamina. 2007. "Judge Dependence, Epistemic Modals, and Predicates of Personal Taste." *Linguistics and Philosophy* 30 (4): 487–525. doi: 10.1007/s10988-008-9023-4.
- Weatherson, Brian. 2011. "No Royal Road to Relativism." *Analysis* 71 (1): 133–43. doi: 10.1093/analys/anq060.
- Weatherson, Brian, and Andy Egan. 2011. "Epistemic Modals and Epistemic Modality." In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson, 1–18. Oxford: Oxford University Press.

Published in *Blackwell Companion to the Philosophy of Language*, 2017, pp. 787-803.