

Relativism

December 7, 2011

Relativism, in the sense at issue here, is the view that some assertions or beliefs can be characterized as accurate, or objectively correct, only relative to a context of assessment. We will begin by looking at a few cases that make relativist views tempting, then consider the technical and philosophical issues such views raise.

MOTIVATION

Epistemic modals

“Joe might be in Boston,” says Maria. In doing so, she expresses a certain state of mind, making it clear that she does not take her information to rule out Joe’s being in Boston. But does she *assert* this? There are reasons to think not. George, who knows that Joe is just down the hall, might disagree with what Maria has said. And, after Maria learns that Joe is in fact in Berkeley, she too may think that she was wrong, and retract her claim: “I said he might be in Boston, but that’s wrong—he can’t be in Boston, since George just saw him down the hall.”

Noticing these things, one might conclude that Maria hasn’t asserted anything, but has only expressed her uncertainty. But suppose one does want to think of her speech act as an assertion. What are the options? One could explain the disagreement and retraction facts by taking her to have made a claim about what is left open by the information of a broader group, including her possible interlocutors. But then it becomes hard to understand how she could have taken herself to be warranted in making this claim. We need to understand how it is that she can legitimately assert that Joe might be in Boston on the basis of nothing more than her own ignorance, while others—and her own later self—can evaluate what she has said in relation to *their* information.

All of this could be explained if we could make sense of the idea that the accuracy of Maria’s assertion is not an absolute matter, but varies with perspective. We could then say that it is the *assessor’s* information, not the speaker’s, that matters for the accuracy of an epistemic modal claim. That would explain why,

when Maria learns that Joe is right down the hall, she takes her original claim that he might be in Boston to have been inaccurate, and retracts it. In this way we might capture the subjectivity of epistemic modal claims—the way in which their accuracy depends on what is known, not on any objective feature of the world—without construing them (as contextualist views do) as claims about what some particular person or group knows. (See Egan, Hawthorne, and Weatherson 2005, Egan 2007, Stephenson 2007, MacFarlane 2011; for critical remarks, see Wright 2007, von Fintel and Gillies 2008, Dietz 2008.)

Taste predicates

“This vegemite sandwich is tasty,” says Joe. In doing so, he expresses his liking for the sandwich, but does he do anything more? Does he assert something, and if so, on what does the accuracy of his assertion depend?

Contextualists hold that Joe has asserted something about the *relation* between the food and himself, or between the food and a certain body of tastes. But then it is puzzling why Sasha, whose tastes are very different, should reject Joe’s assertion. On a relativist view, by contrast, Joe’s assertion can only be classed as accurate or inaccurate relative to a context of assessment, and it is the assessor’s tastes that matter for its accuracy, rather than the speaker’s. If Sasha despises the taste of vegemite, while Joe loves it, then Joe’s assertion will be inaccurate as assessed by Sasha, but accurate as assessed by Joe. And, if Joe later comes to hate the taste of vegemite, his earlier belief will be inaccurate relative to his current context, and he will rightly consider it to have been wrong. In this way, the relativist tries to capture the subjectivity of taste claims—the way in which their accuracy depends on a sense of taste—without reducing them to claims about how various foods relate to the tastes of a person or group. (See Kölbel 2002, Lasersohn 2005, 2009, Stephenson 2007, MacFarlane 2007; and for criticism, Glanzberg 2007, Cappelen and Hawthorne 2009, ch. 4.)

Future contingents

“There will be a sea battle tomorrow,” says Themistocles, at a time when it is objectively indeterminate whether there will be a sea battle the next day. Is his assertion accurate or inaccurate? The question can only be answered, according to a relativist view, relative to a particular context of assessment. Themistocles’ assertion is inaccurate as assessed from the moment at which it is made (m_o), accurate as assessed from a moment one day later on a possible future history with a sea battle (m_1), and inaccurate as assessed from a moment one day later on a possible future history without a sea battle (m_2). In this way the relativist splits the difference between the supervaluationist, who holds that the assertion is (absolutely) inaccurate—and has trouble explaining why it need not be retracted by the asserter at m_1 —and the “true futurist,” who holds that the assertion is either accurate or inaccurate absolutely, even though no

objective feature of the situation breaks the symmetry between the two possible histories. (See MacFarlane 2003, 2008; and for criticism, Heck 2006.)

Such proposals raise two foundational questions in semantics:

The technical question What revisions are needed to standard truth-conditional semantic frameworks to make room for relativist accounts of this kind?

The philosophical question How can we make philosophical sense of the idea that the accuracy of an assertion or belief is assessment-relative? Do we really understand what such proposals *say*? How do relativist theories differ from nonrelativist theories in their predictions about language use?

In what follows, we will briefly consider how each of these questions might be answered.

THE TECHNICAL QUESTION

Index and context

The task of a semantic theory, according to Lewis (1980), is to define truth at a context of use for sentences of a language L . Why is this our target notion? Because if we know the conditions under which sentences of L are true at a context, we know enough to understand and communicate with speakers of L . One counts as a speaker of L in virtue of participating in a convention of uttering only true sentences of L and trusting other L -speakers to do the same. Thus, if we know that a speaker of L utters S in context c , and we know that in L , S is true at c if and only if p , we can presume that the speaker believes that p and intends for us to believe this too. And if we want to communicate that q to another speaker, we have only to select a sentence T that is mutually known to be true in our context if and only if q .

Because languages contain infinitely many sentences, and we want a finite characterization of our semantic knowledge, our definition of ‘true at c ’ must be recursive. However, it is not in general possible to state the condition for a sentence to be true at a context in terms of the conditions for its constituents to be true at a context. One reason is that the constituents of some complex sentences are not sentences but open formulas. Another reason is that a proper treatment of modal and temporal operators requires considering the truth value of the embedded sentence at times and worlds other than the time and world of the context. For example, to evaluate \ulcorner Yesterday: ϕ \urcorner we need to consider the truth value of ϕ relative to the day before the present day. If we are recursively defining truth at a context, we will need to evaluate ϕ at a context c' that differs

from our current context c only in the time of the context. However, as Lewis points out (1983, 29), there is no guarantee that there will be such a context. Contexts represent possible occasions of use, and are not arbitrary collections of parameters. Since the agent of the context must exist at the time and world of the context, if we shift the time back before the agent of the context existed, there will be no corresponding context. Moreover, even if there *were* a context c' that differed from c only in the time of the context, we would not want to say that \ulcorner Yesterday: $\phi \urcorner$ is true at c iff ϕ is true at c' . For ϕ might contain temporal indexicals like ‘today’ or ‘now’, and these must be evaluated with respect to the original context c . ‘Yesterday it was warmer than it is now’ can be true, even though the embedded sentence ‘It is warmer than it is now’ is not true relative to *any* context (Kamp 1971, Kaplan 1989).

The standard solution to these problems is to recursively define truth at a context and an *index*, which is simply a collection of independently shiftable parameters (world, time, assignment of values to the variables, and so on). The clause for ‘Yesterday’ could be:

Yesterday \ulcorner Yesterday: $\phi \urcorner$ is true at $c, \langle w, t, a \rangle$ iff ϕ is true at $c, \langle w, t', a \rangle$, for some t' belonging to the day before t .

Although what we really care about at the end of the day is truth at a context, a recursive definition of truth at a context and index can serve our purposes, as long as we can recover a definition of truth at a context from it. And we can:

Lewisian postsemantics A sentence ϕ is true at c iff ϕ is true at $c, \langle w_c, t_c, a \rangle$, for any assignment a . (Here w_c and t_c are the world and time of the context c .) (cf. Lewis 1983, 31)

MacFarlane (2003, 2005a) calls this step the *postsemantics*, to distinguish it from the recursive definition of truth at a context and index, which we might call the *compositional semantics*.

Assessment sensitivity

As Lewis makes clear, the role of the index in such a theory is entirely technical. If there is no operator that shifts a parameter, then the parameter is not needed in the index. And, if there *is* an operator that shifts the parameter, no further justification is needed for putting that parameter in the index. The index is just part of the machinery for giving a systematic definition of truth at a context; it plays no further role and is subject to no further extrasystematic constraints.

So the presence of outlandish parameters of the index (taste, information state, etc.) does not itself make a semantic theory “relativist” in any philosophically interesting sense. It does not prevent the theory from assigning truth values

to sentences at contexts, or making absolute judgements of the accuracy of assertions.

We can see this more clearly if we change our target notion to make room for the possibility that accuracy is not absolute. Instead of defining truth at a context of use, let us define truth at a context of use and *context of assessment*—a possible situation in which a speech act or mental attitude might be assessed. We can then distinguish two kinds of contextual sensitivity. An expression is *use-sensitive* if its extension depends on features of the context of use, and *assessment-sensitive* if its extension depends on features of the context of assessment. Or, making finer discriminations: an expression is *F-use-sensitive* if its extension depends on the *F* of the context of use, and *F-assessment-sensitive* if its extension depends on the *F* of the context of assessment. A “relativist” semantic theory is one that takes some expressions to be assessment-sensitive. This amounts to taking accuracy to be relative to contexts of assessment, since

Accuracy An assertion of *S* at *c* is *accurate*, as assessed from *c'*, iff *S* is true as used at *c* and assessed from *c'*.

A compositional semantics that employs a taste parameter in the indices (perhaps to make sense of an operator like ‘on any standard of taste’) need not take any expression to be taste-assessment-sensitive. For such a semantics can be combined with either a contextualist or a relativist postsemantics:

Contextualist postsemantics *S* is true as used at *c* (and assessed from *c'*) iff *S* is true at *c*, $\langle w_c, t_c \rangle$, where w_c is the world of *c* and t_c is the taste of the agent at *c*.

Relativist postsemantics *S* is true as used at *c* and assessed from *c'* iff *S* is true at *c*, $\langle w_c, t_{c'} \rangle$, where w_c is the world of *c* and $t_{c'}$ is the taste of the agent at *c'* (i.e., the assessor).

According to the contextualist theory, the accuracy of an assertion of ‘This is tasty’ will depend on the physical disposition of the demonstrated food at the time of the assertion and on the speaker’s tastes at the time of the assertion. According to the relativist theory, by contrast, it will depend on the physical disposition of the demonstrated food at the time of the assertion and the *assessor’s* tastes at the time of assessment. Since the same assertion can be assessed from indefinitely many perspectives, there are only perspective-relative answers to the question whether the assertion was accurate.

Content and circumstance

Lewis (1980) argued that, for purposes of doing semantics, one need not talk about contents or propositions at all. In that respect he differs from Kaplan

(1989), who adopts a two-stage picture of semantics. In the first stage of such a semantics, we define the *expression* relation, specifying for any sentence S and context c the *structured proposition* that S expresses at c . In the second stage, we give a recursive truth definition of truth at a *circumstance of evaluation* for the structured propositions. In Kaplan’s framework, circumstances are time/world pairs, but Kaplan acknowledges that “other features,” such as location, may also be needed (1989, 504). Though it is not clear that Kaplan would go so far, one could in principle take subjective features like *taste* or *body of information* to be parameters of circumstances.

The truth of a sentence relative to a context is the joint product of these two stages:

Kaplanian postsemantics A sentence S is true at a context c iff the content expressed by S at c is true at the circumstance of c , that is, $\langle w_c, t_c \rangle$. (cf. Kaplan 1989, 522)

Here, the context plays two distinct roles in determining a truth value for an occurrence of a sentence. It plays a *content-determinative* role in fixing which content is expressed, and a *circumstance-determinative* role in fixing at which circumstance this content is to be evaluated. Hence, there will in general be two distinct ways to take some expression to be F -use-sensitive. One can take the expression’s *content* to depend on the F of the context of use, or one can take the F of the context of use to play a role in picking out the “circumstance of the context”—the circumstance one looks at in determining whether the sentence is true at the context, and in determining whether an assertion or belief in the context is accurate.

Contextualist views typically use the first model. Contextualists about ‘tasty’, for example, often hold that the sentence ‘vegemite is tasty’ expresses different propositions depending on the tastes of the speaker. But the second model is also an option. One could take the content of ‘vegemite is tasty’ to have truth values relative to worlds, times, and *tastes*, and let the tastes of the speaker (or more generally the taste relevant at the context of use) help determine which triple is the “circumstance of the context.” On such a view, which following MacFarlane (2009) we can call *nonindexical contextualism*, sentences have truth values relative to contexts of use, and accuracy is absolute, not assessment-relative. An assertion that vegemite is tasty is accurate, simpliciter, just in case the taste of vegemite is good according to the taste of the speaker (or whatever taste is relevant at the context of use). On this view, ‘vegemite is tasty’ is taste-use-sensitive, not taste-assessment-sensitive.

Some writers in the literature use the term “truth relativism” for any view that relativizes propositional truth to “nonstandard” features like tastes or information states. However, as the possibility of nonindexical contextualist views shows, such relativism at the level of circumstances of evaluation is compatible with the traditional view that the accuracy of particular assertions and beliefs

is an absolute matter. (Conversely, one can have assessment sensitivity without any nonstandard parameters of circumstances. For an example in which circumstances are just worlds, see MacFarlane 2008.) Since most of the traditional worries about truth relativism are worries about the rejection of this traditional view, and not about the individuation of propositions, we follow MacFarlane (2005a) in reserving the label “truth relativism” for views that countenance assessment sensitivity. (Lasersohn 2005 does not explicitly relativize truth to contexts of assessment, so it might appear that his view is not relativist, in the sense distinguished here. However, that would be misleading. For Lasersohn, a “context of use” is not a concrete possible situation in which a sentence might be used, but an abstract sequence of parameters. How we should set these parameters in interpreting an utterance of ‘This is tasty’ depends, Lasersohn holds, on both features of the concrete speech situation, which help determine the reference of ‘this’ and the relevant world and time, and on features of the concrete assessment situation, which determine the “judge” of the context. So both the concrete use situation and the concrete assessment situation play a role, even though the distinction between them is not made notationally salient.)

If we want a view on which ‘vegemite is tasty’ is taste-assessment-sensitive, then again there are two options in a Kaplanian framework. If we give the context of assessment a content-determinative role, then we have a form of *content relativism*.

Content-relativist postsemantics S is true as used at c and assessed from c' iff the content expressed by S at c , as assessed from c' , is true at the circumstance of c , that is, $\langle w_c, t_c \rangle$, where w_c and t_c are the world and taste of c .

On a content-relativist view, there is no absolute answer to the question which proposition is expressed by a sentence at a context of use. Different propositions are expressed, relative to different contexts of assessment. (Content relativism has had few advocates in the literature, as compared with truth relativism. But see Cappelen 2008 and Weatherston 2009.)

If we give the context of assessment a circumstance-determinative role, then we get a

Truth-relativist postsemantics S is true as used at c and assessed from c' iff the content expressed by S at c is true at the circumstance $\langle w_c, t_{c'} \rangle$, where w_c is the world of c and $t_{c'}$ is the taste of c' .

Here the context of use and context of assessment jointly determine which circumstance is relevant to the truth of the sentence relative to these contexts. According to this view, an assertion or belief that something is tasty will be accurate, as assessed from another context c' , if the way the thing tastes at the time of the assertion or belief is good according to the tastes of the assessor at

c' . So the very same assertion or belief may be accurate, as assessed by one observer, and inaccurate, as assessed by another. This is one way of trying to capture the idea that tastiness depends on our evaluative reactions, without reducing all ascriptions of tastiness to claims about how foods strike some particular person's or group.

Variables

A relativist who adopts a Lewisian or Kaplanian approach can keep the idea that an adjective like 'tasty' is semantically a one-place predicate. There is no extra argument place for a taste or judge. This may be counted an advantage, but it makes accounting for binding phenomena more awkward. We must analyze

- (1) That's tasty to Sam

as

- (2) To Sam: tasty(that),

where 'To Sam:' is an operator that shifts the taste parameter of the index (or, in Lasersohn's 2005 version, a predicate-modifier). And we must analyze

- (3) Everyone got something tasty

(on the reading where it means that everyone got something tasty to them) as

- (4) [Everybody x][Something y](x got y & to x : tasty(y)).

A tempting alternative is to treat 'tasty' as semantically *two-place*, with an implicit argument place for a taste. We could then analyze (1) as

- (5) Tasty(that, Sam's taste)

and (3) as

- (6) [Everybody x][Something y](x got y & tasty(y , x 's taste))

Following Stephenson (2007), we could take the extra argument place to be filled, when not bound, by a special unpronounced pronomial element, PRO_J . On Stephenson's account, PRO_J denotes the judge of the index; here, we could take it to denote the taste of the index, or (if we had no other need for a taste

parameter in the index) we could take its value to be set directly by the taste of the context of assessment.

Thus, although it is sometimes assumed that relativist views require operators to motivate novel parameters in the index (Cappelen and Hawthorne 2009, ch. 3; Stanley 2005, ch. 7), the issue of assessment sensitivity can be pried apart from questions about the parameters of indices, and from related questions about whether there are operators that shift these parameters. This should not be surprising. As we will see below, assessment sensitivity is a way to explain certain features of the use of non-embedded occurrences of sentences. Thus it ought to be largely independent of the issues about *embedded* occurrences of expressions that motivate indices.

THE PHILOSOPHICAL QUESTION

What motivates work on assessment sensitivity is an

Empirical claim There are assessment-sensitive expressions in natural languages. Some of the things we say and think are assessment-sensitive.

Defending this claim would require detailed examination of the use of particular bits of language and careful consideration of alternative explanations. However, the empirical claim presupposes a weaker, conceptual claim that can be established—if it can be established at all—through philosophical clarification:

Conceptual claim The possibility of assessment-sensitive expressions is coherent and intelligible. We know what it would be for an expression to be assessment-sensitive, whether or not this possibility is realized in natural languages.

The ground-clearing of the last section has helped us see what assessment sensitivity is not, but it has not yet told us what it is. If we are given a definition of *truth at a context of use* for a language, then we have a reasonable grip on how its expressions are to be used, since we know that (to a first approximation) speakers try to assert things that are true at their current contexts, and expect others to do the same (see Dummett 1981, 302; Lewis 1980, §2; 1983, §III). But what if we are given a definition of *truth at a context of use and context of assessment*? How does this notion connect to the use of language? What is the significance of assessment sensitivity for our linguistic practices? Something must be said here, or it is not clear that in talking of truth relative to a context of use and context of assessment, we are saying anything meaningful at all.

It is important to realize that both ‘true as used at a context *c*’ and ‘true as used at *c* and assessed from *c*’ are theoretical terms. Ordinary speakers don’t

talk of sentences being true relative to contexts. When they say things like, ‘What Joe said is true’ and ‘None of the claims in Sally’s report are true’, they are using a monadic predicate that applies to propositions. So, we need to say something about what these relativized truth predicates mean if we are to use them in theories of meaning. It might seem tempting to define ‘is true as used at c ’ as ‘would express a truth if it were used at c ’, but this doesn’t quite work: it would render the sentence ‘An English sentence is being used’ true at all contexts, for example. In addition, if the ordinary monadic truth predicate is nothing more than a logical device for expressing generalizations, as many philosophers have claimed, then a definition that fixes its extension cannot serve as a theory of meaning (see Dummett 1959 for the point, Williams 1999 for dissent, and Patterson 2005 for a recent defense). A more promising approach is to let the meanings of these technical terms be fixed (as with many technical terms) by their role in a larger theory—here, a theory of language use. If we understand what is implied about the use of language by claims made using the relativized truth predicate, then we understand this predicate well enough.

In the case of ‘true as used at c ’, we can get pretty far by saying, with Dummett and Lewis, that speakers understand each other (in the absence of special reasons) to be trying to assert only truths; that is, they follow a norm

T1 Assert ϕ in c only if ϕ is true as used at c .

Given that knowledge, plus a definition of truth at a context of use for a language, you will be in a position to understand utterances others make, and to communicate your own thoughts to others using the language. Simplifying somewhat: when you want to communicate that you are hungry, you will find a sentence whose truth in your current context depends on your being hungry; when someone else utters a sentence that your definition tells you is true in their context just in case they are hungry, you will take them to be communicating that they are hungry.

But suppose we have a doubly relativized truth predicate: ‘true as used at c and assessed from c' ’? What shall we say about how ascriptions of assessment-relative truth relate to language use?

We might try quantifying over contexts of assessment in our norm for assertion:

T2 Assert ϕ in c only if ϕ is true as used at c and assessed from c' , for some/most/all contexts of assessment c' .

Or, more plausibly, we might simply privilege a particular context of assessment—the one the speaker is in when making the assertion:

T3 Assert ϕ in c only if ϕ is true as used at c and assessed from c .

But if this is *all* we say, assessment sensitivity will be an idle wheel, with no real use in semantic theorizing. For, given any definition T of ‘true as used at c and assessed from c' ’ for a given language, according to which certain expressions are assessment-sensitive, we can construct another definition T^* that yields the same norms for language use as T does (when combined with T3), but does not make any expressions assessment-sensitive. Here’s the recipe:

T* ϕ is true as used at c and assessed from c' iff according to T , ϕ is true as used at c and assessed from c .

T3 cannot help us distinguish these two definitions, because T3 only tells us the significance of the assessment-relative truth predicate for the special case where the context of assessment is the same as the context of use, and in such cases T and T^* will always agree. If T3 is all we have to say to connect the assessment-relative truth predicate to language use, then, we will never have evidence that favors T over T^* , so we will never have good reason to embrace assessment sensitivity (see MacFarlane 2005a, which draws on Evans 1985, 348).

However, we can distinguish theories that posit assessment sensitivity from theories that do not if we consider not just norms for making assertions, but norms for retracting, challenging, and/or justifying them. For, when one is considering a challenge to an assertion, or considering whether to retract an assertion made earlier, there are potentially two relevant contexts: the context in which the assertion was made and the context in which it is now being assessed. There is room here, then, for the context of assessment to play a significant role in the norms for these speech acts.

For example, suppose we combine T3 with the following norm for retraction:

R Retract at c' an assertion of ϕ made at c if ϕ is not true as used at c and assessed from c' .

Now we can distinguish between relativist and nonindexical contextualist views about ‘tasty’. Both will say that one may assert ‘this is tasty’ at c only if the demonstrated food is good according to one’s tastes at c . But, given R, the two views will differ about the condition under which such an assertion must be retracted at some later context c' where one’s tastes are different. According to the nonindexical contextualist version, ‘this is tasty’ is not assessment-sensitive; so it is true as used at c and assessed from c' just in case it is true as used at c and assessed from c . If it was okay to make the assertion in the first place, then, it need not be retracted. According to the relativist view, however, the truth of ‘this is tasty’ (at a context of use and context of assessment) depends on the taste of the context of assessment; so, even if ‘this is tasty’ is true as used at c and assessed from c , it may be false as used at c and assessed from c' . In such a case, the assertion would have to be retracted, even though one did not violate any norm in making it.

Consider a concrete example on which the relativist and nonindexical contextualist views differ. Suppose that at age ten, I asserted ‘fish sticks are tasty’. Now my tastes have changed; I find fish sticks disgusting, and I deny that fish sticks are tasty. Am I obligated to retract my earlier assertion? The nonindexical contextualist says that I am not. I can let this earlier assertion stand, since, although it was not specifically *about* my tastes, it concerned my tastes at the time. The relativist says that I must retract; I cannot coherently reject the proposition I asserted as false while holding fast to my earlier assertion.

Of course, claims about what speakers are obligated to do are not themselves claims about what speakers *will* do. Still, this is enough to give us an empirical grip on the distinction. After all, speakers can be presumed to be trying to satisfy the norms, except when there is good reason not to, and they may have intuitions about their obligations to retract. If assessment-sensitive and non-assessment-sensitive theories make different predictions about proprieties for language use, they will also make different predictions about language use. Arguably, this is enough to vindicate the conceptual claim, and to tell us what to look for in evaluating the empirical claim. (See MacFarlane 2003, 2005a, 2007 for the general strategy pursued here.)

REFERENCES

- Cappelen, H. (2008) “Content Relativism and Semantic Blindness,” in García-Carpintero and Kölbel (2008), 265–286.
- Cappelen, H. and J. Hawthorne (2009) *Relativism and Monadic Truth* (Oxford: Oxford University Press).
- Dietz, R. (2008) “Epistemic Modals and Correct Disagreement,” in García-Carpintero and Kölbel (2008), 239–263.
- Dummett, M. (1959) “Truth,” *Proceedings of the Aristotelian Society* n.s. 59: 141–62.
- Dummett, M. (1981) *Frege: Philosophy of Language*, 2nd ed. (Cambridge: Harvard University Press).
- Egan, A. (2007) “Epistemic modals, relativism, and assertion,” *Philosophical Studies*, 133: 1–22.
- Egan, A., J. Hawthorne, and B. Weatherson (2005) “Epistemic Modals in Context,” in G. Preyer and P. Peter (eds.), *Contextualism in Philosophy* (Oxford: Oxford University Press).
- García-Carpintero, M. and M. Kölbel (eds.) (2008), *Relative Truth*, Oxford: Oxford University Press.
- Glanzberg, M. (2007) “Context, Content, and Relativism,” *Philosophical Studies* 136: 1–29.

- Heck, R. (2006) “MacFarlane on Relative Truth,” *Philosophical Issues* 16: 88–100.
- Kamp, H. (1971) “Formal Properties of ‘Now’,” *Theoria* 37, 227–273.
- Kaplan, D. (1989) “Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals,” in J. Almog, J. Perry and H. Wettstein (eds.), *Themes from Kaplan*, Oxford: Oxford University Press, 481–566.
- Kölbel, M. (2002) *Truth Without Objectivity* (London: Routledge).
- Lasersohn, P. (2005) “Context dependence, disagreement, and predicates of personal taste,” *Linguistics and Philosophy* 28: 643–86.
- Lasersohn, P. (2009) “Relative Truth, Speaker Commitment, and Control of Implicit Arguments,” *Synthese* 166: 359–74.
- Lewis, D. (1980) “Index, context, and content,” in Stig Kanger and Sven Öhman (eds.), *Philosophy and Grammar*, Dordrecht: Reidel, 1980. Reprinted in Lewis (1983).
- Lewis, D. (1983) *Philosophical Papers*, vol. 1, Oxford: Oxford University Press, 1983.
- MacFarlane, J. (2003) “Future Contingents and Relative Truth,” *Philosophical Quarterly* 53: 321–36.
- MacFarlane, J. (2005a), “Making Sense of Relative Truth,” *Proceedings of the Aristotelian Society* 105: 321–339.
- MacFarlane, J. (2007) “Relativism and Disagreement,” *Philosophical Studies* 132: 17–31.
- MacFarlane, J. (2008), “Truth in the Garden of Forking Paths,” in Kölbel and Manuel García-Carpintero (eds.) 2008, 81–102.
- MacFarlane, J. (2009) “Nonindexical Contextualism,” *Synthese* 166: 231–50.
- MacFarlane, J. (2011) “Epistemic Modals are Assessment Sensitive.” in E. Egan and B. Weatherson (eds.), *Epistemic Modality*, Oxford: Oxford University Press.
- Patterson, D. (2005) “Deflationism and the Truth Conditional Theory of Meaning,” *Philosophical Studies* 124: 271–94.
- Stanley, J. (2005) *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Stephenson, Tamina (2007) “Judge dependence, epistemic modals, and predicates of personal taste,” *Linguistics and Philosophy* 30, 487–525.
- von Fintel, K. and A. Gillies (2008) “CIA Leaks,” *Philosophical Review* 117, 77–98.

Weatherson, B. (2009) “Conditionals and Indexical Relativism,” *Synthese* 166: 333–357.

Williams, M. (1999) “Meaning and Deflationary Truth,” *Journal of Philosophy* 96: 545–64.

Wright, C. (2007) “New Age Relativism and Epistemic Possibility: the Question of Evidence,” *Philosophical Issues* 17, 262–283.

FURTHER READING

MacFarlane, J. (2005b) “The Assessment Sensitivity of Knowledge Attributions,” in T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* 1: 197–233. (Argues that a relativist treatment of knowledge attributions combines the virtues of contextualism and invariantism.)

Recanati, F. (2007) *Perspectival Thought: A Plea for (Moderate) Relativism*. Oxford University Press. (An extended argument for what is here called non-indexical contextualism, with applications.)

Richard, M. (2004) “Contextualism and Relativism,” *Philosophical Studies* 119: 215–242. (Argues for a relativist treatment of vague terms and knowledge attributions.)

Richard, M. (2008) *When Truth Gives Out* (Oxford: Oxford University Press). (Argues for a relativist treatment of vague terms, and against a relativist treatment of taste predicates.)

Zimmerman, A. (2007) “Against Relativism,” *Philosophical Studies* 133: 313–348. (Argues against a variety of relativist proposals.)

BIOGRAPHICAL NOTE

John MacFarlane is Professor of Philosophy at the University of California, Berkeley. He has written extensively on the philosophy of logic and the philosophy of language, and is currently working on a book on relative truth and its applications.