

Assessing Epistemic Modals

Tristan Haze
The University of Sydney

Draft of 30 November, 2011.

1. Introduction

Contemporary work on epistemic modals is dominated by a rivalry between two kinds of theory: contextualism and relativism. A common feature of these theories is that they are put in terms of the *truth or falsity* of epistemic modals: contextualism says, roughly, that the truth of an epistemic modal depends on epistemic features of its context of utterance. Relativism says, roughly, that the truth of an epistemic modal depends on epistemic features of its context of *assessment*.

Contextualism quite clearly makes sense, but struggles with basic data concerning our use of epistemic modals. Relativism, at least superficially, seems to handle at least some of the data better than contextualism, but doesn't obviously make sense.

In this paper, I want to take a different approach. The basic idea of this approach is that there are (at least) two different ways an epistemic modal judgement can be good or bad, successful or unsuccessful, right or wrong. Clearly distinguishing these *dimensions of assessment* of epistemic modal claims should be a good strategy for investigating their meaning and use, and addressing the puzzlement to which they give rise. When doing this, we shall hold the basic notions of truth and falsity at bay - at least, in connection with epistemic modals themselves - in the interest of clarity, fashioning more specific notions in response to the data. In these terms, we can give an account of the basic norm of assertion governing epistemic modals. That being done, we are free to bring the basic notions of truth and falsity back into the discussion, using the results of the preceding investigation as a solid basis for considering where they might fit.

I will begin with some critical remarks on contextualism and relativism in sections 2 and 3, in order to provide some background and motivation for my preferred approach. (I won't attempt to conclusively refute these theories - especially not relativism - but ultimately hope to make it seem plausible that the dimensions of assessment approach is superior.)

In section 4 I distinguish and characterize two dimensions of assessment: robustness and ultimate correctness. In section 5 I give a minimal norm of assertion for epistemic modals in terms of

robustness. In section 6 I discuss the application of 'true', 'false', 'right', 'wrong' etc. in connection with epistemic modals. In section 7 I show how the dimensions of assessment approach handles the salient data, and I conclude in section 8 with a summary of the respective costs and benefits of contextualism, relativism and the dimensions of assessment approach.

Two terminological preliminaries: (i) I will use ' $\langle \rangle p$ ' as a schema for epistemic modal claims of the forms 'It might be that p ', 'It could be that p ', 'Possibly, p ', etc. (ii) I will use the revived mediaeval term 'prejacent' to refer to the proposition which is operated on by a sentential operator. So the prejacent of ' $\langle \rangle p$ ' is ' p '.

2. Contextualism

Contextualism about epistemic modals says that the content and truth of an epistemic modal statement depends on its context of utterance. It can be divided into two kinds. Following MacFarlane's terminology, I will call these 'solipsistic contextualism' and 'non-solipsistic contextualism'.

2.1. Solipsistic contextualism

This simplest form of contextualism is the proposal that an epistemic modal ' $\langle \rangle p$ ' is true in a context of utterance iff ' p ' is not ruled out what is known by the utterer. 'Not ruled out' may be cashed out further, and in various ways, but for our purposes we can abstract away from whether or how that is done.

Solipsistic contextualism is able to explain our seeming ability to assert epistemic modals *on the basis of ignorance* in some sense. However, it seems unable to make sense of the fact that we often retract epistemic modals, as in the following conversation (from MacFarlane 2009, p. 5):

Sally: "Joe might be in Boston."

George: "No, he can't be in Boston. I just saw him an hour ago in Berkeley."

Sally: "Okay, then, scratch that. I was wrong."

Relatedly, when we make third-person assessments of epistemic modals, we do not seem to be assessing whether the prejacent is ruled out by the utterer's knowledge - rather we seem to be

assessing the prejacent itself, by our best lights, to see if *we* should not rule it out. Furthermore, it seems we often disagree about epistemic modals, and as MacFarlane says:

It is crucial to such disputes that the participants take themselves to be contradicting each other when one says “It might be that *p*” and the other says “No, it can’t be that *p*.” Solipsistic Contextualism cannot make sense of this. For it holds that the first participant’s claim is about what *she* knows, while the second’s is about what *he* knows. (MacFarlane 2009, p. 6)

MacFarlane presents these as three problems: the problems of retraction, third-person assessments and disputes. But they are pretty clearly all manifestations of a single root problem: solipsistic contextualism makes false predictions about the assessment of epistemic modals. The problem of retraction is the first-personal manifestation of this, the problem of dispute is the second-personal manifestation, and that of third-person assessments is the third.

MacFarlane considers the defence that this data can be explained by speakers' failure to understand the context-sensitivity of epistemic modals, and makes the point that this is 'a lot of error to impute to speakers', and that such error is not explained by solipsistic contextualism as it stands. MacFarlane concludes that the defence should be put on the backburner, 'as a last resort should no alternative view prove viable'. I would go further, and suggest that we reject this response out-of-hand, and solipsistic contextualism with it. (Cognitive resources are scarce, after all, and if solipsistic contextualism *is* false, the belief that it is false will be a valuable constraint on our inquiry, more valuable than a merely provisional supposition that it is false.)

2.2. Non-solipsistic contextualism

Contextualisms relativize epistemic modal statements to a knowledge-base. Non-solipsistic contextualism differs from solipsistic contextualism in widening the knowledge-base to include more than just the knowledge of the utterer. There are various ways of doing this - one can add other people's actual knowledge to the knowledge-base, or knowledge which would come to light given investigation, or make both of these moves.

The attractive thing about non-solipsistic contextualism is that it is able to explain negative assessments: retractions, disagreements, and third-person rejections. Recall the conversation

which caused problems for solipsistic contextualism above:

Sally: "Joe might be in Boston."

George: "No, he can't be in Boston. I just saw him an hour ago in Berkeley."

Sally: "Okay, then, scratch that. I was wrong."

Consider the view that Sally's first statement was true iff 'Joe is in Boston' is not ruled out by Sally *and* George's - the conversants' - knowledge. This view can easily explain Sally's retraction: she comes to know that George knows something that rules out Joe's being in Boston, and therefore that the truth-conditions of her statement are not fulfilled. It also seems to make reasonable sense of Sally's making her claim in the first place: we can imagine that Sally had reason to think that George would agree with her statement - that she didn't say it purely on the basis of her own ignorance.

(This may not be plausible for other examples, however - which should worry the non-solipsistic contextualist. We will not press this point now, but instead go for a more general criticism. The issue just identified, that we arguably *do* sometimes assert epistemic modals purely on the basis of our own ignorance, will come out as a limiting case of non-solipsistic contextualism's problems with securing conversational warrant for would-be asserters of epistemic modals.)

To begin to see the problem with this sort of view, consider the fact that Sally's corrector, it seems, needn't have been George, but could have joined the conversation upon hearing Sally's initial claim. For cleanness, let us imagine a case where that the corrector is not known to Sally, and not even in the same room listening sustainedly to the conversation: let us suppose Sally said 'Bill Clinton might be in Boston' to George, mid-conversation, on a busy city street, catching the ear of a passer-by who impulsively turns to Sally and says 'No, he can't be in Boston. I just saw him an hour ago in Berkeley'. Supposing that Sally is (perhaps unusually) friendly and open to strangers joining her conversations, there doesn't seem to be any reason why she shouldn't retract her claim, just as she did with George in the earlier example.

This suggests that non-solipsistic contextualism, in order to be plausible, has to allow for this sort of completely unforeseen conversation-joining. But if the truth of Sally's claim required that the prejacent not be ruled out by the knowledge of anyone who might join the conversation, she should have thought twice about uttering it in a crowded place like that. Since that seems wrong,

we should, by modus tollens, be suspicious of the truth-conditions in question.

The non-solipsistic contextualist could perhaps bite the bullet on this consequence without their view seeming *totally* ludicrous. But the problem cases get worse: their view has an even worse time - if this is possible! - with third-person assessments than it does with retraction and disagreement.

The question of who might join a conversation seems like a bit of a red herring, in that bullet-biting "fixes" which involve such considerations will not generalize to problem-cases involving remote third-person assessments. An eavesdropper could, it seems, correctly judge that Sally's initial claim is wrong, without being anywhere near her, and without doing anything about it. Imagine Sally is on a reality TV show, or being spied on, or imagine someone overhearing her statement and later reporting it to someone who assesses it years later on the other side of the world. And what about forms of theism on which God is always listening? Should such theists refrain from ever uttering nontrivial (i.e. unknown-prejacent) epistemic modals? (And couldn't such a God join the conversation, for that matter?) As MacFarlane vividly puts it,

[T]here is no way to stop this machine The same kind of arguments that motivate expanding the relevant group of knowers to include George (in our example above) will motivate expanding the relevant group of knowers to include anybody who will ever consider the claim. (MacFarlane 2009, p. 9)

This is a serious problem for non-solipsistic contextualism, because if the relevant knowledge-base is expanded to include 'anybody who will ever consider the claim', the majority of epistemic modal assertions, including Susan's, would seem completely unwarranted. (The view that we go around making such unwarranted claims all the time is another one I would recommend we dismiss out of hand for the sake of getting something done.) Thus the 'widening-the-knowledge-base' approach to saving contextualism seems impaled on a dilemma: if it's set up to get warrant right, it makes false predictions about retractions, disagreements and third-person rejections, and if it's set up to get these things right, it makes false predictions about warrant. (The avoidance of this problem is the chief motivation for MacFarlane's relativism, which we are about to consider.)

There are other, more sophisticated forms of contextualism on the market, but we will not discuss any here. The most distinctive one I know of is the cloudy contextualism of von Fintel and Gillies

2008, criticized in a talk by MacFarlane entitled 'Relativism vs. Cloudy Contextualism' (available online as of writing).

3. Relativism

Relativism is the claim that epistemic modals only have truth-values relative to a context of assessment. If that sounds like nonsense, don't worry - it may well be. The claim is modelled on the claim that epistemic modals - or more generally, some class of statements - only have truth-values relative to a context of utterance. Statements involving indexicals, for example, are plausibly characterized in this context-of-utterance-relative way.

Rather than critically investigating the sense of MacFarlane's relativism straight away, let us try to understand it a bit better - *qua* phenomenon, if nothing else - by looking at what he tries to do with it. Then, when we come to criticism, we will have more ammunition: we can critique the sense, *as well as* raising problems which MacFarlane's account seems unable to "handle" - that is, cases which for which a treatment parallel to MacFarlane's ostensible handling of his favourite cases looks bad.

According to MacFarlane, relativism handles Sally's retraction in conversation with George as follows: when Sally first uttered her claim, it was true relative to her context of assessment. But then George increased her knowledge, and so came to judge that claim as false, relative to her new context of assessment. Third-person assessments can be handled in a similar way: when a third-person assessor assigns a truth-value to an epistemic modal, it is assigned relative to *their* context of assessment.

According to MacFarlane, there is no problem of explaining how these claims are warranted, as there is for non-solipsistic contextualism, since what matters for being warranted in asserting an epistemic modal is that it is true relative to the asserter's context of assessment.

Disagreement is less straightforward, but MacFarlane has a paper where he tries to make sense of disagreement about assessment-sensitive matters in general, which in his view include judgements of taste as well as epistemic modals. Cf. MacFarlane 2007. This could well be a weak spot, but we will not concentrate on that here.

3.1. *Over-generation of retractions and past-tensed epistemic modals*

MacFarlane's account faces a kind of problem which is raised in different ways by Wright 2007 and von Fintel and Gillies 2008. Here is a case inspired by Wright. S has a ticket in a lottery and asserts: 'It is possible for me to win'. S doesn't win. Still, it seems they could reasonably say: 'I was right when I said it was possible for me to win'. But if we apply MacFarlane's theory, this seems like an error.

Furthermore, it would be bizarre for S to say: 'Oh, it turns out I was wrong: it's not possible that I won', or for that matter 'Oh, it turns out I was wrong: it wasn't possible for me to win'. The first seems like a (reversed) non-sequitur, while the second seems like a valid reversed argument, but with a false premise. But it seems that MacFarlane's theory would predict just some such thing - there seems to be no provision for anything else. Wright reports that MacFarlane responds to this sort of lottery case by challenging the assumption that the relevant modality is epistemic in this case. Wright makes no attempt to respond to this, but says he thinks it's unconvincing.

von Fintel and Gillies provide the following additional problem-case, for which the 'not epistemic modality' reply seems, if anything, less convincing:

The first fact we want to draw attention to is that—*pace* the CIA [MacFarlane's relativism]—not all *mights* are retracted or rejected in the face of new evidence. Speakers can quite often *resist* the invitation to retract even if they have become better informed. Billy is looking for her keys. Alex is trying to help.

- (10) a. Alex: The keys might be in the drawer.
b. Billy: (*Looks in the drawer, agitated.*) They're not. Why did you say that?
c. Alex: Look, I didn't say they were in the drawer. I said they *might be* there—and they might have been. Sheesh.

These sorts of cases raise interesting and very difficult issues about the interaction of tense, "mood" and modals. Hopefully, the account sketched in section 4 will shed *some* light on what is

going on in the cases above, but many questions will remain. For present purposes, all that matters is that MacFarlane's account gives no way of explaining these cases - or at least, no convincing way. (I do not know if MacFarlane would make the move of denying epistemicity here.)

3.2. Coherence

As we said before, the claim that epistemic modals have their truth-values relative to a context of assessment is modelled on the claim that they have their truth-values relative to a context of utterance. But it is far from clear that this is appropriate: many of the things we say about utterance-relative truth are such that nothing corresponds to them, or they become nonsense, when carried over to MacFarlane's notion of assessment-relative truth.

To begin with: it is very natural to say that claims which have their truth-values relative to a context of *utterance* also have their *content* relative to that context. When we assess context-of-utterance-relative assertions, we are assessing those contents.

It seems MacFarlane doesn't want to say the analogous thing with respect to contexts of assessment. If epistemic modals had their content only relative to a context of assessment, then different contexts of assessment would mean different contents, and so we could never be able to assess the *same* epistemic modal claims we made before, once our context changes - and then it seems there would be no way to account for retraction and the like.

Thus, things must be like this: epistemic modal claims have their content, in some sense, relative to their context of *utterance*. This content then has different truth-values when assessed in different contexts. Here the problem of coherence seems particularly sharp.

In addition to complaining about coherence here, it seems we can ask a cutting question of the form 'Why can't we X?', namely: If these claims already have their contents upon utterance, why do we *have to* assess them relative to our context? Why can't we also assess them relative to the context they were uttered in?

It seems like MacFarlane would have to admit that we *can* do this, by using the resources of his

theory, and saying 'Such and such a claim is true relative to the context of assessment which is identical to its context of utterance'? Don't we feel that we say *effectively* this quite often, and much more easily? Isn't that what is going on in the problem cases above, where asserters of epistemic modals can defend - "converting" them into the past tense - when the prejacent turns out false? MacFarlane's account, as it stands, provides for no such possibility.

4. Dimensions of assessment

Let us begin with the observation that epistemic modal judgements seem to be able to fail, or turn out wrong, in two different ways. Consider the following two cases:

(1) A speaker says 'John might be a Taoist', only to be immediately reminded of something they already knew, causing them to retract their statement.

(2) A speaker says 'John might be a Taoist', learning only much later, due to new developments, that John is not - couldn't be - a Taoist.

It seems quite natural to say something like: both of these judgements ultimately fail, in a sense, but the second one fails in a lesser sense.

There seems to be a difference in kind here, not just in degree. We can try to capture this with a contrast between being *wrong to think* (or assert) something and being *wrong in thinking* it. In the former case, it is natural to say that the speaker was wrong to think (and to assert) that John might be a Taoist. In the latter case, it is natural to say that the speaker was in a sense right to think and assert that John might be a Taoist, but ultimately wrong in thinking so.

We shall try to get clearer about this by distinguishing and characterizing two dimensions of assessment of epistemic modals, *robustness* and *ultimate correctness*. To anticipate: it will be natural to say that the speaker's assertion in (1) above fails of robustness as well as ultimate correctness, whereas the speaker's assertion in (2) might fail only with respect to ultimate correctness.

4.1. Robustness

Epistemic modal statements and questions may be thought of as coming with a contextually determined robustness parameter, indicating the kind and degree of inquiry through which the statement, or answer to the question, is supposed to be able to hold up. Often, context will make this quite definite, for instance when a particular course of inquiry is under discussion. Other times it may be very fuzzy.

We will say that an epistemic modal statement ' $\langle \rangle p$ ' *holds up* through a course of inquiry iff, during that course of inquiry, ' $\sim p$ ' does not come to be known.

Finally, we will say that an epistemic modal statement is *robust* iff it holds up through all possible courses of inquiry of the contextually specified kind and degree. (Here we are, roughly speaking, holding fixed everything except the course of inquiry, *especially* the stuff being inquired into.)

4.2. *Ultimate correctness*

Robustness is a matter of an epistemic modal holding up in a contextually given way, to a contextually given degree. The idea of ultimate correctness is that an ultimately correct epistemic modal is one which never gets "falsified" (I use shudder quotes because we are not officially using the notions of truth and falsity in connection with epistemic modals at this point). That said, an ultimately correct epistemic modal may *appear* be falsified, but if this is due to false beliefs, fallacious reasoning or the like, it doesn't count.

A natural first thought is to define ultimate correctness of an epistemic modal in terms of the truth of the prejacent proposition: ' $\langle \rangle p$ ' is ultimately correct iff ' p ' is true. But this definition will arguably undergenerate with respect to the intuitive idea of ultimate correctness indicated above: there may be false propositions ' p ' which cannot be known to be true or false, and for these ' $\langle \rangle p$ ' will always hold up, by definition, and thus should be classed as ultimately correct.

Let us therefore consider this definition: ' $\langle \rangle p$ ' is ultimately correct iff p is true or it is unknowable whether p .

We can simplify this to:

' $\langle \rangle p$ ' is ultimately correct iff it is not knowable that $\sim p$.

What does 'knowable' mean here? What is the relevant modality? I am not sure exactly what to say about this, but I don't mean anything like 'possible for a human to know using modes of inquiry which we can imagine'. I suggest something like: ' p ' is knowable iff ' p ' is known by some knower' cannot be ruled out *a priori*. Or: iff ' p ' is known by some knower' is metaphysically possible. I don't have anything to favour either of these proposals over the other, nor am I opposed in principle to a proposal involving a more restricted modality.

For those who suspect that the truth or falsity of *every* expressible proposition ' p ' is logically or metaphysically possibly known, 'knowable' (as in 'knowable that', not 'knowable whether') will be coextensive with 'true', in which case the definition will still be alright, but might as well be: ' $\langle \rangle p$ ' is ultimately correct iff ' p ' is true.

I will not try to develop this notion further here - I am proposing an approach to epistemic modals, rather than a worked-out theory - but I hope enough has been said about it to convey the basic idea, and to make it seem plausible that we are not without options for working it out in greater detail. (The puzzling thing in my mind is that there are *several* options.)

5. *The norm of assertion*

What I want here is a minimal notion constraining assertion to statements which it is reasonable to believe are correct in some sense. Ordinarily, a norm playing this role can be given in terms of truth: it is appropriate to assert a proposition ' p ' only if one has good reason to believe that it is true. But recall that, at this stage, we are for clarity's sake refraining from explicit employment of the notion of truth in connection with epistemic modals. I accordingly propose the following minimal norm of assertion:

It is appropriate to assert an epistemic modal only if one has good reason to believe that it is robust (where the relevant kind or degree of inquiry is given contextually).

This will, in the next section, be seen to comport with the simple truth-based norm, provided 'true' is given the appropriate one of two possible readings.

6. *Two applications of 'true', 'false', 'right', 'wrong' etc.*

We use terms like 'true', 'false', 'right', 'wrong', 'correct', 'incorrect', 'so', 'not so', etc. to report our assessments of claims. Since we have distinguished two dimensions of assessment of epistemic modals, it is natural to take the step of supposing that, with respect to epistemic modals, these terms are used with two different applications. In one application, we use these terms to attribute robustness and non-robustness to epistemic modals, in the other, we use them to attribute ultimate correctness and incorrectness.

I speak of two applications rather than two meanings. For a start, we clearly don't have an out-and-out case of ambiguity of the river-bank/money-bank sort. The relationship between the two applications is close: ultimate correctness amounts, after all, to robustness through all possible inquiry.

What should be said about 'the notion of truth', 'truth itself' etc. in light of this, I will not get into here on pain of opening an enormous can of worms. Some may prefer a quasi-realism-like approach, on which applications of the truth-predicate to epistemic modals (and, as a natural accompaniment, indicative conditionals) are somehow secondary, not part of the "downtown" of language, etc. My own inclinations are more Wittgensteinian: there are many kinds of propositions, 'proposition', 'language', etc. are open-textured concepts. It is time to stop this gesturing, however, for the can is opening.

(One further thing: with respect to the traditional conception of different theories of truth - correspondence, coherence, pragmatism, deflationism - I have no idea what to say. I think that whole conception is deeply problematic, but this is not the place to go into it.)

7. *Data explained*

I have given a schematic account of epistemic modals on which they have two dimensions of assessment: robustness and ultimate correctness. Furthermore, on this account, evaluative terms like 'true', 'false', 'right', 'wrong' can be used to signal either of these things (or their opposites), and it is appropriate to assert an epistemic modal only if one has good reason to believe it is robust. Let us now see how this account handles the data of warrant and rejection, the Wright-Fintel-Gillies 'could have'-data against relativism, and finally, the fact that we sometimes answer

'Yes' when asked an epistemic modal question such (e.g. a question of the form 'Could it be that p ?'), other times 'I don't know'.

7.1. Warrant and rejection

As we saw in section 4, epistemic modal claims can fail of robustness, or merely of ultimate correctness. Thus, contradictions, retractions and third-person negative assessments can be thought of as coming in (at least) two corresponding kinds.

Recall the case of Susan and George, and how it would be strange for Susan to respond to George's correction by standing by her assertion (as she would be able to if solipsistic contextualism were right). We can explain this in terms of robustness: clearly Susan's remark was supposed to be robust enough not to be ruled out immediately by George. So when she retracts her claim, saying 'I was wrong', this can be understood as saying not merely that she was not ultimately correct, but also that her claim wasn't robust.

In other cases, however, we will let ourselves be corrected, without wanting to say we were wrong in the robustness-sense. In that case, while we might say that the claim we did make *is* wrong (the present tense seems more natural here than the past), it would also be natural for us to stand by our assertion in a sense, saying that it *could have been* that p . (We will consider this sort of case immediately below in 7.2.)

Neither sorts of negative assessment generate any puzzles, on this view, about how we could be warranted in asserting epistemic modals: even in the case where my assertion fails of robustness, I might still have had good reason to think it would succeed - these things happen sometimes.

7.2. Could- and might-haves

Recall the 'keys might be in the drawer' case from 3.1. When they immediately turn out not to be in the drawer, we no longer think they might be in the drawer. In that sense, we think that we said has turned out to be false. Whereas we can say that the keys *might have been* in the drawer, and so in sense, what we said was true (the past tense seems much more natural here than the present).

The dimensions of assessment approach, on which 'true', 'false' etc. have a dual application,

captures this as follows: the first sense, in which what we said has turned out to be false, arises from 'false' being used to talk about ultimate correctness. The second sense, in which what we said was true, arises from 'true' being used to talk about robustness. (The claim that the keys might be in the drawer was, we might suppose, robust to the contextually understood degree: in this case, roughly, the requirement is that it hold up through pre-drawer-opening inquiry.)

7.3. 'Yes' vs. 'I don't know'

When someone asks us an epistemic modal question, i.e. whether something might be the case, we sometimes answer 'Yes', other times 'I don't know', often before explaining our answers. Solipsistic contextualism plainly cannot account for this. Non-solipsistic contextualism can: if we think the relevant context involves knowledge we don't have, we might not know whether that knowledge rules out the prejacent. It is not clear that relativism can make sense of this: when asked whether something could be the case, aren't I being asked to assess this question by myself? Thus, isn't the context of assessment my own epistemic situation? In that case, the only way we could make sense of the 'I don't know' answer would be to suppose that the answerer is in effect saying that they do not know *whether they know* whether the prejacent is true or false - a breach of the controversial KK principle. But there are cases where we give an 'I don't know' answer which are plainly not of this kind - where we *know* we do not know.

The dimensions of assessment approach can make sense of 'I don't know' answers to epistemic modal questions by means of the notion of robustness: when asked an epistemic modal question, context supplies the relevant robustness parameter. So, if someone asks 'Could it be that p ?', and I don't know whether p , two sorts of things (among others) may happen: (i) I may believe the questioner is asking for an answer which is more robust than I can give, and therefore say 'I don't know', or (ii) I may believe that my answer will hold up to the relevant extent or in the relevant way, and therefore say 'yes'.

8. Conclusion - costs and benefits

We have now seen three views about epistemic modals: we have looked critically at contextualism and relativism, and proposed a third alternative, the dimensions of assessment approach. This latter is still quite undeveloped compared to contextualism and relativism, but I suggest that what we have seen makes the following tentative conclusion plausible:

Contextualism is the most straightforward of the three approaches with respect to the basic ideas involved, but it seems to handle the data worst of all. Relativism is "ideologically" the most problematic of the three. It arguably handles the data better than contextualism (many cases of retraction and warrant can be made sense of, at least superficially), but not as well as the dimensions of assessment approach (relativism seems to predict inappropriate retractions). The dimensions of assessment approach, then, is less straightforward than contextualism, considerably more straightforward than relativism, and explains the data best of all.

References

- von Fintel, Kai & Gillies, Anthony S. (2008). Cia leaks. *Philosophical Review* 117 (1):77-98.
- John MacFarlane (2005). Making sense of relative truth. *Proceedings of the Aristotelian Society* 105 (3):321–339.
- John MacFarlane (2007). Relativism and disagreement. *Philosophical Studies*, 132:17–31.
- John MacFarlane (2009). Epistemic modals are assessment-sensitive. In Andy Egan & B. Weatherson (eds.), *Epistemic Modality*. Oxford University Press.
- Crispin Wright (2007). New Age Relativism and Epistemic Possibility: The Question of Evidence. *Philosophical Issues* 17 (1):262--283.