

## Attitudes and Relativism\*

Brian Weatherson, Rutgers/St Andrews

July 2008

Incomplete in many, many respects

Data about attitude reports provide some of the most interesting arguments for, and against, various theses of semantic relativism. This paper is a short survey of three such arguments. First, I'll argue (against recent work by von Fintel and Gillies) that relativists can explain the behaviour of relativistic terms in factive attitude reports. Second, I'll argue (contra Glanzberg) that looking at attitude reports suggests that relativists have a *more* plausible story to tell than contextualists about the division of labour between semantics and meta-semantics. Finally, I'll offer a new argument for invariantism (i.e. against relativism and contextualism) about moral terms. The argument will turn on the observation that the behaviour of normative terms in factive and non-factive attitude reports is quite unlike the behaviour of any other plausibly context-sensitive term. Before that, I'll start with some taxonomy, just so as it's clear what the intended conclusions below are supposed to be.

### 1. *How Not to be a Strawman*

Here are three mappings that we, as theorists about language, might be interested in.

- **Physical Movements → Speech Acts.** I'm currently moving my fingers across a keyboard. In doing so, I'm making various assertions. It's a hard question to say just how I manage to make an assertion by moving my fingers in this way. Relatedly, it's a hard question to say just which assertions, requests, questions, commands etc people make by making physical movements. This mapping is the answer to that question.
- **Speech Acts → Contents.** For some speech acts, their content is clear. If I assert that grass is green, the content of my assertion is that grass is green. For other speech acts, including perhaps other assertions, it is not immediately obvious what their content is. This mapping answers that question.
- **Contents → Truth Values.** Given a content, it isn't always clear whether it is true or false (or perhaps something else). This mapping provides the truth values of every content of every speech act.

The details of each of these mappings is interesting, to say the least. Indeed, a full description of those mappings would arguably include an answer to every question ever asked. We're not likely to know that any time soon. But we can ask, and perhaps answer, interesting questions about the

---

\* Thanks to Herman Cappellen, Andy Egan, Carrie Jenkies, Thony Gillies, Ishani Maitra and the participants at the Contextualism and Relativism workshop at Arché, St Andrews, May 2008.

topology of the mappings. Here are two distinct questions we can ask about each of the three mappings.

- **Is it one-one or one-many?** The mapping from a natural number to its natural predecessor is one-one. That's to say, every number is mapped to at most one other number. The mapping from a person to the children they have is one-many. That's not to say that everyone has many children, or indeed that everyone has any children. It's merely to say that some things in the domain are mapped to many things in the range.
- **Is the mapping absolute or relative, and if relative, relative to what?** Consider a mapping that takes a person A as input, and relative to any person B, outputs the first child that A has with B. This mapping is relative; there's no such thing as the output of the function for any particular input A. All there is, is the output relative to B1, relative to B2, and so on. But note that in a good sense the mapping is one-one. Relative to any B, A is mapped to at most one person.

A very simple model of context-sensitivity in language says that all three of these mappings are one-one and absolute. It will be helpful to have a character who accepts that, so imagine we have a person, called Strawman, who does. There are six basic ways to reject Strawman's views. For each of the three mappings, we can say that it is one-many or we can say that it is relative.

This way of thinking about Strawman's opponents gives us a nice taxonomy. The first mapping is about *speech acts*, the second about *contents* and the third about *truth*. Someone who disagrees with Strawman disagrees about one of these three mappings. They might disagree by being a *pluralist*, i.e. thinking that the mapping is one-many. Or they might disagree by being a *relativist*, i.e. thinking the mapping is relative to some thing external. The respective defenders of Strawman on these questions are the *monists* and the *absolutists*.

So the *speech act pluralist* is the person who thinks the first mapping is one-many. The *content relativist*, is the person who thinks the second is relative to some other variable (perhaps an assessor). And the *truth monist* is the person who thinks that contents have (at most) one truth value. All these terms are a little bit stipulative, but I think it the terminology here somewhat matches up with regular use. And it's the terminology I'll use throughout this paper.

One other nice advantage of this taxonomy is that it helps clarify just what is at issue between various opponents of Strawman. So Andy Egan (200x) and Josh Parsons (200x) have some data about uses of "you" in group settings that suggest such utterances pose a problem for Strawman. But it's one thing to have evidence that Strawman is wrong, another altogether to know which of his views is, on a particular occasion, wrong. I think separating out Strawman's various commitments helps clarify what is needed to isolate Strawman's mistake on an occasion.

It is, I think, more or less common ground that the first of Strawman's commitments, speech act monism, is false. The King can, by uttering "It's cold in here" both assert that it's cold in here, and command his lackey to close the window. Those look like two distinct speech acts

that he's made with the one physical movement. Herman Cappellen and Ernest Lepore (200x) have many more examples to show that Strawman is wrong here. Once we go beyond that though, it's less clear that Strawman is mistaken. Perhaps by thinking about cases where, by the one physical movement, we intend to communicate  $p$  to one audience member, and  $q$  to another, we can try to motivate speech-act relativism. That's an issue I'll leave for another day. In contrast to what he says about speech acts, what Strawman says about content and truth is, if not universally accepted, at least popular. So I'll call *orthodox contextualism* the view that Strawman is right about the content mapping and the truth mapping; they are one-one and absolute.

It is worthwhile noting two very separate models for content that lead to two quite distinct ways in which we can reject Strawman's last two absolutist views. John MacFarlane's paper on Non-Indexical Contextualism (200x) was extremely useful in setting up the relevant distinctions here, but the particular models for content I'm describing here are both set out in greatest detail in recent work by Andy Egan.

The first is the centred worlds model for content. This is the idea that for some utterance types, any token of that type expresses the same content. But that content is a set of centred worlds, that is true at some centres and false at other centres in the same world. So we might think that the content of "Beer is tasty" is, roughly, the set of possibilities who have pro-attitudes to the taste of beer. More precisely, it is the set of world-centre pairs such that the agent at (or perhaps closest to) the centre has pro-attitudes towards the taste of beer. On this view, content monism will be maintained – what an utterance of "Beer is tasty" says is invariant across assessors. (I'm assuming here that assessor-relativity is the only relativity we're interested in.) But truth absolutism will fail, since whether that content is true for  $a_1$  and  $a_2$  will depend on what their attitudes are towards beer. This kind of centred worlds model for content is what Egan has developed in (Egan 200x), (Egan 200x) and (Egan 200x).

The second model lets assessors get into the content-fixing mechanism, but says the content that is fixed is a familiar proposition whose truth is not assessor relative. This is easiest to explain with an example involving second-person pronouns. For some utterances of "You are a fool", the content of that utterance, relative to  $x$ , is that  $x$  is a fool. Now whether  $x$  is indeed a fool is a simple factual question, and whether it is true isn't assessor relative. But if some people are fools and others are not, whether the utterance is true or false depends on who is assessing it. So content relativism is true, while truth absolutism is preserved. This is a view Egan has defended for some tokens of second person pronouns (Egan 200x).

In "Conditionals and Indexical Relativism" (Weatherson 2008), I called the combination of content relativism and truth absolutism "indexical relativism", and defended such a view about indicative conditionals. I called something similar to the combination of truth relativism and content absolutism "non-indexical contextualism". More precisely, I followed MacFarlane in using that phrase for the combination of truth relativism and content absolutism and the view that whether a speaker's utterance is true (relative to an assessor) is a matter of whether the proposition they express is true relative to their context. I like the name "indexical relativism",

but it has also been used for theories that aren't even heterodox in the above sense, so perhaps persisting with it would just invite confusion. (And the name implies a particular view about how the relativity works; namely that there is something like an indexical element in what's asserted that gets its value from the context of assessment.) In other contexts I've used "relativism" as the label for all and only heterodox views, but this label is potentially quite confusing. Indeed, it's a possible worry about the arguments in "Conditionals and Indexical Relativism" that they really just support heterodoxy; a separate argument would be needed (and might not be easy to supply) against pluralist alternatives to content relativism about indicative conditionals.

## 2. *Factive Verbs and Relativism*

In "CIA Leaks" (von Fintel and Gillies 2007), Kai von Fintel and Thony Gillies raise a problem for heterodox theories about 'might'. (Actually they raise several, but I'm only going to deal with one of them here.) Their primary target is what I called truth relativist theories, but the argument they raise is interesting to consider from all heterodox perspectives. The problem concerns embedding of 'might'-clauses under factive attitude verbs. They argue as follows:

1. *S realizes that p* presupposes that *p*.
2. This presupposition is carried over when the sentence is used as the antecedent of a conditional. So, for instance, *If S realizes that p, then q* presupposes that *p*.
3. But, on standard heterodox proposals, we can properly say *If S realizes that it might be that p, then q*, even though it isn't true that it might be that *p*.
4. So heterodox proposals are false.

Here is the example they use to make the case.

Bond planted a bug and some misleading evidence pointing to his being in Zuurich and slipped out. Now he and Leiter are listening in from London. As they listen, Leiter is getting a bit worried: Blofeld hasn't yet found the misleading evidence that points to Bond's being in Zurich. Leiter turns to Bond and says:

(34) If Blofeld realizes you might be in Zurich, you can breathe easy—he'll send his henchman to Zurich to find you. [cite needed]

Now the problem is that for the heterodox theorist, "You might be in Zurich", as uttered by Leiter to Bond, expresses (relative to Bond), a proposition that is true iff for all Bond knows, Bond might be in Zurich. Just how it does this will differ for different heterodox theorists, but so far they all agree. But that isn't the case; since Bond knows he is in London. So (34) should sound defective, since it contains a presupposition failure. But it isn't defective, so heterodoxy is mistaken.

Before we look at how heterodox theorists might respond to this case, it's worth looking thinking about how Strawman might respond to it. The simplest idea is to say that in *It might be that p*, there is a hidden variable *X*. The value of *X* is set by context. And the sentence expresses the proposition that for all *X* knows, *p* is true. (Perhaps we might use some epistemic relation other than 'knows', but that's not relevant here.)

Now, and this is crucial, the variable *X* might be either free or bound. If there is nothing around to bind it, as in a simple utterance of *It might be that p*, then it will be free. And typically if it is free, *X* denotes a group consisting of the speaker and perhaps those in the same conversation. But when the might-clause is embedded under a propositional attitude ascription (factive or not), the variable *X* will be bound to the subject of the attitude ascription. So in *Y believes that it might be that p*, the value of *X* will simply be *Y*. So in *Blofeld realises you might be in Zurich*, the value of *X* is Blofeld. And hence the embedded might claim is true, since that claim is simply that for all Blofeld knows, Bond is in Zurich. Which, in the story, is true.

The reason for going through all of this is that the theorist who accepts truth absolutism, but rejects content absolutism, can say exactly the same thing. There is a variable *X* in the structure of what's asserted. Strawman thinks that you only get a determinate assertion when you fill in the value of *X*. We can disagree with that; we can say that an assertion can literally contain a variable, one that potentially gets its value from assessors. That way the content of a particular assertion can be different for different assessors. Once we've made this move, we can rejoin Strawman's story. This variable is either free or bound. If it is bound, we tell exactly the same story as Strawman. But, we insist, if it is free, the value of *X* is sometimes set by contextual features of the hearer as well as of the assessor. In the standard case, *X* is a group consisting of the speaker, the hearer, and perhaps some people who get in the group in virtue of their proximity to the speaker or hearer.

So we end up saying the same thing about the acceptability of (34) as Strawman. The content of *You might be in Zurich*, as embedded in (34), is quite different to the content those words would have if uttered as a standalone sentence, because the value that a key variable takes is different. For us, the value that variable takes differs for different assessors, but that's completely irrelevant to the explanation of the acceptability of (34).

For the truth relativist (who is also a content absolutist) things are a little more interesting. Such a theorist will typically reject the presence of a variable like *X* in the structure of what is said. So they cannot appeal to the kind of explanation that we've offered (twice over) for how (34) may be acceptable. The solution is to simply reject the generalisation about factive verbs. Let's start with some seemingly distant examples, in particular examples about fiction. It seems that (5) doesn't have any false presuppositions.

- (5) Watson realised that private detectives were (in late 19th Century London) better at solving murder mysteries than police.

I'll leave off the parenthetical in what follows. Now I take it that it simply isn't true that private detectives were better at solving murder mysteries than police. But it doesn't matter; this was true in the fiction and that's what is relevant. Note that neither (6) nor (7) has a false presupposition.

- (6) Had Watson realised earlier that private detectives were better at solving murder mysteries than police, he would have liked Holmes more than he did.
- (7) Had Watson realised earlier that private detectives were better at solving murder mysteries than police, the early chapters of the book would have been more interesting.

What's interesting about (7) is that it's clearly meant to be a claim about this world. When we say (6), it's naturally interpreted as making a claim about the world of the Holmes fiction. But that's not how we interpret (7); what matters is that the book would have been more interesting to us.

Note that this isn't anything particular to do with subjunctive conditionals. Imagine that we are settling down to watch a new adaptation of the Holmes stories that we are told won't be particularly faithful to the books in detail. I might properly say (8).

- (8) If Watson realises that private detectives were better at solving murder mysteries than police, the early scenes will be more interesting.

The lesson we take away from sentences like (8) is that the generalisation about factives and presupposition that von Stechow and Gillies rely upon isn't strictly true. When the subject of the attitude ascription is in another possible world, all that is presupposed is that the proposition they believe is true in their world. We should all agree to that restriction to the principle.

But now it is easy to see the way out of the argument for the proponent of the centred world view. The crucial thing about Watson isn't, such a theorist will say, that he's in another possible world. The crucial thing is that some propositions that are false relative to us (e.g. the proposition that private detectives were better at solving murder mysteries than police) are true relative to him. The true generalisation seems to be that  $S$  Vs that  $p$ , where  $V$  is factive, presupposes that  $p$  is true relative to  $S$ . And that's true in the cases that von Stechow and Gillies describe. So it's not true that the centred world theorist should predict that these utterances have false presuppositions. And that's all to the good, because of course they don't.

### 3. Glanzberg on Metasemantics

Papers arguing for relativism about some term  $t$  frequently, perhaps typically, start with a survey of reasons why (orthodox) contextualism about  $t$  cannot be correct. And such a survey will frequently include a sojourn through some quite specific contextualist theories, with some fairly obvious counterexamples. Egan et al (2005) sticks to the script as far as this goes. We note that  $a$

*might be F* can't, for instance, mean that *For all S knows, a is F*, where *S* is the speaker. And we note that some other simple theories along the same lines can't be true.

It's interesting to think through what kind of force such tours through the possible contextualist theories could have. We might think that there's a tacit argument that if some contextualist theory were true, it would be one of these simple ones, and none of the simple ones is true, so no contextualist theory is true. I'm not going to take a stand on exegetical debates here, so I'm not going to consider who may or may not have been making such an argument. I don't think that it was the intended argument in Egan et al (2005), but that's beside the point, because it is an argument that's now being debated. The argument under consideration isn't, or at least isn't directly, that the contextualist's *semantic* proposal is mistaken. Rather, the argument is that the accompanying *metasemantic* theory, i.e. the theory of how semantic values get fixed, is intolerably complicated. Slightly more formally, we can argue as follows.

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

The problem with this argument, as Michael Glanzberg (2007) has argued, is that premise 2 seems to be false. If we just look at some so-called 'automatic' indexicals, like "I" or "here" or "now", it just may be plausible. (But even in those cases things are tricky when we look at recordings, as in Weatherson 200x.) Once we widen our gaze though, we see that there are examples of uncontroversially context-sensitive terms, like "that", for which the accompanying metasemantic theory is, by any standard, hideously complicated. So the prospects of getting to relativism from metasemantic complexity are not, I think, promising.

That isn't the only kind of metasemantic argument for relativism though. A better argument for relativism turns on the fact that metasemantics *is* generally complicated. The contextualist, I think, has to make metasemantics too systematic at a key point. (Again, I'm not doing exegesis, but I do think something like this argument was intended in Egan et al (200x). I'm largely here highlighting something that I think has been thus far overlooked in the commentaries on that piece.) Consider the following pair of sentences.

- (9) Those guys are in trouble, but they don't know that they are.
- (10) ??Those guys are in trouble, but they might not be.

Something has gone wrong in (10). I conclude that (10) can't be used to express (9). That is, there's no good interpretation of (10) where those guys can be the denotation of *X* in the theory I attributed to Strawman in the previous last section. That's to say, the value of *X* can't just be the

most salient individual(s) in the context, since the guys are being made rather salient. A natural move at this stage is to adopt what some writers [cite?] have called the *Speaker Inclusion Constraint* (hereafter SIC). That is, in unembedded uses of “might” the group *X* always includes the speaker. Now the explanation of the problem with (10) is that for the speaker to assert the first clause, she must know that the guys are in trouble, but if that’s the case, and she’s in group *X*, then the second clause is false.

If the SIC really holds, it looks like it should hold in virtue of the meaning (in some sense of “meaning”) of “might”. As a rule, tidy generalisations like this should be part of semantics, not metasemantics. Compare two possible theories about “we”. Both theories say that “we” is a plural pronoun. One theory goes on to say that it is part of the meaning of “we” that it picks out a group that includes the speaker. That is, it puts this version of the SIC into the semantics. Another theory says that the SIC for “we” is just a nice metasemantic generalisation. I take it that the second position is very unattractive; it’s part of the meaning of “we” that it picks out a group including the speaker. And I think the relevant point generalises. At least it generalises as far as another SIC, namely the one that holds for “might”.

Semantic constraints on indexical terms hold, as a rule, for both embedded and unembedded uses of those indexicals. You can’t use “she”, even as a bound variable, to denote (relative to any variable assignment) a male human. There’s something badly wrong with *Every student thinks she will win*, if some students are female and others male. And you can’t interpret *O’Leary thinks that I am a fool*, as uttered by Daniels, to mean that O’Leary self-ascribes foolishness. In short, it seems that there are three related conclusions we can draw here. First, there are semantic constraints on the possible values of context-sensitive expressions. Second, any interesting generalisation about the possible value of a context-sensitive expression is traceable to such a semantic constraint. Third, these constraints remain in force when the expression is used in an attitude report, or as a bound variable.

The problem for contextualists about “might” is that it doesn’t behave at all this way. The SIC holds for unembedded uses. That implies that it is part of the meaning of “might”. So it should hold for embedded uses. But it does not. Indeed, for many embedded uses of “might”, a reading compatible with the SIC is simply unavailable. For instance, we can’t interpret *Every student thinks that they might have failed* as meaning that every student thinks that, for all I know, they failed. My knowledge just doesn’t matter; we’re talking about those students’ fears. This all suggests the following argument.

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

4. So, contextualism is not true.

Glanzberg argued, correctly, that it's no problem for the contextualist if, according to their theory, metasemantics was complicated and messy. It's not a problem because, well, metasemantics is complicated and messy. But this cuts both ways. And it is a problem for the contextualist that they have to put the SIC into metasemantics. It's just not messy enough to go there.

There are two objections to this argument (both of which were pressed when this paper was presented at Arché) that are worth considering together.

*Objection One: There are other generalisations that do go into metasemantics*

It's very odd, to say the least, to use third-person pronouns to denote oneself. But this doesn't seem to go into the meaning of the pronoun. Relatedly, it is possible to use bound third-person pronouns that take (relative to some variable assignments) oneself as value. For instance, an Australian boy can say "Whenever an Australian boy goes to the cricket, he cheers for Australia." So probably premise 1 of the above argument, requiring that generalisations be semantic, is false. If not, premise 2, requiring that semantic constraints on unbound pronouns also constrain bound pronouns is false.

*Objection Two: The SIC is false*

Egan et al (2005) note that the SIC seems to fail in 'game-playing' and 'advice' contexts. So in a game of hide and seek, where Billy is looking for something Suzy has hidden, if he asks "Is it in the basement?", Suzy can truly say "It might be", even if she knows it isn't true. And a parent can tell their child "Wash that strawberry before you eat it; it might be contaminated" even if the parent knows that the strawberry has been washed.

The simplest response to the first objection is that the purported generalisation, a third-person pronoun does not denote the speaker, isn't really a universal generalisation at all. It's possible to refer to oneself by name; certain people in the news do it on a regular basis. For example, a famous footballer Smith might say "Smith deserves a pay raise". In such a context, it isn't at all odd (or at least any odder than it already is) to use third-person pronouns, e.g. by continuing the above utterance with "and if he doesn't get one, he's not going to play".

The second objection is a little trickier, but I think it's possible to understand these utterances as a kind of projection. The speaker is speaking from the perspective of the hearer. This isn't an unattested phenomenon. Something like it must be going on when speakers use "we" to denote the hearer. Imagine, for instance, a policeman coming across a person staggering out of a pub and saying "We've had a bit much to drink it seems". The policeman certainly isn't confessing to dereliction of duty. Nor is this case sufficient to throw out the idea that "we" is a *first person* plural pronoun. Rather, the policeman is speaking from the drunk's perspective. I suspect the same thing is going on in both of the examples above.

So I think both objections can be answered. But neither answer is completely convincing. And there's a worry that the two responses undermine each other a little. If we accept projection is a widespread phenomenon, then perhaps self-denotation with a third person pronoun is a kind of projection. It might be nice then to have another way to restate the argument, without assuming there's a response to this pair of objections.

Let's step back from the details of the SIC. What we started with was a simple fact, that (10) can't be used to express (9). That's not threatened by counterexamples to the SIC, and it still needs explanation. The SIC is a proposed semantic explanation, and perhaps, if it has counterexamples, it fails. But I think we can be independently confident that the explanation here will be semantic, not pragmatic. We can be confident of this because there just doesn't look to be anything like a pragmatic explanation available.

Compare the discussion of third-person pronouns. Even if we can't use third-person pronouns to pick out ourselves (when not speaking projectively), there is an obvious pragmatic explanation for this. We have first-person pronouns available, and if we mean to denote ourselves, using a first-person pronoun will do so in the clearest possible way. Since it is good to be clear, when we pass up the chance to use a first-person pronoun, the obvious conclusion for a hearer to draw is that we don't mean to denote ourselves. The details of this explanation could use some filling out, but it at least has the form of an explanation. It simply doesn't seem that any such explanation will be available for why (10) can't be used to express (9). It's not that (9) isn't a thought that we might be interested in expressing. And it's not that if we wanted to express it, we would have had an obviously preferable form of words to (10). It's true that we have the words in (9) itself, but (a) they are longer, and (b) on the contextualist view whenever we use an epistemic modal there is some such paraphrase available, a paraphrase that typically does not defeat the acceptability of epistemic modals.

If there isn't a pragmatic explanation of why (10) can't be used to express (9), then there must be a semantic explanation. But the only semantic explanation that looks plausible from a contextualist perspective, is a semantic restriction on *X*. And we know, from considering the data about embedded epistemic modals, that there is no such restriction. So we have a problem for contextualism. Slightly more formally, we can offer the following argument against orthodox contextualism about epistemic modals.

1. Whatever acceptability data can't be explained pragmatically must be explained semantically
2. There is no pragmatic explanation for why (10) can't be used to express (9)
3. If 1 and 2 then the meaning of "might" must explain why (10) can't be used to express (9)
4. If the meaning of "might" must explain why (10) can't be used to express (9), and contextualism is true, there must be a restriction, provided by the meaning of "might" on the relevant group *X* that excludes groups like those guys

5. If there is a restriction, provided by the meaning of “might” on the relevant group *X* that excludes groups like those guys, then when “might” is embedded under an attitude verb, the group *X* still can’t be those guys.
6. In “Those guys believe that they might be in trouble”, the relevant group *X* just is those guys
7. So, contextualism is false

#### *4. Against Moral Relativism*

I’m going to close with an argument here that any kind of contextualism, whether orthodox or heterodox, about moral terms, especially “wrong”, does not fit our usage of those terms. The argument is going to be that in order to offer a contextualist-friendly account of the behaviour of “wrong” in belief ascriptions and knowledge ascriptions, we have to suppose that it behaves quite differently in those two settings. But other context-sensitive terms do not behave that way, and we have good theoretical reasons to believe that this is not in general how context-sensitive terms behave. So “wrong” is not context-sensitive, either to contexts of usage or assessment.

At first glance there seems to be very little pattern to the way that contextually sensitive terms behave in attitude ascriptions. We can find terms, like *we*, whose denotation inside a belief ascription is not particularly sensitive to the context of the ascriber. So in (11), *we* is naturally interpreted as denoting the speaker and those around her.

(11) Otto believes that *we* are fools.

We can find terms, like *tasty*, whose denotation inside a belief ascription seems to vary quite a bit depending on the sentence being used. Sometimes *tasty* seems to mean *tasty to the ascriber*, as in (12).

(12) Vinny the Vulture believes that rotting carcasses are *tasty*.

And sometimes it seems to denote something like *tasty to a contextually salient taster* as in (13).

(13) Suzy believes that this kind of dog food is *tasty*.

It is easy to set up a circumstance where this means that Suzy thinks that the salient dog food is *tasty for dogs*. So the relevant taster can be the ascriber, but could also be given by context. It’s not impossible, but not as easy as perhaps it should be, to get the relevant taster to be the speaker, or the speaker and those around them.

On the other hand, when we use epistemic modals in belief reports, the relevant ‘knower’ is always the ascriber. Consider, for example, (14).

(14) Jack believes that Smith might be happy.

That can only mean that Jack believes that for all Jack (and perhaps his friends) knows, Smith is happy. It can't, for instance, mean that Jack believes that for all the speaker knows, Smith is happy. (Unless the speaker is Jack or one of his friends.)

So we have a progression of cases, where in (11) the contextually sensitive term 'we' has to get its denotation from the context of utterance, in (14) the contextually sensitive term 'might' gets its denotation from the context of the ascriber, and (12) and (13) show that 'tasty' can behave in either of these ways. I've been putting this all in terms that will make most sense if we are accepting truth absolutism, but the same points can be made without that assumption if we so desire.

As I said, at first it might look like there is no pattern here at all. But if we look at other attitudes, we see that there is an interesting pattern. The way that 'we', 'tasty' and 'might' behave in belief reports is just the same as they behave in knowledge reports. We can see this in the following examples.

(11a) Otto knows that we are fools.

(12a) Vinny the Vulture knows that rotting carcasses are tasty.

(13a) Suzy knows that this dog food is tasty.

(14a) Jack knows that Smith might be happy.

In each of these sentences, the contextually sensitive term behaves just as it does in the parallel belief report. This isn't too surprising. It would be a real shock if some term *t* behaved quite differently in belief and knowledge reports. If that were the case it would be possible in principle to find a passage of the form of (15) that's true.

(15) S believes that ... *t* .... Indeed S knows it. But S doesn't know that ... *t* ....

It's impossible to survey every instance of (15) to see whether they all sound contradictory. But I suspect that they will sound contradictory. So we'll assume in what follows that context-sensitive terms behave the same way in belief ascriptions and knowledge ascriptions, whether or not the kind of context-sensitivity at issue is orthodox.

The problem for contextualism about "wrong" is that it is forced to violate this principle. Assume that *X* is wrong means that *X* is wrong relative to the standards of some salient group *G*. We'll leave aside for now the question of whether *G* is determined by the speaker's context or the assessor's context, as well as the question of whether the sentence expresses a proposition involving *G*, or a proposition that is true or false relative to groups. We'll also leave aside the question of just what it means for something to be wrong relative to the standards. (Does it mean that *G* actually disapproves of it, or would disapprove of it under reflection, or that it doesn't have

properties that G wants to promote, or something else?) We'll simply assume that there have been people whose standards are different to ours in ways that make a difference for the wrongness of actions. If that isn't the case, we hardly have a contextualism worthy of the name. It's obviously controversial just what could be an example of this, but I'll take as my example Jefferson Davis's belief that helping fugitive slaves was wrong. It seems true to say Davis had this belief, so (16) is true.

(16) Davis believed that helping fugitive slaves was wrong.

Now (16) clearly doesn't mean that Davis believed that helping fugitive slaves was wrong by my standards. And that's not just because he didn't have any de re attitudes towards me. Whatever (false) proposition I would (according to the contextualist) express by saying "Helping fugitive slaves was wrong" is not what Davis believed. He believed something that was made true (if it was) by his moral standards. Now compare (17).

(17) Davis knew that helping fugitive slaves was wrong.

It seems to me that that's just false. And it's false because helping fugitive slaves wasn't, in fact, wrong. That's not to deny that it was wrong by Davis's standards. I suspect that by Davis's standards, helping fugitive slaves was wrong. Maybe his dispositions to accept various general claims about moral standing, in appropriate condition, combined with some facts about the nature of fugitive slaves, meant that he wouldn't in ideal conditions reject helping fugitive slaves. But I doubt it. It seems to me he had deeply ingrained, deeply misguided standards, and by those standards it was wrong to help fugitive slaves. Perhaps he even knew that about his own standards. But that's neither here nor there to the truth of the English sentence (17), which sounds simply false in virtue of the rightness of helping fugitive slaves.

Now neither the truth of (16), nor the falsity of (17) is, on its own, sufficient to undermine contextualism about "wrong". The truth of (16) is consistent with the claim that "wrong" behaves like "might". So in attitude ascriptions, what matters is the ascriber's context. And the falsity of (17) is consistent with the claim that "wrong" behaves like "we", and (17) is false because what helping fugitive slaves was wrong expresses in our context is false.

Rather, the problem is that an adequate account of "wrong" has to account both for the truth of (16) and the falsity of (17). And that doesn't seem to be possible. At least it isn't possible without supposing that "wrong" behaves differently in knowledge reports and belief reports. And we've seen some reasons above to believe that that's not how context-sensitive terms behave, whether the term is one like "we", for which an orthodox theory seems best, or like "might", for which a heterodox theory seems best.

I'll end with some objections that I've encountered since I've started discussing this argument, with my replies to each of them.

*Objection: This is question-begging against the moral relativist.*

This is the most common reaction I've heard, but I do find it hard to make sense of. It is hard to see just which premise is question-begging. Nothing in moral relativism as such prevents us accepting the truth of *Davis believed that helping fugitive slaves was wrong*, and nothing in moral relativism prevents us from rejecting *Davis knew that helping fugitive slaves was wrong*. There is, I say, a problem with doing both of these things, as we should want to do. But if an argument is going to be rejected as question-begging because the other side can't simultaneously accept *all* of its premises, well we won't have many arguments left to work with.

A little more seriously, the relativist theories that I'm opposing here are semantic theories. If we can't reject them because they commit us to endorsing sentences that we (the opponents of the view) can see to be false, then it is hard to know what could count as an argument in semantics. It's no defence of a view to say that its proponents cannot see it is false, if the rest of us can see it.

*Objection: We would see the knowledge claim (17) is true, if only we didn't have anti-relativist prejudices.*

This might well be right; it's certainly hard to know when one is prejudice free. Perhaps all that's going on is that we don't want to be committed in any way to saying that it's wrong to help fugitive slaves, and we're worried that accepting (17) would, in some way, so commit us.

But note how much I've done to stack the deck in favour of pro-relativist intuitions, and dissipate this worry. The argument is coming at the end of a whole paper defending relativism. Earlier in this very paper I defended some relativist views from arguments using factive attitude verbs by noting that it is tricky to state just what factivity comes to. In particular, I noted that we can sometimes say *A knows that S* in circumstances where we would not, indeed could not, utter *S*. And I repeated this observation at the start of this section. I think I've done as much as possible to (a) overcome anti-relativist prejudice, and (b) frame the argument in such a way as to make it as easy as possible to accept (17). But even in this frame, I still say we can see that it is false.

*Objection: This is only an objection to one kind of context-sensitivity in moral terms, the kind we associate with traditional moral relativism. But it doesn't show that moral terms are in no way context-sensitive. We'd expect that they are, since some moral terms are gradable adjectives, and gradable adjectives are context-sensitive.*

There's a really interesting philosophical position around here. Start with the idea that we should be invariantists, perhaps realists in some quite strong sense, about moral comparatives. Perhaps this could be tied to the idea, gaining currency in some circles, that comparatives are what's crucial to morality. Then say that terms like "right" and "wrong" pick out, in a context-sensitive way, points on the moral scale. So some kind of contextualism, presumably orthodox, is right for those terms. This position is immune to the objection given above, because (16) turns out to be

true, and (17) false, if we interpret “wrong” to mean *above the actually contextually-salient level of wrongness, on the objectively correct wrongness scale*.

But I think a similar pair of examples show that this won’t work. Assume that we’re talking about various people’s charitable giving in a context where we don’t hold people to super-high standards. So the charitable actions of, say, Bill Gates count as laudable in our circumstances. (I assume that on the merits Gates’s donations are for the good; determining whether this is true is well outside the scope of this paper.) Now consider a philosopher, call him Peter, who doesn’t believe in the moral supererogatory, so he thinks anything less than the best you can do is morally wrong. It seems to me that, as uttered in our context, (18) expresses a truth and (19) a falsehood.

- (18) Peter believes that Bill Gates’s level of charitable donation is wrong.
- (19) Peter knows that Bill Gates’s level of charitable donation is wrong.

And I don’t think there’s a good contextualist explanation for this pair of judgments. If “wrong” was just a simple context-sensitive term in the way suggested, then (18) should be false, because Peter doesn’t believe that Bill Gates’s level of charitable donations do rise to the level of wrongness that is, as it happens, is salient in our context. But intuitively, (18) is true.

The same kind of objection can be raised to a more prominent kind of theory that takes a certain kind of normative standard to be set by context, namely classical contextualism about “knows”. Assume it is common ground that George has excellent, but not irrefutable, evidence that he has hands. Assume also that we’re in a low standards context for knowledge. And assume that René thinks knowledge requires objective certainty. Then it seems that we can truly say (20), but not (21).

- (20) René believes that George doesn’t know he has hands.
- (21) René knows that George doesn’t know he has hands.

Again, the pattern here is very hard to explain on any kind of contextualist theory, be it orthodox or heterodox.

*Objection: Normative terms might be sui generis. Perhaps they are the only counterexamples to the pattern in (15).*

Anything could be the exception to any rule we like. But it’s bad practice to assume that we have an exception on our hands. If we heterodox theorists simply responded to von Fintel and Gillies’ argument from factive verbs by saying that we had an exception to a general pattern here, we would, quite rightly, not be taken seriously. The contextualist about normative terms should be held to the same standard.

*Objection: Relativism does have counterintuitive consequences, but all theories have some counterintuitive consequences. Arguably everyone is going to have to accept some kind of error theory, and this is a relatively harmless kind of error to attribute to the folk.*

If we were convinced, perhaps by one or other of the contemporary developments of Mackie's argument from queerness, that no non-relativistic moral theory is possible (apart from a Mackie-style moral error theory), that would be an interesting argument for moral relativism. Certainly I'd be more willing to accept that (16) and (17) don't have the same kind of context-sensitivity than I would be to accept that, say, it's not wrong to torture babies.

It is well beyond the scope of this paper to adjudicate such debates in any detail, but I am a little sceptical that we will ever face such a choice. Generalising wildly, most of the time our choice is between (a) accepting a moral error theory, (b) accepting some odd semantic consequences, as outlined here, or (c) rejecting some somewhat plausible claim about the nature of moral judgment, such as motivational internalism. Arguments that present a trilemma such as this deserve to be judged on their merits, but my feeling is that we should normally take option (c). That's not to say this objection couldn't work, or that the argument I've offered is a long way from a refutation of relativism. It clearly is not. But I think it is a challenge that moral relativists and contextualists should face up to.