No Royal Road to Relativism

Brian Weatherson

January 18, 2010

Relativism and Monadic Truth is a sustained attack on 'analytical relativism', as it has developed in recent years. The attack focusses on two kinds of arguments. One is the argument from the behaviour of operators, as developed by David Lewis (1980) and David Kaplan (1989). The other kind of argument takes off from phenomena concerning speech reports and disagreements. Such arguments play central roles in arguments by, among others, Andy Egan (2007), Max Kölbel (2009), Peter Lasersohn (2005), John MacFarlane (2003, 2007), Mark Richard (2004) and Tamina Stephenson (2007). These arguments also play a role in a paper that I co-authored with Andy Egan and John Hawthorne (Egan, Hawthorne, & Weatherson, 2005).

As the reader of *Relativism and Monadic Truth* can tell, John Hawthorne no longer much likes the arguments of that paper, nor its conclusions. And I think he's right to be sceptical of some of the arguments we advanced. The objections that he and Herman Cappelen raise to arguments for relativism from speech reports and from disagreement are, I think, telling. But I don't think those are the best arguments for relativism. (For what it's worth, I don't think they're even the best arguments in the paper we co-authored.) The primary purpose of this note will be say a little about what some of these better arguments are. The core idea will be that although there is some data that is *consistent* with non-relativist theories, the best explanation of this data is that a kind of relativism is true. In short, we should be looking for inductive, not deductive, arguments for relativism. I'm going to fill in some details of this argument, and say a little about how it seemed to slip out of the main storyline of *Relativism and Monadic Truth*.

In Chapter 2 of *Relativism and Monadic Truth*, Cappelen and Hawthorne attempt to develop diagnostics for when an utterance type *S* has invariant content. They note that some relativist arguments presuppose a diagnostic based on speech reports. The idea behind the presupposed diagnostic is that if we can invariably report an utterance of *S* by *A* by saying *A said that S*, then *S* is semantically invariant. And they note that this diagnostic isn't particularly reliable.

What they aim to replace it with is a diagnostic based on agreement reports. Cappelen and Hawthorne are more careful on the details than I'll be, but the rough idea is easy enough to understand. The diagnostic says that if whenever *A* and *B* utter *S*, we can report that by saying *A and B agree that S*, and the basis for our saying this is that they made those utterances, then *S*'s content is invariant. The idea behind the test is that if there isn't a single proposition that *A* and *B* endorse, then it would be odd to say that they agree.

I don't think the diagnostic is particularly plausible. The next couple of paragraphs won't come as much of a surprise to the authors of *Relativism and Monadic Truth*, since the ideas come from a talk Herman Cappelen did at the Arché Summer School in July 2009. But they are central enough to the story I'm telling that they are worth including here. The core problem for this agreement based diagnostic is that sometimes we can report parties as agreeing even though they don't agree on the truth value of any proposition. So while (1) has a disambiguation where it is true only if something like (2) is true, it also has a disambiguation where it is true as long as something like (3) is true.

(1)   Alec, Pierre and Franz agree that they were lucky to be born where they were actually born.

(2)   Alec, Pierre and Franz each consider each of Alec, Pierre and Franz lucky to be born where they were actually born.

(3)   Alec, Pierre and Franz each consider themselves lucky to be born where they were actually born.

Given that sentences like (1) could mean something like (3), there is little reason to think that agreement diagnostics will provide us clear evidence of sameness of content. Indeed, Cappelen and Hawthorne should hope that this diagnostic doesn't always work, because the diagnostic seems to entail relativism about epistemic modals. Imagine that a detective and a psychic are both investigating a murder. They both conclude that their evidence entails that Tatort did it, and that their evidence is consistent with Tatort being dead. They are, however, ignorant of each other's evidence, and indeed of the fact that the other is working on the investigation. Still, if each utters (4), it seems we are in a position to endorse (5).

(4)   Tatort must be guilty, and might be dead.

(5)   The detective and the psychic agree that Tatort must be guilty, and that he might be dead.

This will be very hard to explain on a contextualist theory of epistemic modals, if we accept the agreement diagnostic. That's because there's no proposition (other than the proposition that Tatort is guilty) that they agree about. I think this is *some* evidence in favour of relativism, but if the contextualist wanted to argue that we should understand (4) the same way we understand (1) (on its 'distributive' disambiguation), it would be hard to conclusively show they were wrong.

In any case, it is hard to see why we should expect there to be a diagnostic of the kind Cappelen and Hawthorne are aiming for. Such diagnostics are the exception, not the rule, in social sciences. There's no simple diagnostic for whether a particular state is democratic or not. (Is modern-day Afghanistan a democracy? What about modern-day Alabama?) Nor is there a simple diagnostic for whether a particular rule is a law. (Are internal revenue regulations laws?) But political science and jurisprudence don't collapse in the absence of such diagnostics. Nor should philosophical semantics collapse in the absence of a simple test for context-sensitivity.

Indeed, the situation is political science and in jurisprudence is in one sense worse than it is in semantics. We can state, admittedly in theory-laden terms, what it is for the content of a sentence type to be context-invariant or context-sensitive. It is much harder to state, even in theory-laden terms, what it is for a state to be democratic, or for a rule to be a law. The problem with thinking about the questions I asked in the previous paragraph isn't that there's some hidden piece of evidence we haven't yet uncovered. It's that the concepts do not have clear application conditions, and the hard cases fall between the clear instances and non-instances of the relevant property. In semantics we have, to a first approximation, a mere epistemic challenge.

Even if the hunt for a diagnostic for context-sensitivity is bound to be futile, as I think it is, that doesn't mean it is harmless. I think the structure of Cappelen and Hawthorne's inquiry, which starts by looking for a test and then goes on to apply it, pushes us towards the wrong kind of argument. The effect of this structure is that we end up looking for *deductive* arguments for or against relativism, and the absence of deductive arguments for relativism is taken to be a big problem for the relativist. But we should have been looking for *inductive* arguments. The best case for relativism, I think, will be a kind of inference to the best explanation. For instance, a relativist might try to clean up this argument.

1. Our best theory of mental content is that the contents of beliefs and desires do not satisfy Simplicity.[1]

---

[1] *Simplicity* is Cappelen and Hawthorne's name for the conjunction of theses they want to defend against the rel-

2. The role of language is to express thoughts, so if the contents of belief and desire do not satisfy Simplicity, the contents of sentences and utterances probably don't either.

3. Simplicity is false as a theory of linguistic content.

This argument clearly isn't valid. That's by design; it's meant to be an abductive argument against Simplicity about linguistic content. And of course both premises are controversial. There's one argument for premise 1 in (Lewis, 1979), and another in (Perry, 1979). Both arguments are controversial. Indeed Cappelen and Hawthore spend some time (pages 50 to 54) responding to the Lewisian arguments, though they spend less time on Perry's arguments.

I'm not going to try to advance the debate here over whether premise 1 is true or not. I suspect the solution will turn on much bigger issues than can be covered in a note of this length. And that's because I think the judgment about whether premise 1 is true will turn on quite global features of our best theory of mental content. For instance, Daniel Nolan (2006) argues that there are certain desires that we cannot understand on the modal Lewis offers. That doesn't *entail* that Lewis is wrong about the nature of belief, but it does make Lewis's theory of belief look less attractive. From the other direction, many authors working on the Sleeping Beauty problem, dating back to the problem's introduction to the philosophical community in (Elga, 2000), have felt that the problem was best approached in Lewis's Simplicity-unfriendly framework. That doesn't entail Simplicity is wrong, but it is I think evidence against it. On the other hand, Robert Stalnaker (2008) has recently argued that this is not the best framework for thinking about the Sleeping Beauty problem, and I've argued (Weatherson, forthcoming) that Stalnaker's approach lets us see things about the Sleeping Beauty puzzle that are hidden on the standard, Lewisian, approach. So if we're going to evaluate this kind of argument for relativism, the issues are going to get far removed from familiar disputes about distributions of words and phrases. That's not too surprising. In general, the hard thing about abductive reasoning in philosophy is that we have to start looking at all sorts of different kinds of evidence. But that's no reason to think that the most telling arguments won't, at the end of the day, be abductive arguments.

A quite different kind of argument comes from thinking about property ascription and ignorance. It's a somewhat frequent occurrence that modern science discovers that some of our thoughts

---

ativist. For our purposes, Simplicity about mental content is the view that the contents of beliefs and desires are propositions, and these propositions are simply true or simply false, not merely true or false relative to some or other parameter. Simplicity about linguistic content is the view that these same propositions, the ones that are simply true or false, are the contents of declarative utterances.

seem to depend for their truth on more variables than we realised. So it isn't true that two accelerating objects simply have the same mass or different masses; rather, their relative mass might be different relative to different inertial frames. Or two colour patches might not be simply the same colour or simply different colours. If the colours are metamers (relative to human vision) then they will be in a good sense the same colour relative to human vision, and different colours relative to more discriminating detectors. Such cases raise challenges for the project of interpreting a language.

Assume that the community uses terms like 'mass'. Indeed, assume they are sophisticated enough to distinguish mass from weight, for they know that weight is relative to a gravitational field, and gravitational fields vary in strength. But they are not sophisticated enough to know that masses are relative to inertial frames. The members of this community frequently go around saying things like "Those two objects have the same mass," referring to a and b. Call that sentence M. We assume that the members are in a particular inertial frame, call it $\mathcal{F}$. Let's assume (just for a few paragraphs) that the propositions that satisfy Simplicity are structured, and assume that we can represent the relation *has the same mass as* by a somewhat unstructured relation *SameMass*. (In other words, ignore whatever internal structure *SameMass* has, since it won't be relevant to this example.) Then it seems to me that there are three live options around.

1. By M, the speakers express the pseudo-proposition $SameMass(a,b)$, and this pseudo-proposition is not capable of being true or false, since *SameMass* is a three-place relation (between two objects and an inertial frame) and only two places are specified.
2. By M, the speakers express the proposition $SameMass(a,b,\mathcal{F})$, and this proposition is (capable of being) true.
3. By M, the speakers express the proposition $SameMass(a,b)$, and this is (capable of being) true relative to $\mathcal{F}$, although it might be false relative to some other inertial frame $\mathcal{F}'$.

If option 3 is correct, then it seems Simplicity fails.[2] So if there are compelling arguments against options 1 and 2, and those are all the options, then Simplicity is in trouble. And it seems the

---

[2] I say 'seems' since I'm not sure exactly what it takes for there to be a notion of truth *simpliciter*. The argument on page 96 against the conjunction of Simplicity, Eternalism and Temporalism suggests that Cappelen and Hawthorne believe the following principle: If p is true in C_1, and false in C_2, and C_1 and C_2 both exist, then p is not either simply true or simply false. It's not obvious to me why p couldn't, in C_1, be simply true, but I take it Cappelen and Hawthorne are using 'simply' in such a way as to exclude that. So option 3 is inconsistent with Simplicity.

relativist might make progress by pushing back against both of those options.

The simplest argument against option 1 is that it violates even a very weak form of the Principle of Charity. Obviously there are very many different kinds of charity principles. For instance, there are three different versions endorsed in (Davidson, 1970), (Lewis, 1974) and (Williamson, 2007, Ch. 8). But any kind of Charity will imply that options 2 or 3 are preferable to option 1, since option 1 will imply that the subjects don't even have beliefs about the relative masses of objects, whereas the other options will imply that their beliefs may well be true, and rational, and even in some cases amount to knowledge. An alternative argument against option 1 is that the members of that community would have been right to take it as a Moorean fact that some things have the same mass.[3] So option 1 doesn't look overly plausible.

One argument against premise 2 is that it is impossible for the members of the community, given their powers of individuation, to make singular reference to such a thing as an inertial frame. If they don't know what an inertial frame is, then we might be sceptical of claims that they can refer to it. (Note that the thought here isn't merely that some individuals don't know what inertial frames are; the imagined case is that even experts don't know about the kind of things that we would need to put into the propositions to give them simple truth values.) Another argument is that competent speakers of the language should be able to identify the number of argument places in the properties they use.

Neither of the arguments just offered is completely compelling, though I think both are at least promising avenues for research. But both arguments do look notably weaker if we drop the assumption that the relevant propositions are structured. In an unstructured propositions framework, we perhaps don't need to worry about the members making singular reference to things like inertial frames. We just need to have the speakers pick out (in a perhaps imperfect way) the worlds in which their beliefs are true. And in an unstructured propositions framework it isn't clear that being unable to identify the number of arguments places in the properties they use is any more of a sign of linguistic incompetence than not knowing the individuals to which they refer. But it is a commonplace of semantic externalism that speakers can refer without knowing who it is they are referring to.

The arguments in the previous two paragraphs have been sketchy, to say the least. But if they can be developed into compelling arguments, then it might turn out that the case against option 2

---

[3]Compare the discussion of Moorean facts in (Lewis, 1994, 489).

succeeds iff propositions are structured. In that case the argument for Simplicity will turn on a very large question about the nature of propositions, namely whether they are structured or not. Again, the take home lesson is that debates in this area are not susceptible to easy resolution.

I'll end with a more narrowly linguistic abductive argument for relativism and against Simplicity. I think you can find the core ingredients of this argument in (Egan et al., 2005), though it isn't as well individuated as it might have been. The argument takes off from what looks like a somewhat misleading claim in Cappelen and Hawthorne's book. The context is a discussion of autocentric and exocentric uses of predicates.[4] The distinction between autocentric and exocentric uses is important for thinking about the way various predicates are used, though it isn't easy to give a theory-neutral characterisation of it. Assuming contextualism, Cappelen and Hawthorne note that it is easy to explain the distinction: "a use of a taste predicate is autocentric iff its truth conditions are given by a completion that indexes the predicate to the subject" and exocentric iff "its truth conditions are given by a completion that indexes it to a person or group other than the speaker, which may, however, include the speaker." (104) The core idea here is clear enough, I hope, though as they say it requires a slightly different gloss if we assume relativism.[5]

The problem is what they go on to say about epistemic modals. In a footnote they say,

> [I]t is worth noting that there is a similar contrast between autocentric and exocentric uses of epistemic modals. If I see Sally hiding on a bus then I might in a suitable context say 'She is hiding because I might be on the bus' even though I know perfectly well that I am not on the bus. ('Must' is harder to use exocentrically, though we shall not undertake to explain this here.) (104n7)

The parenthetical remark seems mistaken, or at least misleading, and for an important reason. It is true that it is very hard to use 'must' exocentrically in a sentence of the form *a must be F*. But that doesn't mean that it is hard to use 'must' exocentically. In fact it's very easy. Almost any sentence of the form *S believes that a must be F* will have an exocentric use of 'must'. That's because almost

---

[4]The terminology is from (Lasersohn, 2005).

[5]As they also go on to note, things get very complicated in cases where the truth conditions turn on the nature of an idealised version of the speaker. An example of such a theory is the theory of value in (Lewis, 1989). One of the key points that Cappelen and Hawthorne make, and I think it is a very good point against a lot of claims for relativist theories concerning predicates of personal taste, is that this kind of case is very common when it comes to evaluative language.

any use of an epistemic modal in the scope of a propositional attitude report will be 'bound' to the subject of that report. (I put 'bound' in scare quotes because although contextualists will think of this as literally a case of binding, non-contextualists may think something else is going on.) This suggests an argument for relativism about epistemic modals, one that seems to me to be quite a bit stronger than the arguments for relativism discussed in *Relativism and Monadic Truth*.

1. Unembedded uses of epistemic modals are generally autocentric (except in the context of explanations, like 'because I might be on the bus').
2. Epistemic modals embedded in propositional attitude reports are generally exocentric.
3. There is a good, simple relativist explanation of these two facts.
4. There is no good, simple explanation of these facts consistent with Simplicity.
5. So, relativism is true, and Simplicity is false.

Note that I'm not for a minute suggesting that there is no Simplicity-friendly explanation of the facts to be had; just that it won't be a very good explanation. Nor am I suggesting that the phenomena obtain universally, rather than just in most cases. But they obtain often enough to need explanation, and the best explanation will be relativist. And that, I think, is a reason to like relativism.

The simplest relativist explanation of premises 1 and 2 uses the idea, derived from (Lewis, 1979), that contents are λ-abstracts. So the content of *a must be F* is roughly $\lambda x.(x$'s evidence entails that a is $F)$. A content $\lambda x.\phi(x)$ is true relative to a person iff they are ϕ, and believed by a person iff they consider themselves to be ϕ, under a distinctively first-personal mode of presentation. Then a typical utterance of *a must be F* will be autocentric because if the asserter thinks it is true, they must take themselves to satisfy $\lambda x.(x$'s evidence entails that a is $F)$. So assuming they are speaking truly, the hearer can infer that the speaker's evidence does indeed entail that a is F. But a typical utterance of *S believes that a must be F* will be true just in case S takes themselves to satisfy $\lambda x.(x$'s evidence entails that a is $F)$, and hence will be true as long as S's evidence, or at least what S takes to be their evidence, entails that a is F. There's no reference there to the speaker's evidence, so the use of 'must' is exocentric. There are, to be sure, many details of this explanation that could use filling in, but what is clear is that there is a natural path from the view that contents are λ-abstracts to the data to be explained. And that explanation is inconsistent with Simplicity.

Is there a good, simple Simplicity-friendly explanation of the data around? I suspect there is

not. There are two obvious places to look for a Simplicity-friendly explanation. We could look for an explanation that turns on the meaning of 'must', or we could look for an explanation in terms of salience. On closer inspection, neither avenue is particularly promising.

It does seem likely that there is an available explanation of premise 1 in terms of the meaning of 'must'. The contextualist about pronouns has an easy explanation of why 'we' almost always picks out a group that includes the speaker. The explanation is just that it is part of the meaning of 'we' that it is a *first-personal* plural pronoun, so it is part of the meaning of 'we' that the group it picks out includes the speaker. We could argue that something similar goes on for 'must'. So *a must be F* means, roughly, that x's evidence entails that a is F, and it is part of the meaning of 'must' that x either is the speaker, or is a group that includes the speaker. The problem with this explanation is that it won't extend to premise 2. And that's because meanings (in the relevant sense) don't change when we move into embedded contexts. For example if Jones says "Smith thinks that we will all get worse grades than she will get", Jones isn't accusing Smith of having the inconsistent belief that she will get lower grades than what she gets. Rather, the reference of 'we' is still a group that includes the *speaker*, not the subject of the propositional attitude report. On this model, you'd expect the truth condition of *S believes that a must be F* to be that S believes that x's evidence entails that a is F, where x is the speaker, or a group containing the speaker. But that's typically not at all what it means. So this kind of explanation fails.

The problem for salience based explanations of premises 1 and 2 is that salience is too fragile an explanatory base to explain the data. Let's say that in general we think *a must be F* means, roughly, that x's evidence entails that a is F, and x is generally the most salient knower in the context. Then we'd expect that it would be not too hard to read (6) in such a way that its truth condition is (6a) rather than (6b), and (7) in such a way that its truth condition is (7a) rather than (7b).

(6)   Jones's evidence must settle who the killer is.

    a. Jones's evidence entails that Jones's evidence settles who the killer is.

    b. Our evidence entails that Jones's evidence settles who the killer is.

(7)   Smith believes that Jones's evidence must settle who the killer is.

    a. Smith believes that Jones's evidence entails that Jones's evidence settles who the killer is.

    b. Smith believes that her evidence entails that Jones's evidence settles who the killer is.

After all, (6) and (7) make Jones's evidence *really salient*. That evidence settles who the killer is! But, it seems, that isn't salient enough to make (6a) or (7a) the preferred interpretation. That seems to be bad news for a salience-based explanation of the way we interpret epistemic modals.

Like all abductive arguments, this argument is far from conclusive. One way for a proponent of Simplicity to respond to it would be to come up with a neater explanation of premises 1 and 2 in our abductive argument, without giving up Simplicity. Another way would be to argue that although there is no nice Simplicity-friendly explanation of the data, the costs of relativism are so high that we should shun the relativist explanation on independent grounds. I don't pretend to have ready responses to either of these moves. All I want to stress is that these abductive arguments are generally stronger arguments for relativism than the arguments that are, correctly, dismissed in *Relativism and Monadic Truth*. Those arguments try to take a quick path to relativism, claiming that some data about reports, or disagreement, or syntax, entails relativism. I doubt any such argument works, in part because of the objections that Cappelen and Hawthorne raise. There is, as my title says, no royal road to relativism. But I doubt there's a quick road away from relativism either. If the relativist can explain with ease patterns that perplex the contextualist, we have good reason to believe that relativism is in fact true.

References

Davidson, D. (1970). Mental events. in L. Foster & J. W. Swanson (Eds.), *Experience and theory* (p. 79-101). London: Duckworth.

Egan, A. (2007). Epistemic modals, relativism and assertion. *Philosophical Studies*, *133*(1), 1-22.

Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic modals in context. in G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (p. 131-169). Oxford: Oxford University Press.

Elga, A. (2000). Self-locating belief and the sleeping beauty problem. *Analysis*, *60*(2), 143-147.

Kaplan, D. (1989). Demonstratives. in J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (p. 481-563). Oxford: Oxford University Press.

Kölbel, M. (2009). The evidence for relativism. *Synthese*, *166*(2), 375-395.

Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, *28*(6), 643-686.

Lewis, D. (1974). Radical interpretation. *Synthese*, *27*(July-August), 331-344.

Lewis, D. (1979). Attitudes *De Dicto* and *De Se*. *Philosophical Review*, *88*(4), 513-543.

Lewis, D. (1980). Index, context, and content. in S. Kanger & S. Öhman (Eds.), *Philosophy and grammar* (p. 79-100). Dordrecht: Reidel.

Lewis, D. (1989). Dispositional theories of value. *Proceedings of the Aristotelian Society*, *Supplementary Volume 63*,

113-137.

Lewis, D. (1994). Humean supervenience debugged. *Mind*, *103*(412), 473-490.

MacFarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly*, *53*(212), 321-336.

MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies*, *132*(1), 17-31.

Nolan, D. (2006). Selfless desires. *Philosophy and Phenomenological Research*, *73*(3), 665-679.

Perry, J. (1979). The problem of the essential indexical. *Noûs*, *13*(4), 3-21.

Richard, M. (2004). Contextualism and relativism. *Philosophical Studies*, *119*(1-2), 215-242.

Stalnaker, R. (2008). *Our knowledge of the internal world*. Oxford: Oxford University Press.

Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, *30*(4), 484-525.

Weatherson, B. (forthcoming). Stalnaker on sleeping beauty. *Philosophical Studies*.

Williamson, T. (2007). *The philosophy of philosophy*. Oxford: Blackwell.