

Propositional Attitudes

PHIL 434/634, Week 8: 28 Feb. 2012

Bruno Whittle

ROSS AND SCHROEDER ON BELIEF (CONTD); AND BELIEF, INFORMATION AND REASONING

1. Ross and Schroeder

The Ross and Schroeder account: believing that *p* involves having a defeasible disposition to treat *p* as true in reasoning.

The hope is that (a) this account can explain pragmatic encroachment; and (b) it compares favourably with 'Pragmatic Credal Reductivists' (e.g., Weatherson).

1.1. Correctness

Intuitively, if you believe that *p*, and *p* turns out to be false, then you were *wrong* about *p*.

But (the objection to PCR goes) how can a view on which believing that *p* can be constituted merely by having a high degree of belief in *p* accommodate this thought?

How can the R&S account do better? The idea is that to believe that *p* is to have a disposition to treat *p* as true; but that is to 'operate with a conception' that corresponds to the way the world is iff *p* is true; it is thus in *this* sense in which you would be *wrong* if *p* turned out to be false.

1.2. Stability

This criticism of PCR uses the following principle:

Stability. A fully rational agent does not change her beliefs purely in virtue of an evidentially irrelevant change in her credences or preferences. (p. 20)

PCR falls foul of this principle: e.g., in the low stakes sandwich case, according to PCR, a fully rational agent would go from believing to not believing if they increased their credence that their guest had an allergy.

Perhaps the reasoning disposition account can uphold stability. But is it slightly funny for a defender of pragmatic encroachment to wield stability? After all, if we can go from knowing to not knowing purely on the basis of changes in 'evidentially irrelevant pragmatic factors', then, if you rationally adjust your credences about such factors, and if your beliefs aim to be knowledge, then is it not rational to go from believing to not believing? (This may seem contentious however!)

But one might also worry that pragmatic encroachment might more directly violate principles in the ballpark of stability. E.g.:

An agent does not go from knowing to not knowing (or vice versa) purely in virtue of evidentially irrelevant changes in their practical situation.

Or perhaps a line could be drawn here?

3.3. *Inconsistency*

PCR has issues with inconsistency.

How do R&S avoid this? They say that having inconsistent beliefs would be ‘cognitively costly’. ?

BELIEF, INFORMATION AND REASONING

1. A Puzzle

The following seem (to me) to be prima facie natural ideas about belief.

(I) Beliefs are our means of storing information.

E.g., the way in which we store information about what snow is like is by having beliefs about it (such as the belief that snow is white).

(R) If S believes that p then she is willing to use p in reasoning.

E.g., if Rachel believes that snow is white, then she will be willing to use this proposition in her reasoning.

((R) would ultimately have to be qualified to allow for ‘over-riding’ factors; but I won’t try to state a more precise version.)

(I) and (R) cannot both be true, however, as the following example shows.

Part 1. Rachel is thinking about what she is going to do next summer. She would love to go on safari. But, after checking her bank balance, and thinking about how much she is going to earn between now and then, she reluctantly concludes that she is going to have to stay in the US. I.e., she comes to believe that she will be in the US next summer.

Part 2. She goes to a deli to buy a sandwich. While she is there she thinks about buying a lottery ticket. She thinks: if I win, then I will be on safari next summer; and she decides to buy the ticket.

In this situation (i.e., in the deli), Rachel is not willing to use the proposition that she will be in the US next summer in her reasoning.

- If she was, she would conclude that there is no point in buying the ticket
- (for she would reason that if she buys the ticket and it wins, then she will be on safari, and not in the US, next summer;
- thus, since she *will* be in the US, if she buys the ticket it won’t win).

So it follows from (R) that she no longer believes that she will be in the US next summer. But then it follows from (I) that she has lost the information she acquired at the start of the example:

- for if our store of information is simply our beliefs, and Rachel has lost the belief she acquired, then it follows that she has lost the information she acquired.

Part 3. But of course that's not right. Because if when Rachel gets home she takes a call from a friend, who wants to know if she is going to be around next summer, then Rachel will be able to answer 'Yes' immediately.

Thus, she has *not* lost the information she acquired in part 1: because it is still available for her to use to guide her answer; after all, it is not as if she has to revisit the facts about her bank balance and so on; rather, she simply uses the result of her earlier deliberations.

So (I) and (R) cannot both be true. (So that is the puzzle!)

Note also that this is not an isolated example. E.g., a case when someone is deciding when to buy a new computer, then thinking about insurance, then taking a call; etc.

2. A New Picture of Belief

The solution, I want to argue, is a new picture of belief. On this picture, there are two distinct components:

- on the one hand, there is our information; and,
- on the other, there is what we are taking seriously at any given point.

Each of these can be thought of as a set of worlds (i.e., a set of metaphysically possible worlds; but they can also be thought of in terms a more fine-grained notion of content; it is just that we don't the 'added granularity' to solve the puzzle, and so I put things initially in terms of metaphysically possible worlds).

Thus, the first component on this picture is our 'information set' = intuitively, the set of worlds our information leaves in play.

The second component is our 'relevance set' = the set of worlds we are taking seriously at that point (i.e., taking seriously as candidates for how the world is).

E.g., in a typical situation in which Rachel is thinking about what she is going to do this weekend, she will take seriously (typical) worlds where she goes to visit her family, and also (typical) worlds where she stays at home; but she will not take seriously worlds at which she is hit by a bus on Friday morning, or at which there is a nuclear holocaust on Friday afternoon.

(I use the term 'ruling out' for eliminating from our information set.)

These two components then give rise to two kinds of belief, as follows.

Consider again Rachel's deliberations about what she is going to do this weekend. And suppose that, in the end, she decides to visit her parents. On the picture proposed, what is going on in this sort of case is that Rachel is ruling out (i.e., eliminating from her information set) all *relevant* worlds (i.e., worlds that she is taking seriously) at which she does *not* visit her parents this weekend. So, the only relevant worlds *remaining* are those at which she *does* visit them.

Thus, in one sense, she believes that she will visit her parents this weekend: i.e., in the sense that this proposition is true at every relevant world in her information set. But, in another sense, she does not believe this: because there are still worlds in her information set that this proposition is not true at.

Thus, more generally, say that we *I-believe* p if it is true at every world in our information set; and say that we *R-believe* p if it true at every *relevant* world in our information set.

(Cf. David Lewis's account of knowledge, on which 'S knows that p' is true in context C iff she has eliminated all of the not-p-worlds that are relevant in C. Thus, very roughly, the idea behind the proposal is that belief has something like the structure Lewis thought knowledge had.)

And the idea is *then* that information goes with I-belief, while reasoning goes with R-belief. Thus (I) and (R) are replaced by:

- (I') I-beliefs are our means of storing information.
- (R') If S R-believes that p then she is willing to use p in reasoning.

So, in the example of Rachel thinking about this weekend:

- it's not part of her information that she will visit her parents this weekend;
- but—despite this—she will be willing to use this in her reasoning.

And the idea is that this is how things go more generally: we store information by having I-beliefs; but we are nevertheless willing to use in our reasoning anything that we merely R-believe.

(One might wonder: are these really kinds of *belief*? I think it is natural to so describe them, on the basis of their connections to information and reasoning; but all I really want to claim is that we have these two mental states, one of which is tied to information, and one of which is tied to reasoning.)

How does this picture solve the puzzle?

Start with part 1: here (the idea is) Rachel is taking seriously worlds at which she has saved enough money to go on safari; and she is also taking seriously worlds at which she has not saved enough, and at which she instead spends the summer in the US.

But (the idea goes) she is *not* taking seriously worlds at which she is kidnapped and forced to work as a safari guide; or worlds at which she wins the lottery and uses the winnings to go on safari.

Then, what Rachel does (in this part of the example) is to rule out those worlds that she is taking seriously and at which she is not in the US next summer. (But she does *not* rule out any of the more far-fetched worlds that she is not taking seriously—so these remain in her information set).

Thus: by the end of part 1, Rachel R-believes that she will be in the US next summer. So, at this point, Rachel R-believes that she will be in the US next summer (even though she does not I-believe it).

Next part 2: what happens (according to the proposal) is simply that certain worlds that were irrelevant become relevant. E.g., worlds at which Rachel wins the lottery and uses her winnings to go on safari next summer.

Thus, since has not ruled such worlds out, she no longer counts as R-believing that she will be in the US next summer. So (R') no longer gives the result that she is willing to use this proposition in her reasoning.

But—crucially—we get this change in what Rachel is willing to use in her reasoning despite the fact that (on this picture) her information remains the same. So: we no longer get the unacceptable result that Rachel has lost the information that she acquired in part 1.

That is, unlike when we had (I) and (R), we can have a change in what Rachel is willing to use in reasoning without a subsequent change in her information (because (R') and (I') now tie reasoning to R-beliefs, while information is instead tied to I-beliefs).

Consequently, on this picture, this information *is* still available for Rachel for to use in part 3. More precisely, what happens (on the proposal) in part 3 is simply that the worlds that *became* relevant in part 2 go back to being *irrelevant*. And, thus, Rachel goes *back* to R-believing that she will be in the US next summer (and so (R') now gives the *desirable* result that Rachel is now willing to use this proposition in her reasoning).

In summary: on this picture, we seem to get all the right results about when Rachel acquires and loses information, and about what she is willing to use in her reasoning—while retaining versions of (I) and (R). So the picture seems to solve the puzzle.

3. Solving the Puzzle with the New Picture?

The natural thing to try is to use degrees of belief. (So this would then promise to be a more economical solution: because we seem to think we have these *anyway*.)

One option would be PCR approach: but we've seen how problematic *they* are. (!!)

Alternatively, one might think that a simple 'varying threshold' picture might be able to handle the puzzle.

So the idea would be: in any given situation, there is some number x with $0 \leq x \leq 1$, and you count as believing something (in that situation) iff you believe it to degree $\geq x$. But the thought would be that different situations have different x 's.

So the idea would be one would replace (I) with:

(I*) Degrees of belief are our means of storing information.

But the hope would be that the threshold account of belief would allow one to hold onto (R) unchanged.

So the idea would be that all that happens in the example is that Rachel's threshold is higher in part 2 than in parts 1 and 3; but her degrees of belief (i.e., her information) is constant.

The problem however is that there are variants of the example that this solution won't work with: because we can easily change things so that Rachel *is* willing to use in her reasoning in part 2 things that she believes to degree less than the degree to which she believes that she will be in the US next summer. Since she is willing to use these things in her reasoning just as she uses her beliefs, and she would assert them, etc., we presumably want to say that she believes them.

But then her threshold must be sufficiently low that she *also* would count as believing that she will be in the US next summer—if she retained the same degree of belief.

So, by (R), it once again follows that she has lost the information that she acquired—contrary to fact.

Thus it is far from clear that degrees of belief offer an alternative solution of the puzzle.