

Another puzzle about belief

Graham Seth Moore

In this paper I would like to present another puzzle about belief. This puzzle is distinct from the classic puzzles of Frege, Russell, and Kripke. There are three parts to my puzzle and I will present each one of them in order. They are: (1) the existence of semantic indeterminacy, (2) that a subject can have a belief-like attitude without determinate content, and (3) the traditional analysis of belief that relates the subject to a proposition. I will not, in this paper, advocate for any one solution to this puzzle. [At least not in this draft. I want to see how well I can motivate the problem before attempting a resolution.]

1.1 Semantic indeterminacy, examples

Case 1: The pope is a bachelor. Fred and Bill are buddies who frequently engage in friendly debate. One day, Fred declares that “the pope is a bachelor.” Bill hears Fred and is outraged. “The pope is not a bachelor!”, Bill exclaims. Fred replies and insists that he is. The two of them enter into heated debate over whether the pope is a bachelor. Both participants, let’s say, have full knowledge of the pope’s marital status and his eligibility for marriage, and both of them intend to use the term “bachelor” in the same way that it is commonly used in English.

It is tempting to say that Fred and Bill’s debate is not *really* a substantial one. Their dispute is “merely verbal.” That is to say that, whether or not the pope counts as a ‘bachelor’ depends on what the term “bachelor” means. On one definition, it means *unmarried man*, in which case the pope counts; on another—equally viable—definition, it means *man that is eligible for marriage*, in which case the pope doesn’t count.

In his *Writing the Book of the World*, Ted Sider endeavours to characterize the phenomenon of nonsubstantial questions, which includes the question at issue, “is the pope a bachelor?”. According to Sider’s account, questions like this one are *nonsubstantial* when (i) some of their key expressions could have had different meanings (that are consistent with the ‘semantic purpose’ of the expression) and (ii) the answer to the question turns on which of the candidate meanings we select. Applying this to our example, the question “is the pope a bachelor?” is nonsubstantial because (i) the term “bachelor” has multiple candidate meanings that could have been meant by the term (for all intents and purposes of the term), and (ii) on *some* of those meanings, “the pope is a bachelor” will express a true proposition; on other meanings, it will not.

This is not to say that the nonsubstantial question “is the pope a bachelor?” must lack an answer. On Sider’s view, the meaning of an expression is determined by (a) our linguistic community’s usage of the term, (b) the interpretation of that usage that maximizes truth for our community’s pronouncements which use that term (c) constrained by *reference magnetism*—the idea that *natural properties*, that ‘*carve nature at its joints*’, make for the most eligible interpretations of a term or predicate. The term “bachelor” may not carve nature at its joints very well, but it’s always possible to settle the question by pointing to our history of usage of the term to privilege one content over another.

Nonetheless, there will still be some nonsubstantial questions that lack answers. This will occur when neither the history of usage of a crucial term nor reference magnetism succeed in singling out a unique content. I want to assume, for the sake of argument, that Fred and Bill’s debate is like this. That is,

I want to assume that Fred and Bill's debate does not have a determinate resolution; it is not the case that one of them is right and the other is wrong. The reason why it does not have a resolution is because both *<unmarried man>* and *<man eligible for marriage>* are equally viable candidates for the meaning of "bachelor." This assumption is an empirical assumption about common usage of "bachelor" in English; I'm assuming that our usage is generally consistent with both meanings. (Perhaps I'm wrong. If you disagree, then that's probably because you side with either Fred or Bill. In that case, please replace the example with another example of a debate that you regard as 'merely linguistic' and lacking a determinate answer.)

I want to make another controversial stipulation about this case. It is natural to expect that both Fred and Bill have one of the candidate meanings *in mind* when they assert their position. Surely Fred *intends* "bachelor" to mean *unmarried man* and Bill *intends* "bachelor" to mean *man eligible for marriage*. But I want to stipulate that this is not the case. I want to stipulate, for the sake of argument, that both Fred and Bill intend to use the term "bachelor" with whatever common meaning it has in public language. Neither of them has explicitly entertained the thought that "bachelor" has any particular precisification.

Perhaps this last assumption will sound implausible. For *surely* if Fred is so stubborn in insisting that "the pope is a bachelor", he must have some precisification of "bachelor" *in mind*. But if you think it's implausible to assume that he doesn't, then don't worry—I have two more cases to present. I expect that the temptation to think that Fred *must* have a precisification in mind is tied to the fact that, as a native English speaker, he has the conceptual resources to precisify his use of the term if he wanted to. But in my next two examples, the subject's ability to precisify their terms will not be present. Hence, in the next two examples, the claim that the subject doesn't have a precise content in mind will be nearly irresistible.

The next example comes from Hartry Field's paper "Theory change and the indeterminacy of reference." In this paper, Field is concerned to argue there are plausible historical cases of *referential indeterminacy* and so, very plausibly, some of our current terms are referentially indeterminate. Here is his case.

Case 2: Newtonian Mass. Isaac Newton's physics uses a single term "mass." Newton and his followers treated this term as univocal, and accepted many sentences involving it; e.g. "mass is conserved in all interactions." Given the advent of the theory of relativity, we can now see, from our perspective, that there is no unique property that corresponds to Newtonian 'mass.' Rather, the Newtonian term 'mass' approximately corresponds to two distinct physical quantities, *rest mass* and *relativistic mass*.

Field forcefully argues that we should not think of Newton's term "mass" as simply denotationless. If we do that, then we would be forced to assign the intuitively incorrect truth values to many of Newton's pronouncements. For instance, if Newton says "mass is F", and this is true whether we assign *rest mass* or *relativistic mass* to "mass", then intuitively we should think of his pronouncement as true. Field proposes instead that we should think of Newton's term "mass" as *referentially indeterminate*. There is no fact of the matter as to whether *Newton's word 'mass' denoted rest mass* or *Newton's word 'mass' denoted relativistic mass*. In his paper, Field develops a supervaluationist truth theory for sentences with referentially indeterminate terms. Basically, a referentially indeterminate sentence like "m is F" is true if and only if m_1 is F and m_2 is F and ... and m_n is F for each m_1, \dots, m_n that "m" partially denotes. "m is F" is

false if and only if m_1 is not F and m_2 is not F and ... and m_n is not F . “ m is F ” is otherwise truth-value-less. Naturally, Field develops the truth theory to handle sentences with greater complexity, but the higher orders of complexity need not concern us here. Suffice to say that a semantically indeterminate sentence can be true (or false) if it is true (or false) under the selection of each of its candidate contents.

Consider one last example.

Case 3: Slow-switching Twin Earth. Oscar is a resident of Earth. As such, Oscar has a long history of interacting with the substance with chemical composition H₂O, colloquially known as “water.” Owing to his history of causal interaction, let’s say that his usage of the term “water” refers to the substance H₂O. One day, over night while Oscar is asleep, he is abducted by humanoid aliens and transported to Twin Earth. Twin Earth appears on the surface to be indistinguishable from Earth, but instead of containing the substance H₂O, its lakes, rivers, ponds, and drinking fountains are filled with clear, odourless liquid with chemical composition XYZ. After residing on Twin Earth for some time, Oscar builds a causal rapport with XYZ (and loses his rapport with H₂O). Eventually his uses of the term “water” will begin to refer to XYZ rather than H₂O. (Although all of this is unknown to him; as far as he can tell, the new ‘water’ is indistinguishable from the old.)

As I understand this case, the switch between the two contents will not be sudden. As Oscar’s previous causal history with H₂O becomes more and more distant, and he gains more and more history with XYZ, the contents of his ‘water’ utterances will gradually change from one to the other. As such, there will be a period of transition in which his ‘water’ utterances will not determinately refer H₂O or determinately refer to XYZ. They will partially refer to both.

1.2 Semantic indeterminacy characterized

All three of the cases just described are instances of the phenomena known as *semantic indeterminacy*. Here is how to define it:

Semantic indeterminacy. A sentence S is semantically indeterminate if and only if a token of S (used in a context C) does not express a unique proposition P ; rather, there are multiple distinct propositions P_1, \dots, P_n such that P_1, \dots, P_n are all ‘partially expressed’ by the use of S in C (that is, there is no fact of the matter as to which of P_1, \dots, P_n is uniquely expressed by S).

It is important to note that, as I defined it, semantic indeterminacy is a property of sentence tokens taken in context. This distinguishes it from the more mundane phenomenon of *ambiguity*, which is a feature of sentence *types*. The sentence “I am near the bank” is ambiguous in the straightforward sense that it could be used to mean <I am near the river bank> or <I am near the financial institution>. But which of these propositions is expressed by a given token of the sentence will usually be disambiguated by context (for example, by how the speaker makes their referential intentions manifest). Semantic indeterminacy is more extreme. Unlike ambiguity, semantic indeterminacy will not typically get resolved by the semantic mechanisms typical of Kaplanian character (functions from contexts to contents) or by the speaker’s referential intentions. A semantically indeterminate sentence will retain its indeterminacy even after all of the relevant contextual factors and referential intentions have done their job to narrow down its content.

Another important feature of my definition is that semantic determinacy is characterized as a feature of *sentences* and *not the contents themselves*. Semantic indeterminacy occurs when a sentence (tokened in context) has multiple distinct propositional contents. It is *not*, on my definition, a case of *semantic* indeterminacy when one sentence expresses a unique content, which in turn suffers from some other sort of indeterminacy. For instance, let's suppose that we can make sense of the idea that there is 'vagueness in the world', and that the property of *being a hominid*, when viewed diachronically, is an instance of a vague property. In that case, when someone utters "my great-great-great-...-great-grandfather was a hominid" they would express a unique proposition (namely, <S's great-great-great-...-grandfather was a hominid>), and this proposition would be indeterminate in a sense, owing to the metaphysical vagueness of *being a hominid*. But it is not a case of *semantic indeterminacy*, on my definition, because there is only one proposition that is expressed. I do not, in this essay, want to take a stand on whether cases like this are possible, or whether there is 'vagueness-in-the-world' (or what that consists in).¹ My only aim is to point to the existence of indeterminacy of the distinctively semantic kind, which is characterized by a one-to-many sentence-content relationship.

My interest in this section was to motivate the claim that this phenomena *exists*. Hopefully my three examples have made that plausible. I, for one, find the existence of semantic indeterminacy to be eminently plausible *prima facie*. Surely our natural language, with all of its flaws, will be unlikely to have perfectly determinate content, all of the time.

By saying this, I do not take myself to be advocating for anything extreme. I am *not* here advocating for anything as radical as W.V.O. Quine's thesis of radical indeterminacy of translation, where indeterminacy is the rule, rather than the exception. Quine was motivated to this extreme view because he could not see any facts (which for him, meant strictly behavioural or physical facts) that could narrow the space of possible contents for ordinary sentences. My thought, rather, is that in most typical cases, our sentences and expressions will have a unique, determinate meaning. I'm not sure how exactly that happens—or what kind of facts underly content-determination—but nonetheless there must be *some way* in which our words get associated with contents. And however this happens, it seems plausible to think that it will sometimes break down, and fail to uniquely determine a content for some expressions. This is plausible, I think, because we have apparent examples where it does happen.

Why did I use three examples, whereas I could have made my point with one? Because I wanted to make plausible the idea that semantic indeterminacy can arise under a variety of different metasemantic assumptions. When I discussed Case#1, the assumed metasemantic background was Sider's *interpretationalism* with *Lewisian reference magnetism*. On this view, the key ingredients to metasemantics are (i) the usages of the term in the community of speakers, (ii) an interpretation that maximizes truth for these usages, and (iii) a constraint on this interpretation to maximize the naturalness of candidate contents. In this case, we assume that no candidates for 'bachelor' "carve at the joints", and the history of our usages of "bachelor" would not tell one way or another as whether it should be interpreted to include the pope; hence the semantic indeterminacy. As for Field's example, *Newtonian mass*, Field does not himself advocate any particular metasemantic view in that paper. However, during that period of his intellectual development, Field expressed his affinity for causal theories of reference.

¹ Arguably, the existence of semantic indeterminacy entails the existence of metaphysical indeterminacy. Taylor, David E., and Alexis Burgess. "What in the World Is Semantic Indeterminacy?." *Analytic Philosophy* 56, no. 4 (2015): 298-317.

Under the assumption of a causal theory, perhaps the semantic indeterminacy of Newton's term "mass" could be cashed out as Newton's usages of the term being *causally overdetermined* by both rest mass and relativistic mass. (Although this explanation strikes me as *very sketchy*.) Alternatively, if we keep Sider's metasemantic perspective, the semantic indeterminacy will be explained by (i) a roughly equal amount of Newton's "mass"-utterances coming out as true on either interpretation, and (ii) rest mass and relativistic mass being *equally highly joint-carving* (i.e. natural). Lastly, the slow-switching twin earth example assumes a causal-interactionist metasemantic background. On such a view, semantic indeterminacy can be cashed out as the equal amount of causal interaction Oscar has had with both H₂O and XYZ in his recent history. In conclusion, the semantic indeterminacy thesis is not tied to any one picture of content-determination; it is explicable by a variety of views.

2. *Belief*

Let's continue to tell the stories.

Case 1: The pope is a bachelor. Fred is utterly insistent that the pope is a bachelor. Whenever asked, he responds in an assertive tone of voice, "the pope is a bachelor." He does not regard the term "bachelor" as indeterminate (although he is wrong about this, according to our assumptions), and he intends for his usage of the term to share the same meanings as it usually has in public discourse. If we may speak of *mental representations*, let's say that Fred has, in his mind, a representation of the pope that is semantically equivalent to the public sentence "the pope is a bachelor." Moreover he positively assents to this representation and regards it as true. This representation of Fred's can also be explanatory of his behaviour. Say, for whatever reason, that Fred needs to take a survey of all the bachelors, Fred would invariably include the pope in his survey.

Case 2: Newtonian mass. In many ways, Isaac Newton's relation to "mass" shares many similarities with the fictional Fred's relation to "bachelor." He also utters many pronouncements concerning 'mass' that sound just like assertions, such as "mass is conserved in all interactions." We can also presume that Newton had an inner mental representation to the effect of "mass is conserved in all interactions", which he positively assented to and regarded as true. This mental representation was manifest in his behaviour: for example, he would use this representation to make real world predictions about the behaviour of physical objects. There is however one difference between Isaac Newton and Fred: namely, that unlike Fred, the historical Newton was the creative author of the concept of Newtonian mass. He was not deferring to public usage of the term. Newton thus had some authorship over the mechanisms that would determine the content of his term 'mass.' Unfortunately for him, they didn't determine a unique content.

Case 3: Oscar. Imagine Oscar during the period of semantic transition. His 'water' utterances do not determinately refer to H₂O, nor do they refer to XYZ; they partially refer to both. Imagine Oscar running his bath. A clear liquid flows out of the facet and fills his tub. He dips his toe in and silently says to himself "this water is too hot; better wait another 10 minutes." He waits another ten minutes before fully emerging himself in the bath. What goes on in Oscar's head (what's accessible to him from his point of view) would be indistinguishable to him whether the tub was filled with H₂O or XYZ.

Here is the central claim that I want to make in this section. Under these circumstances, it seems nearly irresistible to think of these subjects as having a belief. Indeed, the most natural conclusion is that *Fred believes that the pope is a bachelor*, *Newton believes that mass is conserved in all interactions*, and *Oscar believes that the water in his bath is too hot*. The justification for these claims is that (i) the subjects assertively utter these things,² (ii) the subjects represent things in a way that's semantically equivalent to how the relevant belief-ascription complement clauses represent things, and (iii) their representations explain their behaviour just as beliefs explain behaviour in a belief-desire psychology.

3. *The orthodox view on propositional attitudes*

If we put the first two pieces of the puzzle together, we get apparent examples of subjects that have beliefs without determinate content. The third piece of the puzzle, then, is that this result clashes with the orthodox view of belief. According to the orthodox view, what it is to have a belief is to stand in a relation *to a propositional content*. This conflicts with the foregoing examples because the subjects in them seem to have beliefs, and yet they aren't related to *unique propositional contents*. In this section I will try to make this conflict more precise. In order to do so, I need to divide the orthodox view of belief into its two components—its semantic thesis and its metaphysical thesis—and discuss each of them in turn.

First, the semantic thesis of the orthodox view concerns *belief ascriptions*, that is, sentences of the form “S believes that P.” According to the orthodox view of belief ascriptions, the proper way to carve up the expressions in these sentences into their appropriate syntactic categories is given by “S” / “believes” / “that P.” In belief ascriptions, “S” is a singular term that refers to a subject, “believes” functions as a two-place predicate expressing a relation, and “that P” forms a complex singular term that refers to a special kind of object called a *proposition*. The truth conditions of belief-ascriptions are then explained by their compositional structure by assimilating them into the broader class of atomic sentences with syntactic structure “aRb.” Namely, “S believes that P” is true if and only if the pair <S, that P> is a member of the extension of “believes.”

The metaphysical thesis of the orthodox view concerns the sort of *facts* that get expressed by belief ascriptions. In particular, it holds that propositions really exist—there is such an entity that is the primary bearer of truth and falsity, the meaning of sentences, the content of attitudes, etc., and it holds that to have a belief is to have a certain kind of attitude towards a proposition. Belief is thus essentially relational, and its relata are subjects and propositions.

The first two sections of this paper presented us with subjects that seemed to hold beliefs without determinate content. This, I take it, is primarily in conflict with the *metaphysical thesis* of the orthodox view. That is because, according to the orthodox view, if Fred believes that the pope is a bachelor, then Fred must stand in an attitudinal relation to the proposition <the pope is a bachelor>. The problem with this is that there is no such proposition; remember, the sentence “the pope is a bachelor” does not uniquely express a proposition.

² Compare this claim to the assent-belief link in Kripke's “A puzzle about belief”: “If a competent speaker x of a language L sincerely and reflectively assents to (or assertively utters) an indexical-free sentence s of L, and if p is a proper English translation of s, then x satisfies [y believes that p].” This rendition is from page 61, Scott Soames, “Direct reference, propositional attitudes, and semantic content.” *Philosophical Topics* 15, no. 1 (1987): 47-87.

It is a less straightforward as to how my observations relate to the orthodox view's semantic thesis. It does not conflict with the syntactic analysis of belief reports as a binary relation, or the claim that *that*-clauses function as referring terms. The following three claims are *not in genuine conflict*:

- (1) The sentence "Fred believes that the pope is a bachelor" is true.
- (2) "that the pope is a bachelor" functions as a singular term.
- (3) The sentence "the pope is a bachelor" is semantically indeterminate.

The reason these are not in conflict is because (1) can be explained in light of (2) and (3) by appealing to supervaluationist semantics. In that case, the sentence "Fred believes that the pope is a bachelor" is true if and only if Fred believes X_1 and Fred believes X_2 and ... Fred believes X_n for each candidate referent ($X_1 \dots X_n$) of the expression "that the pope is a bachelor." This shows that certain aspects of the orthodox view's semantic thesis are not question.

However, even though we can appeal to supervaluationist techniques to explain the truth of "Fred believes that the pope is a bachelor", the *metaphysics* of the orthodox view still faces a major problem. Recall that, when we were discussing *Newtonian mass*, we noted that we could use supervaluationist techniques to assign a truth value to many of Newton's utterances using "mass." In particular, his utterance "mass is conserved under all interactions" comes out *true* because it is true under an assignment of each candidate referent to the term "mass." But despite the possibility of assigning *truth* to this utterance, there is still a major problem with the *metaphysics* of Newtonian mass: namely, it doesn't exist—in reality, there is no such thing. Similarly, despite the fact that we can assign truth to the sentence "Fred believes that the pope is a bachelor", there is still a major problem for the orthodox view's metaphysical claim that this reports a relation between Fred and the proposition <the pope is a bachelor>: namely, there is no such proposition. So the availability of supervaluationist techniques to ascribe truth to belief reports does nothing to alleviate the conflict with the orthodox metaphysics of belief.

Solutions

Is there a genuine problem here? If so, then (at least) one of the following three claims must be false (corresponding to each of the three previous sections):

- (1) There exist semantically indeterminate sentences.
- (2) It is possible for a subject to have a *belief* that corresponds to their assent to a semantically indeterminate sentence.
- (3) Belief is a relation to (unique) propositional content.

But which one?

Here is a quick list of solutions that I am aware of.

I. Views that deny (1)

II. Views that deny (2)

II.a Robert Matthews, *The Measure of Mind*

II.b Gareth Evans, *The Varieties of Reference*

III. Views that deny (3)

III.a In his paper “Mental Representation”, Hartry Field analyzes belief as “X believes that p if and only if there is a sentence S such that X believes* S and S means that p ”³ where S is a sentence in one’s internal system of representation. This view posits a *belief** relation that relates a subject to a *sentence in mentalese*. According to this view, on the *sentential* sense of “believes”, (2) is true and (3) is false. (Of course, Field isn’t the only thinker to propose the existence of a sentential belief relation. So did Davidson and Fodor, and I’m sure there are many others.) I’m unaware of any paper that has used the puzzle presented here as a consideration in favour of mental representation thesis and the sentential belief relation.

III.b In “Is belief a propositional attitude?”, Ray Buchanan uses considerations that are very similar to mine in order to argue that belief is not (always) a relation to a proposition. (Buchanan gives a separate treatment to the Fregean and Russellian view of propositions. I regard this as unnecessary, since the existence of semantic indeterminacy is independent of which view of propositions one takes. [I could defend this claim in a later version of this paper.]) His solution is that, in his problem cases, belief is a relation to *a property of propositions*. Suppose that the sentence ‘S’ has $P_1 \dots P_n$ as its candidate meanings, and a subject accepts ‘S’ as a truth. Then the subject bears the belief relation to a *property* that has $\{P_1 \dots P_n\}$ as its extension.

³ Field, “Mental Representation”, page 12.