X-Sender: jpryor@imap.fas.harvard.edu (Unverified)
Date: Sun, 25 Feb 2001 12:26:06 -0500
To: Sara Granovetter <granovet@fas.harvard.edu>,
Angharad Laing <alaing@fas.harvard.edu>,
Ben Dickson <bdickson@fas.harvard.edu>,
Alexis Burgess <aburgess@fas.harvard.edu>,
Ben Jarvis <bjarvis@fas.harvard.edu>,
Aryeh Weinstein <weinst@fas.harvard.edu>,
Vanessa Browder <browder@fas.harvard.edu>,
Marc Wallenstein <mwallens@fas.harvard.edu>,
Henry Rich <hrich@fas.harvard.edu>,
Graham O'Donoghue <rodonogh@fas.harvard.edu>
From: James Pryor <jpryor@fas.harvard.edu>
Subject: What's going on in the Quine article?

A number of you have expressed bewilderment about what's going on in the Quine article. You don't really see what he's trying to do.

Part of the reason for this is that the relation between this Quine article and the Russell stuff we were looking at earlier takes some work to uncover. I don't want to give that away, I want us to discuss it and figure it out together.

Another part of the reason is that Quine is a very spare, economical writer. This makes it harder to see what he's trying to do. This can't be helped. It's true that Quine could have written in a way that's more reader-friendly; but as philosophical writers go there are many who are much worse. Quine is actually a quite good writer, compared with others. So you have to learn how to accommodate difficult writing. I have some guidelines I give to my classes about How To Read a Philosophical Paper. It's on the web at <http://www.people.fas.harvard.edu/~jpryor/general/reading.html>. Some of that advice may already be familiar to you, but it might help to be reminded of it.

I want you to learn for yourselves how to pick apart a paper, how to figure out what is going on in each part. You won't learn how to do this if I tell you in advance what the structure of each paper is, what its main points are, and so on. And you have to accept that picking apart and understanding a philosophical paper takes time and work. It isn't like reading a newspaper article. It's more like doing a crossword puzzle or a jigsaw puzzle.

But if you're not making much progress on your own, that's no reason for despair. That's what our group discussions are for. We can try to figure the paper out together. We can try out ideas: "Maybe he's trying to do this... No, that doesn't seem right because... Does he assume this? I think so, because..."

It's perfectly okay for you to propose discussion questions of that form. Of course, don't just email me questions like "What is Quine up to?" We're going to use your questions to get discussion started. So they should be questions that we can all understand, and that we can all engage with and start discussing. If you said to the group "What is Quine up to?" the group will ponder this silently and scratch its heads. No good! Better to say something like "I'm trying to figure out what Quine is up to. Do you guys think his point is that there are two kinds of belief (and two kinds of wanting, and so on), and that neither can be reduced to the other?" That's more likely to start a discussion.

Since so many of you have said you can't figure out what Quine is up to, though, this one time only I will walk you through the article a little bit. Watch how we do this. This is the kind of thing you need to do yourself, with every article you read.

In the first few paragraphs it's hard to see where Quine is going in the article, or what his main point will be. But it's clear that he's introducing a distinction between two kinds of hunting, and two kinds of wanting. Then halfway down p. 186 he says that this distinction applies to belief as well.

On p. 187 we get the first clue of why Quine thinks any of this is important. He says "However, the suggested formulation of the relational senses...all involve

quantifying into a propositional-attitude idiom from outside. This is a dubious business, as may be seen from the following example..."

Before we go any further we should try to figure out what Quine means by "a propositional attitude idiom" and what he means by "quantifying into" such an idiom. The examples Quine cites should make it clear that he's counting talk about striving, wishing, and believing as "propositional attitude idioms." And they should also make it clear that claims of the form "Exists an x: Witold wishes that...x..." count as "quantifying into" the wishing idiom, and claims of the form "Witold wishes that: (Exists an x:...x...)" doesn't count as quantifying into the wishing idiom. So you only have "quantifying in" when the quantifier "Exists an x" is outside the "wishing, " but the variable "x" occurs inside the report of what's wished.

(Remark about "propositional attitudes." These are mental states that you stand in to propositions. For example, believing that P is standing in a certain attitude to the proposition that P. Similarly with the other propositional attitudes Quine cites: striving that Ernest finds x, wishing that I have x. In both cases we have an attitude towards a certain proposition--ignoring for the moment any peculiarities due to the "x." On the other hand, mental states like hunting x and wanting x are attitudes towards THINGS, not attitudes towards propositions. You hunt lions, you don't hunt that P. So those don't count as propositional attitudes. But Quine thinks that when you have some claim of the "hunting x" variety, you can transform it into some propositional attitude claim. That's why instead of "Ernest is hunting x" he switches over to talking about "Ernest strives that Ernest finds x." Here we have a propositional attitude "Ernest strives that such-and-such." This may be more than you could have figured out on your own from Quine's article, if you had never encountered talk of propositional attitudes before. But you should at least be able to see from his examples what sorts of things he is counting as "propositional attitudes," and what sorts of things he is counting as "quantifying into" an idiom.)

Okay, so back to the spot on p. 187 where Quine says "However, the suggested formulation of the relational senses...all involve quantifying into a propositional-attitude idiom from outside. This is a dubious business, as may be seen from the following example..."

Quine then goes on to give the example of Ralph and Ortcutt. He says:

> "Can we say of this MAN (Bernard J. Ortcutt, to give him a name), that Ralph believes him to be a spy? If so, we find ourselves accepting a conjunction of the type:
>
> (11) w sincerely denies '...' w believes that ...
>
> as true, with one and the same sentence in both blanks. For, Ralph is ready enough to say, in all sincerity, 'Bernard J. Ortcutt is no spy.' If, on the other hand..."

Okay, at that point we should stop, because we see from the "If on the other hand..." that Quine is going on to make a new point. It looks like he's starting to construct some sort of dilemma. That's an argument of the form "If such-and-such then we end up here, and if on the other hand so-and-so, then we end up over there." Before we proceed any further in the article then, we should try to figure out what Quine has just done, and what kind of dilemma he might be starting to construct.

In the little bit of text we're looking at, Quine seems to be saying that, if we use the relational sense of "believing that so-and-so is a spy," then we should count Ralph as believing Ortcutt to be a spy. (There is an x, namely Ortcutt, that Ralph believes to be a spy.) But this is odd, because Ralph might in all sincerity deny the sentence "Ortcutt is a spy." When he thinks of the guy as a spy, he doesn't connect him with the name "Ortcutt." So the relational sense of believing that so-and-so is a spy seems to lead to the odd result that you could have claims of the following sort being true:

(11*) Ralph sincerely denies "Ortcutt is a spy," yet Ralph believes Ortcutt to be a spy.

Quine thinks that's odd or somehow problematic.

Anyway, that's a first stab at what's going on in that bit of text. Maybe we'll want

to come back later and adjust our understanding of it, but as a first stab it seems OK. So let's see what happens next. Quine continues:

> "If, on the other hand, with a view to disallowing situations of the type (11), we rule simultaneously that
> (12) Ralph believes that the man in the brown hat is a spy
> (13) Ralph does not believe that the man seen at the beach is a spy
> then we cease to affirm any relationship between Ralph and any man at all..."

This is puzzling. It probably won't be clear right away what's going on here. But we can chip away at it a bit. It looked like Quine was constructing a dilemma, and the first part was about what happened using the relational sense of "believing that so-and-so is a spy." So it's natural to assume that this second part will be about what happens using the notional sense of "believing that so-and-so is a spy." And that does seem to fit what Quine is saying here, though it's not 100% clear that it's what's going on. Another thing you might notice is that in the previous passage, Quine talked about "Can we say of this MAN (Bernard J. Ortcutt, to give him a name) that Ralph believes him to be a spy?" There we're using a name "Ortcutt" and the pronoun "him." Here we're using descriptions "the man in the brown hat" and "the man seen at the beach." Is that difference significant here? Or not? Can't tell yet.

(See how it's like doing a crossword puzzle?)

Let's see how Quine continues, it might clear some things up. He writes:

> "If, on the other hand, with a view to disallowing situations of the type (11), we rule simultaneously that
> (12) Ralph believes that the man in the brown hat is a spy
> (13) Ralph does not believe that the man seen at the beach is a spy
> then we cease to affirm any relationship between Ralph and any man at all. Both of the component 'that'-clauses are indeed about the man Ortcutt; but the 'that' must be viewed in (12) and (13) as sealing those clauses off, thereby rendering (12) and (13) compatible because not, as wholes, about Ortcutt at all. It then becomes improper to quantify as in (7)...
> [looking back we see that (7) was: (7) Exists an x: Ralph believes that x is a spy.]
> ...'believes that' becomes, in a word, referentially opaque."

Okay, in a way this just makes things more confusing but it helps a little too. What does this "referentially opaque" mean? You may never have encountered that expression before. Well, from the way Quine uses it, it's clear that it has something to do with "quantifying into" belief-talk, as in "Exists an x: Ralph believes that x is a spy." Quine says here that when it's improper to "quantify into" the belief-talk, the belief-talk is "referentially opaque." So that's how we'll understand talk about things being "referentially opaque."

It may be hard to see what some of the individual sentences here are saying. For instance, the sentence beginning "Both of the component 'that'-clauses..." is confusing. But looking at the passage as a whole, and considering it in the context of a dilemma whose other horn concerned the relational sense of "believing that so-and-so is a spy," the following general picture emerges:

> If you're dealing with the relational sense of "believing that so-and-so is a spy" then you're allowed to quantify into the belief-talk, but then you get the odd result that things like
> (11*) Ralph sincerely denies "Ortcutt is a spy," yet Ralph believes Ortcutt to be a spy.
> might come out true. If on the other hand, you're dealing with the notional sense of "believing that so-and-so is a spy," that may enable you to avoid odd results like (11*)--instead we'd express things by saying (12) and (13). But with the notional sense, there's some sense in which Ralph's beliefs aren't relations to the man Ortcutt. (Is this because they're relations to propositions?) And with the notional sense

> Quine says it is "improper" to quantify into the belief talk. (Or as Quine also puts it, the notional belief-talk is "referentially opaque.")

There may still be parts of what Quine has said here which are unclear, but that general picture seems OK. So let's see how things continue from there. He writes:

> "No question arises over (8)...
> [looking back we see that (8) was: (8) Ralph believes that: (Exists x: x is a spy).]
> ...it exhibits only a quantification WITHIN the "believes that" context, not a quantification INTO it. What goes by the board, when we rule (12) and (13) both true, is just (7). Yet we are scarcely prepared to sacrifice the relational construction "There is someone whom Ralph believes to be a spy" which (7) as against (8) was supposed to reproduce."

Well, it takes some work, but piecing this together with the general picture we've extracted so far, the idea seems to be:

> If you go with the relational sense of belief, then you get odd results like:
> (11*) Ralph sincerely denies "Ortcutt is a spy," yet Ralph believes Ortcutt to be a spy.
> So if you want to avoid those odd results you might go with the notional sense of belief, instead. Then we can just say:
> (12) Ralph believes that the man in the brown hat is a spy
> (13) Ralph does not believe that the man seen at the beach is a spy
> and there's no confusion about Ralph's beliefs. Ralph would accept the sentence "The man in the brown hat is a spy" and he would refuse to accept the sentence "The man I saw at the beach is a spy." But the problem is that the notional sense doesn't allow us to quantify in at all, it doesn't let us say things like "There is someone whom Ralph believes to be a spy." So each sense seems somehow inadequate.

So what? Is Quine thinking that we're only allowed to have one of the senses? Then maybe there would be a problem here. Maybe he is assuming that. He goes on to say "The obvious next move is to try to make the best of our dilemma by distinguishing two senses of belief..." That does seem a reasonable move to make. That's the end of the first section.

The second section begins "But there is a more suggestive treatment..." Then Quine goes on to talk about how we might just take one of the notions of belief (the notional or "referentially opaque" notion) and extend it in a way to make it do all the work.

So from what we've seen so far, the general theme of the article seems to be: "There are two senses of belief, we can't just ditch one of them and keep the other, so at first it looks like we need both. But let's try to see if we can give a unified analysis, so that what appear on the surface to be two kinds of belief can really be seen to be instances of a single common underlying relation." And that is pretty much what Quine is up to.

Now that we've figured that out, it's really helpful. We'll come back to section two. Now that we've figured out what Quine's main project is, let's first try to get an overview of the general structure of the rest of the paper. (As I say in the guidelines on reading philosophy papers, it's important to figure out what the main project or point of the article is, and what its general structure is, before you concentrate on the details. Everything we've done so far was just trying to unearth what Quine's main project was.)

Skimming ahead to section three, at the start of that section Quine writes:

> "Striving and wishing, like believing, are propositional attitudes and referentially opaque. (3) and (4)...

> [that is, (3) Exists an x: (x is a lion & Ernest strives that Ernest finds x)
> and (4) Exists an x: (x is a sloop & I wish that I have x)]
> ...are objectionable in the same way as (7)...
> [(7) Exists an x: Ralph believes that x is a spy.]
> ...and our recent treatment of belief can be repeated for these propositional attitudes. Thus, just as (7) gave way to (17)...
> [(17) Exists an x: Ralph believes z(z is a spy) of x]
> ...so (3) and (4) give way to:
> (24) Exists an x: (x is a lion & Ernest strives z(Ernest finds z) of x)
> (25) Exists an x: (x is a sloop & I wish z(I have z) of x)."

So the general point of section three seems to be to apply the account developed for belief, in section two, to other propositional attitudes, like striving and wishing. (There is also some discussion of whether it's OK to stick with notions like hunting x, which aren't relations to propositions, or whether it's better to analyze them in terms of the account Quine has developed for propositional attitudes.)

Skimming ahead to section four, Quine writes there:

> "There are good reasons for being discontent with an analysis that leaves us with propositions, attributes, and the rest of the intensions. Intensions are less economical than extensions (truth values, classes, relations), in that they are more narrowly individuated. The principle of their individuation, moreover, is obscure..."

So the general point of this section seems to be that the analysis Quine has given us so far has made us of some things he calls "intensions." And that for various reasons (he cites two in this little passage) he thinks it would be better to have an analysis that uses some things he calls "extensions" instead. That seems to be the project of section four. (In fact, he ends up giving an account not in terms of extensions but rather in terms of sentences. But that's a detail, we're not paying attention to the details right now.)

If you haven't encountered the notions of "intentions" and "extensions" before, then this section will be hard to understand. Also section two may be hard to understand. You should be able to piece together a little bit of a picture of what the difference is between extensions and intensions, from the things Quine says. As I said, it's like a crossword puzzle or a jigsaw puzzle... But this email is already pretty long, so I'll just tell you what the difference is between intensions and extensions, and what that difference has to do with the details of sections 2 through 4 of the article.

Everything up until now you really would have been able to uncover for yourselves, by digging through the article, putting the pieces together, and discussing it with each other! Really you would. And it's really only the stuff that we've discussed so far that is relevant for our discussion of Russell and descriptions and names. What I'm about to explain is important to Quine, and his overall project, but it gets into details that we don't really need to worry about, for our purposes. Still, you asked, so...

INTENSIONS AND EXTENSIONS: a rough introduction

First, we want to distinguish different parts of language: we have sentences, names, descriptions, predicates, and so on. A predicate is a bit of language like "runs quickly" or "is a key to Jim's apartment" or "is in Jim's pocket on Sunday morning at 11 AM."

Then we have the worldly counterparts of those. The worldly counterpart of a name is some thing in the world. So we might have the name "George" and its worldly counterpart is the person George. We say that the name REFERS TO that person.

With sentences and predicates the story about the worldly counterpart is more

complicated. First of all, when George runs quickly we say that George SATISFIES THE PREDICATE "runs quickly." If Kiki is one of the keys to my apartment (I give all my keys names), then Kiki satisfies the predicate "is a key to Jim's apartment." Basically, a sentence like "Kiki is a key to Jim's apartment" will be true just in case the name "Kiki" refers to some object, and that object satisfies the predicate "is a key to Jim's apartment."

So referring is a relation that names stand in to objects, and satisfying is a relation that objects stand in to predicates, which are bits of language. The object is the worldly, extra-linguistic counterpart of the name. What is the worldly counterpart of the predicate?

Well, one of the things that's a worldly counterpart of a predicate is the predicate's EXTENSION. The extension of a predicate is the set of all the objects that satisfy that predicate. So if Kiki and Jojo are two keys to my apartment, and they're the only two keys, then the extension of "is a key to Jim's apartment" is the set consisting of Kiki and Jojo.

Suppose that, as it happens, Kiki and Jojo are also both in my pocket right now, and they're the only two things in my pocket right now. Then the extension of "is in Jim's pocket on Sunday morning at 11 AM" is also the set consisting of Kiki and Jojo. So the two predicates "is a key to Jim's apartment" and "is in Jim's pocket on Sunday morning at 11 AM" both have the same extension. (Or as we also say, they're COEXTENSIONAL.) At the level of extension, there is no difference between these predicates.

Yet we think that, intuitively, there is SOME difference between these two predicates. As it happens they have the same extension, but they MIGHT HAVE HAD different extensions. There might have been other keys to my apartment than just those two. Or one of those two might never have existed. Or I might have had some coins in my pocket on Sunday morning, in addition to the keys. Or... We capture these differences between the predicates "is a key to Jim's apartment" and "is in Jim's pocket on Sunday morning at 11 AM" by saying that the two predicates have different INTENSIONS, or express different PROPERTIES. The property of being a key to Jim's apartment is not the same property as the property of being in Jim's pocket on Sunday morning. It's POSSIBLE for things to have the one property but lack the other, even if, as it IN FACT happens, the very same things have both properties.

So the intension of a predicate is something like the property the predicate expresses, and the extension of a predicate is the set of objects that actually in fact happen to have that property. These can both be counted as extra-linguistic, worldly counterparts of the predicate.

So far I've just been talking about the intensions and extensions of 1-place predicates like "runs quickly" and "is a key to Jim's apartment." We can also talk about the intensions and extensions of 2-place predicates, like "is taller than"; but let's not worry about that for now. We also sometimes talk about the intensions and extensions of SENTENCES. The extension of a sentence is its truth-value: so all true sentences have the same extension. The intension of a sentence is the proposition it expresses. There can be pairs of sentences that have the same extension (truth-value) but which MIGHT HAVE HAD different truth-values. For instance:

Kiki is a key to Jim's apartment

and:

Kiki is in Jim's pocket in Sunday morning at 11 AM.

have the same extension (they're both true). But Kiki might not have been in my pocket on Sunday morning: it could have been the case that the first sentence was true but the second was false. In such a case the sentences have different intensions (they express different propositions).

We also sometimes talk about the extensions of names. The extension of the name "Kiki" is just the object Kiki, that the name refers to.

Well, that's an introduction to the notions of intensions and extensions. It may help

to know that Quine isn't very fond of intensions (properties, propositions, and so on). He thinks they're creatures of darkness. Extensions are just sets, and they're much more well-behaved. In part this is because if Superman = Clark Kent (and they both really exist), then any set containing Superman also contains Clark Kent, and vice versa. There is no difference between the set {Superman, Kiki, Napoleon} and the set {Clark Kent, Kiki, Napoleon}. On the other hand, it's not so clear that intensions work that way. Lois might believe the proposition that Superman flies, but fail to believe the proposition that Clark flies. She might desire to have the property of standing close to Superman, but not desire to have the property of standing close to Clark. And so on. Well, matters are not so clear here. Some philosophers think those propositions and properties are different, others think they are the same. It's murky and controversial. With extensions, it's clear that there is no difference between extensions that contain Clark and extensions that contain Superman.

This may help make it clearer why Quine says, in section 2 of his article, that you can't quantify into the names of intensions. He's using expressions like "z(z is a spy)" as a name for the property of being a spy (in other words, as a name for the intension of "is a spy"). In section two of his article, he wants to represent the relational claim, that Ralph believes of Ortcutt that he is a spy, as follows:

> Ralph stands in the believing relation to (i) the object Ortcutt, and (ii) the 1-place intension or property: being a spy (in other words, z(z is a spy)).

Here the reference to Ortcutt is not inside the name of the intension, so it is OK to quantify over it, or to replace it with another name or description that picks out the same guy.

Quine wants to represent the notional claim, that Ralph believes (as Ralph would himself put it) that Ortcutt is a spy, as follows:

> Ralph stands in the believing relation to (i) the 0-place intension or proposition: that Ortcutt is a spy.

Here the reference to Ortcutt is all inside the name of the intension, so it can't be quantified over, and we're not allowed to replace it with any old name or description that picks out the same guy.

What Quine's analysis in section 2 involves, then, is a believing relation which you can stand in to a variably-sized set of objects, and a variably-placed intension. (You could have 0 objects and a proposition, or 1 object and a 1-placed property, or 2 objects and a 2-placed property (otherwise known as a relation), and so on.) The part of the sentence that names the intension is "referentially opaque": we can't quantify into it, or substitute names. This part of the sentence is supposed to represent how Ralph thinks of the phenomena his belief is about. But the part of the sentence that is not supposed to be naming any intension, or representing how Ralph thinks of the relevant objects, allows you to quantify and substitute names freely.

That's the gist of the account in section 2. It may not look like much of an advance over the end of section 1, where we had two different senses of "belief." How is it really more informative to be told that there is just one underlying belief relation, it's just that Ralph can stand in that relation to the object Ortcutt and the intension: being a spy, or he can stand in that relation to the intension: that Ortcutt is a spy. By itself, this is not much of an advance. But Quine thinks this is an important intermediate step, because it makes it clear where we're allowed to quantify and where not, and we have a rationale for why. (Quine's rationale is that we're not allowed to quantify into names of intensions.) In section 4, Quine goes on to see if he can fiddle with the account in such a way as to get rid of all the intensions, and replace them with something else, like extensions (that doesn't work) or bits of language (Quine thinks that does work). In Quine's eyes, THAT'S what important about this article. The final account that you get in section 4.

But as I said, it's really not what's central for our purposes. What's central for our purposes is all there on the table already, once we've figured out what Quine's main project is, and what the overall structure of the paper is. As I put it above, the general theme of the article seems to be: "There are two senses of belief, we can't just ditch one of them and keep the other, so at first it looks like we need both. But let's try to see if we can give a unified analysis, so that what appear on the surface to be two kinds of belief can really be seen to be instances of a single

common underlying relation." That's what Quine is up to.

What we want to figure out is how all that relates to the stuff we've been discussing, about Russell and names.

OK, those are all the hints you get for now... :-) See you all this week.