

Nearly final Draft 18/9/

*In preparation for 'Contemporary Philosophy in Focus: Jerry Fodor', ed. Crane, T.,
Cambridge University Press,*

Comments welcome

Not for citation without permission

Gabriel Segal

Keep Making Sense

0. Preramble

"There are various things that you can usefully do when your car gets a ping in one of its cylinders; but declining to quantify over the engine is not one of them" (1987 xi)

"If the semantic and the intentional are real properties of things, it must be in virtue of their identity with (or maybe of their supervenience on) properties that are themselves neither intentional nor semantic. If aboutness is real, it must really be something else." (1987 97)

Why is it that so many undergraduates confuse reductionism with eliminativism when it comes to the mental, while even schoolchildren have no great difficulty with the idea that water can be H₂O yet still exist at the same time? Perhaps it's because our intuitive view of the mind, the one provided by our genes (Fodor 1987 129-33, Segal MoToM) is deeply dualistic (see Wellman CtoM for some evidence). On this view, if something is physical then it is automatically not mental. Anyway, the line between reduction and elimination is shadowy and dim: it can seem arbitrary where the one should end and the other begin. It's OK for water to be H₂O. But why isn't it ok for hypnotists to be magicians or caloric to be heat (ref Papineau, "Theory Dependent Terms")?

Our innately given picture of the mind comes as part of a package deal with a theory: folk psychology. As Fodor argues persuasively at the start of Fodor (1987), this theory is really very good. It works. And it's profound, too. At its core lies the idea that psychological states have both intentional and causal properties and that one can predict the latter from the former. Suppose that someone thinks that no dragon is furry and holds also that Fafnir is a dragon. What will happen if she is forced to consider the

question of whether Fafnir is furry? She will probably conclude that he isn't. What we are told about is the content of the first two beliefs. From this we can infer that they are likely to cause a particular third belief, again specified by its content. And that's just brilliant.

Folk psychology is profound, it works and as Fodor also points out, it out-performs all competitors in its domain: cognition, volition and action. Physics, chemistry, biology and the rest just aren't too helpful when it comes to predicting and explaining things like beliefs, desires and actions. To put it mildly.

But how could the mind as we intuitively conceive of it really be something else? How could intentionality be something else? 'Being' something else means either being type-identical with something otherwise describable or being supervenient on such a something (Fodor 1987, 97 quoted above). And if the latter, not just any old supervenience relation will do, I'd say. I think that Fodor would want the relation to be metaphysically necessary and to be the converse of the makes-it-the-case-that relation that obtains in familiar cases: the way the diamond's crystals hang together makes it the case that the diamond is hard etc...

I think we all struggle to specify the relevant relation: hence maybe the "maybe" in Fodor's parenthetical remark to the effect that maybe supervenience would do. Maybe makes-it-the-case-that is identity of tropes, rather than supervenience, for instance. A trope is a token of a property: the fragility of my favourite cup, considered as distinct from fragility in general or the fragility of my least favourite cup. Suppose that my favourite cup is indeed fragile and that its fragility is due to irregularities in the pattern of the crystals that make it up. The irregular pattern makes it the case that the cup is fragile.

Now it wouldn't be absurd to suggest that the fragility of the cup, the trope, just is the trope of being made of crystals in that irregular pattern. While fragility in general, the universal shared by all the fragile, cannot be identified with irregularity of pattern in general, we might think that when we get to the specifics of the cup case, there is just the one trope that is both its fragility and its irregularity of pattern.

So, the thought is, maybe we can account for making it the case in terms of trope-identity. Since the fragility of my cup just is it's having the irregular pattern, its having the pattern makes it the case that the cup is fragile. Nothing more is required for the cup to be fragile than it's having the irregular pattern.

In any event, I don't propose to fuss about the details of reduction, desirable sorts of supervenience and all that. For convenience, I'll call both makes-it-the-case-that and type-identity stories, if true, "reductions".

So, given that background, we live with a general problematic: any decent shot at providing a reduction of intentionality will very likely force us to give up some of our intuitive preconceptions. That's what typically happens when we do science, after all. Lots of folk physics is badly wrong. How much is one willing to give up? How much is one willing to pay for a reduction? Fodor (1987) starts with what might look like a pretty clear answer to those questions: the reduction has to vindicate the basic apparatus of folk psychology, the idea of items that have both intentional and causal properties such that etc. etc. etc...

But further inspection finds murk: how much of the folk psychological conception of its basic apparatus has to be vindicated and how much is negotiable?

What I plan to do is look at a couple of specific points related to that last question. In his (1987), (1990) and (1994) Fodor offers a sketch of a possible reduction of intentionality, at a specific price. If you buy the story, you have to give up certain aspects of the folk psychological notion of singular concepts. I will give some reasons for thinking that maybe we should hang on to them, instead, and, in this case, think of intentionality as what it is, rather than as something else. In part 2, I'll take issue Fodor's account of non-singular concepts and argue that, in a certain respect it, may be too close the folks' one.

1. Fodor, neo-Russell and quasi-Frege

"Psychologists have no right to assume that there are intentional states unless they can provide, or anyhow foresee providing or anyhow foresee no principled reason why someone couldn't provide naturalistic sufficient conditions for something to be *in* an intentional state." (1994, 5)

"[A]s far as I can see, of the various proposals around for a naturalistic account of content, only the informational ones have a prayer of working. So I'm going to hold onto informational semantics if I can" (1994, 6)

Fodor's favoured candidate for a reduction is informational semantics. Fodor's informational semantics includes what one might call a 'neo-Russellian' notion of informational content. On this view, the contents of propositional attitudes towards particular objects are (kind of) Russellian propositions, their components being the objects and the properties ascribed to them. The content of any singular concept, including definitely descriptive ones (Fodor departs here from the book of Great Grandpa Russell), is the object it denotes. *Mutatis mutandis* for kind concepts, such as WATER and H₂O. The idea behind that is that reduction requires intentional relations to be cashed in terms of real ones, such as causation. But as far as real relations go, there is no difference between Hesperus and Phosphorus. Whatever causal relations Hesperus enters into, Phosphorus enters into too.

Further, on Fodor's view, empty singular concepts have a sort of null content: all empty singular concepts have the same content. I think the reason that he holds this is that the alternative neo-Russellian positions are unprepossessing. These alternatives are: either (a) there aren't empty singular concepts (so it's just impossible for anyone to believe that Fafnir has no fur, no matter how hard they try) or (b) empty singular concepts refer to non-existent objects. One shudders to think.

Folk psychology is not neo-Russellian. It's 'quasi-Fregean'. On the quasi-Fregean outlook, different co-extensive concepts can have different contents. Empty concepts are possible. Lots of them are actual. And different empty concepts have different contents. And the content of general concepts cuts finer than the properties (if any) that correspond to them: being water is being H₂O, but being WATER isn't being H₂O. In other words, quasi-Fregeans believe in content that at least roughly resembles Fregean sense. Call it 'quasi-sense'.

Do be so kind as to considering the following four candidate psychological explanations:

(1) a. Oedipus wants to fly to Hesperus

b. Oedipus believes that the *USS Evening Star* is about to depart for Hesperus.

So c. Oedipus boards the *USS Evening Star*

* (2) a. Oedipus wants to fly to Hesperus

b. Oedipus believes that the *USS Evening Star* is about to depart for Phosphorous.

So c. Oedipus boards the *USS Evening Star*

(3) a. Helga believes that Gunther fought with Hagen

b. Helga believes that Gunther is the King of the Gibichungs

so c. Helga believes that the King of the Gibichungs fought with Hagen

** (4) a. Helga believes that Gunther fought with Hagen

b. Helga believes that Gunther is the King of the Gibichungs

so c. Helga believes that Fafnir fought with Hagen

What would Fodor's attitude be to (1)-(4)? Let us consider the first two first. On his view, these ascribe attitudes to just the same Russellian propositions. So if (1) offers a good explanation in terms of content, then so too does (2). (1) does, so (2) does too. A typical reaction to this, a quasi-Fregean one, would be to think that there is something right about (1) and something wrong about (2). In particular, it would seem that in (1) (a) and (b) appear fit to explain (1)(c), while (2)(a) and (b) don't appear fit to explain (2)(c). Equally, in (1), (a) and (b) would ground a prediction of (c), but in (2), not. So, if the generalisations of psychology fail to distinguish (1) and (2), we had better fix them up. That would be the standard reaction. But Fodor thinks it's time to dump the old hat and buy a new one: "*Don't fix the generalization[s], fix Oedipus*" (1994, 44, emphasis Fodor's).

Fodor's idea is that (1) and (2) both follow from perfectly good psychological generalisations, lawlike generalisations hedged by *ceteris paribus* clauses. *Ceteris paribus*, someone with the attitudes ascribed to Oedipus in either (1) or (2) (a) and (b) would indeed endeavour to board the *USS Evening Star*.

Fodor argues that "Frege cases" will lead to exceptions to the generalisation and legitimate the invocation of the *ceteris paribus* clause. By a "Frege case" Fodor seems to mean a case like that of Oedipus in the original Sophocles play. Oedipus's problem is that for some a, b and F, he does not know that a=b (although she is), he desires that Fa and he also desires (or would desire if he contemplated it) that most definitely not-Fb. Frege cases, Fodor argues, must be aberrations, otherwise one's endeavours would often be thwarted by inconvenient identities of which one was ignorant. He concludes that there must be mechanisms in play that ensure that subjects typically do know all relevant identities. Frege cases arise only when these mechanisms fail.

Fodor's argument is problematic. The sort of Frege cases that are relevant to the generalisations at issue are not merely those where an unknown identity would thwart the subject's plans, but also those where an unknown identity would help the subject attain their goals. Suppose that (2)(a) and (b) are true and that Oedipus doesn't know that Hesperus is Phosphorus. The truth of that identity is not going to cause Oedipus any special problems. If he comes to learn that the *USS Evening Star* is going to Hesperus, and continues not to believe that Hesperus is Phosphorus, he might unintentionally end up visiting Phosphorus. But that shouldn't trouble him. Phosphorus is just as good a place to visit as Hesperus.

If (2)(a) and (b) are true, then knowledge of the identity would help Oedipus attain his goal. But surely there are no special cognitive mechanisms dedicated to the task of discovering those identities knowledge of which would enhance one's capacity to attain one's ends. Rather it is just that we have at our disposal cognitive capacities that we can use to try to find out whatever we think we need to know in order to get what we want. Many relevant truths are within the reach of these capacities. But equally many are not. It's a hard life.

The generalisation that would be needed to underwrite explanations like (2) (seen as Fodor wants to see them, as subsumptions under *ceteris paribus* laws) would be along the lines of (5):

(5) *Ceteris paribus*, if a=b and an agent, x, desires that Fa and believes that [if x does D, then Fb], then x will learn that a=b.

But (5) is implausible. To put it mildly. Suppose that: I am at a roulette table, I have £1,000 stake, I believe that if I place my stake on the winning number, then I will win £36,000 and I desire to win £36,000. The winning number is 2. Am I particularly likely to find out that 2 is the winning number and so proceed to bet on it? No. My cognitive mechanisms are not up to the task of discovering the relevant identity in time. But there is no breakdown of mechanisms that work in normal circumstances. Nothing exceptional is happening. *Cetera sunt all too paria* in this case.

Fodor might deal with the last example, and with the original Oedipus story, which was all about a guy who ended up accidentally marrying his mother, by borrowing from the real Grandpa Russell and allowing that the content of complex, descriptive concepts isn't identical with their denotations. Thus the content of Oedipus's MY MOTHER concept would be a complex structure including Oedipus himself and the property of motherhood. And that complex structure is not the same thing as Jocasta herself.

It's possible that Fodor's resistance to this standard neo-Russellian position, which agrees with Russell on this, is due merely to its relative complexity. If you've got denotation and mental syntax, then why bring in more content? On the other hand, the more content is there to brought in, if necessary, since the conceptual components of complex descriptive concepts still have their contents, even when they are appearing as such components.

But, certainly on Fodor's view, there are still lots of non-descriptive singular concepts, like HESPERUS and PHOSPHORUS. And surely there are no special mechanisms for ensuring that when we need to know which of these are co-referential, we'll be able to find out. Sad, but true.

What about (4)? According to Fodor, that's ok too. That's because thoughts featuring empty concepts take as their objects incomplete propositions. Thus believing that Gunther fought with Hagen and believing that Fafnir fought with Hagen involve doxastic relations with the very same incomplete proposition — something like $\langle\langle x, y \rangle \text{ fought-with} \rangle$. So (4)(a) and (c) say the same thing. And then, of course, (4) is valid.

Well it certainly doesn't look particularly valid, does it? It even makes (2) look good. I think reason for this is that (2) is OK when read as an enthymematic explanation (i.e. one with a missing premise) in de re mode, along the lines of (2')

(2') a. Oedipus desires of Hesperus that he fly to it.

- a. Oedipus believes of Phosphorus that the *USS Evening Star* is about to depart for it.
- b. Hesperus=Phosphorus

So c. Oedipus boards the *USS Evening Star*

But compare (4) with (4*):

***(4) a. Helga believes of Gunther that he fought with Hagen

b. Helga believes of Gunther that he is King of the Gibichungs

c. Gunther=Fafnir

so c. Helga believes of Fafnir that he fought with Hagen

A tissue of falsehood, if ever I wrote one.

I think the motivation behind countenancing examples like (2) is pretty good. It surely is right that someone in the position ascribed to Oedipus in (2)(a) and (b) is not unlikely to end up boarding the *USS Evening Star*. Lots of people know that Hesperus is Phosphorus and Oedipus might be one such, even if (2) doesn't say that he is. Moreover, someone who doesn't know that Hesperus is Phosphorus might find out. Particularly if they want to go to Hesperus. Special mechanisms or no, we do sometimes find things out when we need to. I suggest that it's not really that *ceteris paribus* we will find these things out, rather it's that there is a non-negligible chance that we will. Perhaps the thing to say about (2), understood along the lines of (2'), is that it works as an explanation because the truth of premises increases the likelihood of the truth of the conclusion.

(4) can't be read *de re*, however, so it would be nice if Fodor said something more about why it would be o.k. for us to think that it's o.k.

If we don't think it's ok, then how should we proceed? I think we should acknowledge the wisdom of our elders and believe what Grandpas Frege and/or Russell told us. We should think of HESPERUS and PHOSPHORUS as having different contents, quasi-senses, that might perhaps be accounted for in terms of descriptive contents that they associate with.

Incidentally, I don't think there's much chance of getting Uncle Turing

to pull Fodor's chestnuts out of this fire. The suggestion might cross one's mind that we could distinguish HESPERUS from PHOSPHORUS in terms of their syntactic identities in the language of thought. But I wouldn't be too happy with that suggestion. We want to understand how lots of people can all believe that Hesperus shines in the evening. But it's very unlikely that lots of peoples' HESPERA are syntactically type-identical.

Anyway enough of HESPERUS and PHOSPHORUS. It's time to think about general concepts.

2. General concepts and narrow contents

Fodor takes general concepts express properties (see e.g Fodor 1990 chapter 3). Concepts that express the same property have the same informational content and that's the only content they have. And, according to Fodor, this content is broad, in the sense that Putnamian twins (beings that are exact replicas in all intrinsic, micro-physical respects) could have counterpart concepts with different contents.

In this section, I will discuss whether Fodor's theory of content does (or should, if properly moulded) dictate that the content of general concepts is broad rather than narrow. This will involve me running of a version of an argument for narrow content that I often use, adapted to Fodor's specifics.

According to Fodor, a concept can express a property that is not instantiated in the concept's world (1990, 101). Thus PHLOGISTON and ECTOPLASM can have different contents (unlike FAFNIR and HAGEN). Well, let us just have a look at what sort of a content PHLOGISTON would have, shall. What kind of a property does it express? Presumably there are many twin worlds where twin 18th century scientists hold a 'phlogiston' theory that looks just like the one that was held here. Only the twin theories are true, the twin word 'phlogiston', in each case, successfully corresponding to a natural kind in the theorists' environment. Moreover, some of these theories will be true of different natural kinds, we might be so bold as to presume.

Presuming even more boldly, I claim that a plurality of these worlds with their own particular phlogiston-like substances will be equal closest to our own. Our PHLOGISTON, then, does not apply to any one of those natural kinds in particular, while failing to apply to the others. But, according to Fodor, it still has to express some property or other. What would that be? Well one's best guess would be that it's something like a higher-order functional property, a property shared by all samples of what we might call the 'phlogiston-motley', viz. samples of all the different phlogistonoid kinds (c.p. 1994, 31) — the property of being phlogistonoid.

Fine. Now let us home in on one of those worlds we've been talking about and focus on a particular inhabitant, the famous T-phlogiston theorist, T-Stahl (T-phlogiston being the specific natural kind in that particular world). What is the content of T-Stahl's PHLOGISTON? A normal externalist would say: it extends over samples of T-phlogiston and only those. I take it that Fodor would like to be a normal externalist in that respect and so would claim that it expresses the property of being T-phlogiston.

But then my question is: what is to prevent T-Stahl's PHLOGISTON from extending over all samples of the phlogiston-motley, from expressing the property of being phlogistonoid? And if that's what it expresses, then its content would be narrow and not broad, since he and we and all of our twins would be like T-Stahl in that respect.

Fodor, of course, has a theory of content that could be used to address the question of what T-Stahl's PHLOGISTON refers to. Here is the original (Fodor 1987, 109) version. A sufficient condition for an LOT expression 'A' to express property A and *not* to be true of instances of some other property B, is:

1. A's cause 'A's
2. 'A' tokens are *not* caused by B's in nearby worlds in which A's *don't* cause 'A's.
3. A's cause 'A's in nearby worlds in which B's don't cause A's.

So suppose HORSE expresses the property of being a horse. That's because encounters with actual horses do or would cause HORSE to pop in to your head. Sometimes actual cows cause that too, say, when, on a dark night, you vaguely see one that has wandered into your stables and you rashly assume

it's a horse. But your HORSE doesn't express the disjunctive property of being a horse or a cow on a dark night. And that's supposed to be explained by the 'asymmetric dependency' clauses 2 and 3. The causal connection between cows and HORSEs is only there because of the more fundamental connection between horses and HORSEs (clause 2). But the connection between HORSEs and horses doesn't depend on the connection between HORSEs and cows (clause 3). It's because you think the cow is a *horse* that you think HORSE when you see it. But your thinking HORSE when you see a horse doesn't have anything to do with cows.

Now 1-3 provide only the roughest of sketches of what the proper theory would look like. Which A's, exactly cause, 'A's? Where? When? How often? So now, our question is what does Fodor's theory say about T-Stahl's PHLOGISTON? And the answer hinges on the interpretation of the theory. And that's kind of a tricky question that deserves a little digression.

I digress.

In Segal (2000) I endeavoured to argue that Fodor had gotten confused about his own theory — particularly clause 1 - and that, contrary to what he was saying in his (1994) the relevant counterfactuals about what would cause what dictate that the contents of general concepts are narrow. Actually, though, it's not clear that Fodor was confused. It's possible that I was. Here's what happened.

Fodor was discussing the case of a Swampman (S2) who had just come into being on an XYZ twin Earth. Fodor thinks that when this Swampman thinks WATER he is thinking about XYZ. (A Swampman arising on Earth, by contrast, would in such a state be thinking of H₂O.) And the reason for that, according to Fodor, is that "it's because it's XYZ that would cause his WATER tokens in all worlds that are nearest to Twin Earth, there being, again by assumption, no H₂O on any of them." (1994 118). So Fodor's interpretation of clause 1 was something like this:

(FC) A's would cause 'A's, in the right circumstances, in nearby possible worlds

That struck me as a bit barmy. First of all it's barmy because S2 is a twin of Donald Davidson's. So if you asked S2 about the constitution of what he called 'water' he'd tell you it was H₂O. And he could persuade you that he knew what H₂O is too. But that bit of barminess can be easily avoided by thinking of S2 as a twin of Oscar, who lived in the 18th century and didn't know any chemistry. But more importantly: What's this about nearby possible worlds? Why can't we think about things that exist only in extremely distant possible worlds, like free lunches? How could the mere unlikelihood of S2 bumping into any H₂O stop him from even thinking about it? So then I said that on Fodor's theory, the right way to think of clause 1, the chief content-fixing counterfactual was, rather, along the lines of (CS):

(CS) A's would cause 'A's if A and the thinker were together in the right circumstances

You know: what would you think an A was if, say, you were to bump into one.

Well maybe (CS) could reasonably be read into one of Fodor (1987). And then again, maybe it couldn't. But either way, what I missed was the subtle and complex point that it's Fodor's theory and so *he* gets to choose the kind of counterfactuals that, according to Fodor's theory, fix contents. So he can perfectly well appeal to counterfactuals of the form (CF) if he wants to. And those would render content broad, just as he says.

Fair enough. The theory looks reasonable. It does mean that we can't have atomic concepts that express properties that are instantiated only in distant possible worlds. But that's not necessarily a calamity. (I think my SOUL is one, but I wouldn't like to have rely on that point).

Incidentally, one might wonder whether the question of whether S2's WATER would be true of H₂O, according to Fodor's theory, would depend on the asymmetric dependency clauses. They, you will recall, are supposed to account for error. So if S2 were presented with a sample of H₂O and to think WATER, would he be making a mistake? What does asymmetric dependency say about that, then?

At first blush, you might think, it says that yes, S2 would be making a mistake. For it's true that if XYZ didn't cause his WATER to token, then H₂O wouldn't. But the reverse is not true. Why not? Well, suppose that S2 is on a planet where there is only XYZ and that on another planet in his universe there is some H₂O. Note that we can consider this to be a relevant possibility, since the question is whether asymmetric dependency gives you broad content, so any case in which it distinguishes twins will do. So now, considering S2 in this scenario, is it the case that if H₂O were not to cause his WATER to token, then XYZ wouldn't either? On the face of it, probably not, because were he to travel to the H₂O world, he might well come to discover the difference between H₂O and XYZ and then not think his WATER when confronted with the former.

But actually, I think this won't work for Fodor. It won't work because this is a counterfactual that would require a change in S2. And when you've changed your subject, it's too hard to tell what her concepts were like before the change. To know what a subject's concepts mean you have to look at her current dispositions. Asymmetric dependency has to be understood synchronically, rather than diachronically.

OK, so to return to our thread: we've got two possible candidates for the form of content-fixing counterfactuals, (CS) and (CF) and we have to pick one. How should we decide? At least part of what is driving Fodor's choice here are intuitions about what would mean what in which possible worlds. (That's sort of obvious. But see his (1994, 118) for a hint to that effect). And, indeed, carefully constructed twin Earth thought experiments can be used to reveal that according to certain intuitions of many folk psychologists, content is broad. (I myself think that this gives a misleading picture of folk psychology: there are intuitions that go the other way, etc. etc., but that's not important now.)

Since Fodor is prepared to endorse explanations like (4), direct application of intuitions to particular types of case is evidently a highly defeasible constraint on decisions about how a Fodorian theory should

be shaped to deal with those types of case. And when it comes to the choice between (CF) and (CS) we have a much better source of constraint: methodology. If you want to know what content is like, then, think about the role that content would play in the best psychological theory. And it is arguable that a psychology that deployed a notion of content derivable from (CS) would be better than one using a notion derivable from (CF). For (CS) allows you to generalise over twins in cases where (CF) doesn't. Under (CS), you can generalise over Davidson and swampman. So (CS) gets you greater generality. And (CS) correctly predicts how S2 would behave if he were moved to Earth, too. And, crucially, it doesn't lose you anything. For example, you can still generalise over Davidson and swampman and Fodor and lots of us who all believe that you need to boil water if you're going to make tea. No great explanatory benefit would be gained by not generalising over twins in the cases where (CF) prevents it.

But I don't propose to go into all that now. After all, this is just a digression.

I cease to digress

For T-Stahl's PHLOGISTON to fail to extend over the motley, there would have to be some relevant counterfactual that is true of him and false of Stahl (or vice versa). Well, if (CS) were the relevant sort of counterfactual, then there would be no distinction of content among the Stahl twins. So if Stahl's PHLOGISTON extends over the motley, then so does T-Stahl's.

But look, even if Fodor sticks to (CF), there will be no relevant counterfactual to distinguish the twins. The only extra ingredient that (CF) brings in is distance between worlds. To get the phlogiston motley out of the extension of T-Stahl's PHLOGISTON, while leaving it in the extension of Stahl's PHLOGISTON, the latter would have to be nearer other phlogiston worlds than the former. But that just doesn't seem to be so.

So content is narrow and I win.

Conclusion

So where does all that leave us? Folk psychology is quasi-Fregean. It recognizes quasi-sense, even for concepts of particulars, like HESPERUS. In giving up on any sense-like kind of content, a Fodorian psychology would be giving up a great deal. It might be harder to naturalise quasi-sense than reference. But it looks like psychology can't get very far without it. In this respect, Fodor's naturalism leads him to be too revisionary. It's a baby and bathwater thing. There are various things that you can usefully do when your car gets a ping in one of its cylinders; but declining to quantify over the engine is not one of them.

When it comes to twins and narrow content, I'd suggest, by contrast, that Fodor is too conservative. Fodor is sympathetic to externalist intuitions generated by standard twin Earth experiments and I think it is mainly just that which leads him to try to shape his theory of content into an externalist one. (Indeed an earlier version of theory had an ad hoc clause mainly for that (1990, 121)). Such intuitions do reflect a

strand of externalism in folk psychology. But it's a strand that can easily be extracted without the whole ball unravelling. The discussion above indicates that when it comes to non-singular concepts, the best Fodorian psychology, i.e. the best psychology that could be naturalised along Fodorian lines, would be internalist rather than externalist.

Such conclusions leave one with lots of candidate total packages. But one natural line would be to adopt a two-factor account of content for both singular and general concepts. Both have quasi-sense and reference. The reference of a non-empty singular concept would be an object. And real-world relations between the object and the concept could underpin the explanatory role of reference in psychological explanation. For example the fact that Oedipus's HESPERUS refers to Phosphorus might play a role in explaining why his desire to visit Hesperus and his belief that the *USS Evening Star* is about to depart for Phosphorus are parts of a causal chain that ends up with him boarding the starship. Quasi-senses remain in place to distinguish co-referring concepts and to account for empty ones.

And there's a rub. For the quasi-sense of a singular concept will probably have to be constructed somehow out of general concepts that are associated with it (See Segal (forthcoming a) for discussion). This will almost inevitably lead to holism. And Fodor just hates holism (Fodor 1987, Fodor and LePore 19xx).

For kind concepts, one option would be to let quasi-sense be the narrow content and let the local kind be the reference. Thus the quasi-sense of a prescientific, Earthly WATER could be the property of being hydroid and the reference could be water. But that idea won't generalise to other sorts of non-singular concept, like WITTY, PHILOSOPHERS and UNUSUAL. So it might be better if we were to think of reference merely in terms of the relatively uninteresting notion of actual extension and let quasi-sense take the starring role.

BIBLIOGRAPHY

Davidson, D., 1987 'Knowing One's Own Mind' *Proceedings and Addresses of the American Philosophical Association*, 441-458.

Dretske, F. (1981), *Knowledge and the Flow of Information*, Cambridge MA: MIT Press.

Fodor, J., (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge MA: MIT Press.

- (1990) *A Theory of Content and other Essays*, Cambridge MA: MIT Press.

- (1994), *The Elm and the Expert*, , Cambridge MA: MIT Press.

Fodor, J. and LePore, E., (1992) *Holism: A Shopper's Guide*, Oxford UK and Cambridge MA: Blackwell.

Kim, J., (1998) *Mind in the Physical World: an Essay on the Mind-Body Problem and Mental Causation*, Cambridge MA: MIT Press.

Millikan, R., (1984), *Language, Thought and other Biological Categories*, Cambridge MA: MIT Press.

Papineau, D., (1993) *Philosophical Naturalism*, Oxford: Basil Blackwell.

- (1996) "Theory dependent terms" *Philosophy of Science*, 63, 1-20.

Noonan, H., (1986) "Russellian Thoughts and Methodological Solipsism" in Butterfield, J., (ed) *Language, Mind and Logic*, Cambridge: Cambridge University Press, 67-90.

Salmon, N., (1986) *Frege's Puzzle*, Atascadero, CA: Ridgeview Publishing Company

Segal, G., (1989), "The Return of the Individual", *Mind* 98 35-57

- (1995) "The Modularity of Theory of Mind" in Carruthers A. and Smith, P., (eds) *Theories of Theories of Mind*. Oxford: Oxford University Press.

- (2000), *A Slim Book about Narrow Content*, Cambridge MA: MIT Press.

- Forthcoming a, "Intentionality," in *The Oxford Handbook of Analytic Philosophy*, Jackson, F., and Smith, M., (eds) Oxford: Oxford University Press.

- Forthcoming b, "Ignorance of Meaning," in *The Epistemology of Language* Barber A., (ed) Oxford: Oxford University Press.

Turing, A., (1950), "Computing Machinery and Intelligence" *Mind* 59: 434-460

Wellman, H., (1990), *The Child's Theory of Mind*, Cambridge MA: MIT Press.