

Cambridge University Press  
 978-0-521-10844-7 - Language, Mind and Logic  
 Edited by Jeremy Butterfield  
 Excerpt  
[More information](#)

## Banish disContent\*

JERRY FODOR

It is a curiosity of the philosophical temperament, this passion for radical solutions. Do you feel a little twinge in your epistemology? Absolute scepticism is the thing to try. Has the logic of confirmation got you down? Probably physics is a fiction. Worried about individuating objects? Don't let anything in but sets. Nobody has yet suggested that the way out of the Liar paradox is to give up talking, but I expect it's only a matter of time. Apparently the rule is: if aspirin doesn't work, try cutting off your head.

The latest of these cures for which there is no adequate disease is the suggestion that – largely on account of some semantic puzzles about content – psychological theory should dispense with the attribution of propositional attitudes. In my view, this is a *grotesque* proposal; both because we can't do without propositional attitude psychology and because it's far from clear that the semantic problems are as bad as they are alleged to be. The plan of this paper is to look very briefly at the first of these considerations and rather extensively at the second.

### 1. WHY WE CAN'T DO WITHOUT PROPOSITIONAL ATTITUDE PSYCHOLOGY

We can't do without propositional attitude psychology because, on the one hand, propositional attitude psychology works and, on the other hand, nothing else does. Propositional attitude psychology works so well that the mechanism is practically invisible. It is like those mythical Rolls Royce cars whose engines are sealed when they leave the factory; only it's better because it's not mythical. Someone I don't know telephones me at my office in Cambridge from – as it might be – Arizona. 'Would you like to lecture here next Tuesday?' are the words that he utters. 'Yes, thank you. I'll be at your airport on the 3:00 pm flight' are the words that I reply.<sup>1</sup> That's

\* For useful discussions of earlier drafts of this work, I'm especially indebted to Ned Block, Tyler Burge, David Israel, and Ron McClamrock.

<sup>1</sup> So invisible is the mechanism that one may overlook such facts as that even the notion of a word belongs to commonsense intentional psychology and not to acoustics (say) or to physics.

Cambridge University Press  
 978-0-521-10844-7 - Language, Mind and Logic  
 Edited by Jeremy Butterfield  
 Excerpt  
[More information](#)

JERRY FODOR

*all* that happens, but it's more than enough; the rest of the burden is routinely taken up by theory. And the theory works so well that, several days later (or weeks later, or months later, or years later; you can vary the example to taste) and several thousand miles away, there I am at the airport, and there he is to meet me. Or if I don't turn up, it's less likely that the theory has failed than that something went wrong with the airline and the aeroplane is late. For these sorts of purposes, people are more predictable mechanisms than jets.

The point is that the theory from which we get this extraordinary predictive power is just good old commonsense propositional attitude psychology. That's what tells us, for example, how to infer people's intentions from the sounds they make (if someone utters the form of words 'I'll be at your airport on the 3:00 pm flight' then, with enormously high probability, he intends to be at your airport on the 3:00 pm flight) and how to infer people's behaviour from their intentions (if someone intends to be at your airport on the 3:00 pm flight then, with enormously high probability, he will produce behaviour of a sort which will eventuate in his arriving at that place at that time barring mechanical failure). And all this works not just with people whose psychology you know intimately: your closest friends, say, or the spouse of your bosom. It works with absolute strangers: people you wouldn't know if you bumped into them. And it works not just in laboratory conditions – where you can control the interacting variables – but also, indeed preeminently, in field conditions where *all* you know about the sources of variance is what commonsense psychology tells you about them. Remarkable. If we could do that well with predicting the weather, no one would ever get his feet wet. Yet the aetiology of the weather is surely child's play compared with the causes of behaviour.

Physics, by the way, doesn't begin to compete: though we can, by exploiting our knowledge of their psychology, routinely bring off these extraordinary predictions of the *behaviour* of 'intentional systems', there is no way – practically speaking – that we can hope to predict their *trajectories*. If you think of me as two hundred pounds (less in the jogging season) of peripatetic philosopher whose itinerary you've just been informed of, there are long odds you can predict where I'll be at 3:00 pm on Tuesday. Whereas, if you think of me as an equivalent mass of hydrocarbons which just produced a couple of seconds of local acoustic perturbation, nothing useful follows from any physical science we now have or can ever reasonably expect.

To be sure, that's just *commonsense* propositional attitude psychology; and it's mostly *implicit* commonsense propositional attitude psychology to boot. It might still be that, for scientific purposes, something other – something better – is required. After all, what we want in science is insight and

## Banish disContent

explanation, not just prediction and control. It is now widely said that propositional attitude psychology is seen to be a 'stagnant' theory when viewed in the light of these desiderata. By contrast, what is not stagnant, what is said to be making scientific progress, is computational psychology, in which appeals to the attitudes are dispensed with. (See, for example, Stich 1983.)

I am, I suppose, as good a friend of computational psychology as the next chap; especially if the next chap is a philosopher. But that view of the current scene in cognitive science strikes me as a gross distortion. For, one can say in a phrase what it is that computational psychology has been proving so successful at: viz. *the vindication of generalizations about propositional attitudes*; specifically, of the more or less commonsense sort of generalizations about propositional attitudes illustrated a paragraph or so above. Thus, for example, we've got fragments of a theory of perception, and it makes clear how a computational system could regularly come to believe that *P* in causal consequence of its being visibly the case that *P*. Analogously, we've got fragments of a theory of language, and what it does is to make clear how, in a computational system, the intention to communicate the belief that *P* could eventuate, causally, in the production of an utterance that means that *P*. Or again, we've got fragments of a theory of memory; it explains how a computational system could come to believe that it was once the case that *P* in causal consequence of the fact that it was once the case that *P*. And so forth. The mental processes that such computational theories acknowledge are, no doubt, in some sense 'purely syntactic'; at a minimum, they're supposed to be specifiable without recourse to intentional idiom. But the generalizations that such theories account for are intentional down to their boots. Accordingly, computational theories of mental processes don't replace the commonsense story about propositional attitudes and their behavioural effects. Rather, what a computational theory does is to make clear the mechanism of intentional causation; to show how it is (nomologically) possible that purely computational – indeed, purely physical – systems should act out of their beliefs and desires.

I don't suppose that all that comes to much; I just wanted to remind you that – given standard principles of Scientific Realism – there is an enormous *prima facie* case for taking commonsense propositional attitude psychology to be more or less true. No doubt this case is less than literally conclusive. In principle, somebody might find overwhelming reason for rejecting intentional psychology; so overwhelming that rejecting intentional psychology would be rational even if we had no idea at all what to replace it with. (As, indeed, we don't.) But *overwhelming* reason is surely what rationality would require. If the sky opened up and God told us that there aren't any propositional attitudes, then I suppose that we would

Cambridge University Press  
 978-0-521-10844-7 - Language, Mind and Logic  
 Edited by Jeremy Butterfield  
 Excerpt  
[More information](#)

JERRY FODOR

have to believe Him. But it's not clear that a great deal less than that would do.

Having thus made clear where the burden of proof resides, let's look at the semantic puzzles about content.

## 2. PUZZLES ABOUT CONTENT: TWIN-EARTH

Is there anybody who *hasn't* heard? There's this place, you see, that's just like here except that they've got XYZ where we've got H<sub>2</sub>O. In this place, there's someone who's just like me down to and including his neurological microstructure. The intuition that we're invited to have is that, in virtue of the hydrochemical facts and in spite of the neurological ones, the form of words 'water is wet' means something different in his mouth than it does in mine. And, similarly, the content of the thought that Twin-Me has when he thinks (*in re* XYZ, as one might say) that water is wet is different from the content of the thought that I have when I think that water is wet *in re* H<sub>2</sub>O.

Suppose these intuitions are reliable; what follows? Well, on the one hand, my Twin and I are identical in physical constitution but, on the other hand, our thoughts have different truth conditions (i.e. what makes *his* 'water'-thoughts true is the facts about XYZ, whereas what makes my water-thoughts true is the facts about H<sub>2</sub>O). So, it looks as though we have to say either (a) that thoughts don't have their truth conditions essentially (i.e. that two tokens of the *same* thought can have *different* truth conditions), or (b) that type-identity of thoughts doesn't supervene upon biochemical type-identity of their thinkers.

That, then, is the 'Twin-Earth problem'. Except that so far it isn't a problem; it's just a handful of intuitions together with a commentary on some immediate implications of accepting them. If that were all there is, the right response would surely be 'So what?' What connects the intuitions and implications with the proposal that we give up on propositional attitude psychology is a certain *Diagnosis*. And, while a lot has been written about the intuitions and their implications, the Diagnosis has gone largely unexamined.

*Here's the Diagnosis:* 'Look, on *anybody's* story, the notion of content (or "intention") has got to be a little problematic. For one thing, it is proprietary to the information sciences, and *soi-disant* "emergents" ought, at a minimum, to bear the burden of proof. At the very minimum, if you're going to have contents, you owe us some sort of account of their individuation.

'Now, prior to the Twin-Earth problem, there was some sort of account of their individuation; you could say, to a first approximation, that identity of intention depends on identity of extension. No doubt that story leaked

a bit: thoughts about the Morning Star look to be different in content from the corresponding thoughts about the Evening Star, even though their truth conditions are arguably the same. But at least one could hold firmly to this: “No difference in extension without some difference in intention.” Conversely, it was a test for identity of intention that the extensions came out to be the same. And that was the best test; it was the one source of evidence about intentional identity that seemed surely reliable. Compare the notorious wobbliness of intuitions about synonymy, analyticity and the like.

‘But now we see that after all it’s not true that difference of extension implies difference of intention. Your Twin’s ‘water’-thoughts and your own are intentionally identical – assuming supervenience – but extensionally distinct. So unclear are we now about what intentional identity comes to – hence about what identity of propositional attitudes comes to – that we can’t even assume that if two thoughts are intentionally identical they will be true and false together. The fact is, as we now see, we have no idea at all of what criteria of individuation for propositional attitudes might be like; hence we have no idea at all what counts as evidence for identity of propositional attitudes. (Given which, by the way, it is hardly surprising that propositional attitude psychology so rarely proves to be disconfirmed; once you abandon the extensional constraints, ascriptions of propositional attitudes are arbitrarily available for being gerrymandered.)

‘*To summarize*: Inferences from difference of extension to difference of content used to bear almost all the weight of propositional attitude attribution. That was, however, a frail reed and now it has broken. The Twin-Earth problem is a problem *because it breaks the connection between extensional identity and intentional identity*.’<sup>2</sup>

Now, the Twin-Earth intuitions are fascinating, and if you care about the semantics of kind-terms you will, no doubt, do well to attend to them. But, as I’ve taken pains to emphasize, you need the Diagnosis to connect the Twin-Earth intuitions to the issues about belief/desire psychology, and – fortunately for friends of propositional attitudes – the Diagnosis rests

<sup>2</sup> Cf. Putnam (1983: 148): ‘Once we decide to put the reference...aside...far from making it easier for ourselves to decide whether the representations are synonymous, we have made it impossible. In fact, the first approximation we have to a principle for deciding whether words have the same meaning or not in actual translation practice is to look at the extensions. “Factoring out” differences in extension will only make a principled decision on when there has been a change in meaning totally impossible.’ I’m not, however, claiming that Putnam actually endorses the Diagnosis I sketched in the text; I’m not sure whether he thinks that the Twin-Earth cases show that intentional Realists have to ‘factor out the extension’ from the individuation conditions on meanings. Such passages as the one I’ve quoted make it seem as though he does, but I can’t find a place where he straight-out says it. My claim, in any event, is that without the Diagnosis, it’s unclear that Twin-Earth considerations have any bearing at all on the facticity of propositional attitude ascriptions.

JERRY FODOR

on a quite trivial mistake: THE TWIN-EARTH EXAMPLES DON'T BREAK THE CONNECTION BETWEEN INTENTION AND EXTENSION; THEY JUST RELATIVIZE IT TO CONTEXT.

Suppose that what you used to think, prior to Twin-Earth, is that intentions are something like functions from thoughts to truth conditions. Presumably, a truth condition would then be a function from worlds to truth values; condition *TC* takes world *W* onto the value *T* iff *TC* is satisfied in *W*. For example, in virtue of its intention, the thought that it's raining has the truth condition *that it's raining* and is thus true in every world in which it's raining and false in every other world.

I hasten to emphasize that if you don't – or didn't – like that story, its quite all right for you to choose some other; my point is going to be that if you liked that *kind* of story before Twin-Earth, you're perfectly free to go on liking it now. For, even if all the intuitions about Twin-Earth are right, and even if they have the implications that they are said to have, extensional identity still constrains intentional identity because *intentions still determine extensions* (relative to a context). If you like, intentions are functions from *contexts* and thoughts onto truth conditions.

What, if anything, does that mean? Well, there is presumably something about the relation between Twin-Earth and Twin-Me in virtue of which his 'water'-thoughts are about XYZ even though mine are not. Call this condition that's satisfied by {Twin-Me, Twin Earth} condition *C* (because it determines a Context). Similarly, there must be something about the relation between me and Earth in virtue of which my water-thoughts are about H<sub>2</sub>O even though my Twin's 'water'-thoughts aren't. Call this condition that is satisfied by {me, Earth} condition *C'*. We don't know what sorts of things *C* and *C'* constrain, but causal relations of some kind are the current best bet and, anyhow, it doesn't matter much for the purposes at hand. Because we *do* know this: short of a miracle, it must be possible to satisfy *C* without satisfying *C'* and vice versa. How do we know that? Well, because short of a miracle, the following must be true: if an organism shares the neurophysical constitution of my Twin *and satisfies C*, it follows that its thoughts and my Twin's thoughts share their truth conditions. In particular, the following counterfactual is true: in a world where I am in my Twin's context, given the neurophysical identity between us, my 'water'-thoughts are about XYZ iff his are.<sup>3</sup>

But now we have an extensional identity criterion for intentions: two intentions are identical only if they effect the same mapping of contexts onto truth conditions. Specifically, your thought is intentionally identical to mine only if in every context in which your thought has truth condition *TC* mine

<sup>3</sup> Given shared context, the neurophysical identity of thinkers is, of course, sufficient but not necessary for the intentional identity of their thoughts. This is a matter we will return to.

has truth condition *TC* and vice versa.<sup>4</sup> It is worth noting that, by this criterion, my Twin's 'water'-thoughts are intentionally identical to my water-thoughts; they have the same intentional content even though, since their contexts are *de facto* different, they differ in truth conditions. In effect, what we have here is an extensional criterion for what is sometimes called 'narrow' content. (The 'broad content' of a thought, by contrast, is what you can semantically evaluate; it's what you get when you specify a narrow content *and fix a context*. This makes the notion of narrow content the more basic of the two; which is just what sensible people have always supposed it to be.)

We can now see why we ought to reject both of the following two suggestions found in Putnam (1975): That we consider the extension of a term(/concept) to be an independent component of the (to use Putnam's terminology) 'vector' that is the meaning of the term; and that we make do, in our psychology, with stereotypes *instead of* intentions. The first proposal is redundant since, given a context, intentions determine extensions; and we have – so far at least – no reason for supposing that there aren't intentions. The second proposal is unacceptable because, unlike intentions, stereotypes *don't* determine extensions *even* given a context. And, as the Diagnosis rightly says, we need an extension-determiner as a component of the meaning vector because we rely on 'different extension → different intention' for the individuation of concepts.

There are, no doubt, serious objections to the line that I've been pushing. For example, 'different extension → different intention' gives us only a necessary condition on intentional identity where what we really want is a biconditional. And if we try to parley what we've got into what we want, we run into the familiar troubles about thoughts that are extensionally identical but intentionally distinct. So, for example, it looks like – on the biconditional version – the thought that  $2 + 2 = 4$  would come out to have the same intention as the thought that  $3 + 3 = 6$ . Not good.

However, these are just the old-fashioned, pre-Twin-Earth objections to the reduction of intentional identity to extensional identity, so they needn't concern us here. The main point to bear in mind is that if 'different extension → different intention' substantively constrains the attribution of propositional attitudes, then so too does this same principle when it is

<sup>4</sup> This condition has to be applied with some care when the thoughts are truly indexical (as some thoughts no doubt are). Thus specifying the context which – together with its intention – determines the truth condition of the thought that I am bored includes specifying the thinker of that thought.

Somewhat similarly, 'visiting cases' suggest that Context must be four dimensional: what state of affairs makes your thought true or false probably depends on your history as well as your current surround.

This is all much in the spirit of David Kaplan – or, at least, I hope it is – though, patently, he is not to be blamed for any of it. See also Fodor (1982).



Cambridge University Press  
 978-0-521-10844-7 - Language, Mind and Logic  
 Edited by Jeremy Butterfield  
 Excerpt  
[More information](#)

JERRY FODOR

relativized to context. So, if the worry about propositional attitudes is that Twin-Earth shows that extensions don't constrain intentions, the right thing to do about Twin-Earth is to STOP WORRYING.

Which is not, alas, to say that there is *nothing* to worry about.

*First worry:* It is one thing to have an (e.g. extensional) criterion for the intentional identity of thoughts (which, indeed, we do); it is quite another to be able to say what the intention of a thought is (which, indeed, we can't). What, for example, is the intention of the thought – or, as we might as well as well say, what is the *thought* – such that when I have that thought its truth condition is that H<sub>2</sub>O is wet and when my Twin has it its truth condition is that XYZ is wet? Or, to put it another way (to put it, in fact, the old way) what is the *concept* WATER assuming that the concept WATER is not the concept H<sub>2</sub>O?

I don't know the answer to this question, but I know how to find out: find out what (the word) 'water' means. Perhaps 'water' means something like 'the local, transparent, potable, dolphin-torn, gong-tormented..., etc. stuff one sails on'. If that is so, then the intention of my thought that water is wet (and of my Twin's thought that 'water' is wet) is that the local transparent, potable, dolphin-torn, gong-tormented..., etc. stuff one sails on is wet. My water thoughts are about H<sub>2</sub>O because H<sub>2</sub>O is the local transparent, potable..., etc. stuff in the context in which my water thoughts transpire (and similarly, *mutatis mutandis*, for XYZ in the context of my Twin's 'water'-thoughts). This would mean that a kind concept is some sort of implicit description after all, and it does leave one wondering just which description the concept WATER is. Just as it always used to.

But why does this question have to have an answer? Why, that is, does there have to be a way of expressing the intention of 'water' in English – or in Tw-English, for that matter – except as the intention of 'water'? Does English have to have more than one way of expressing every intention that it can express at all? Does every word (concept) that has an intention have to have a definition? (There is, of course, a way of expressing the intention of 'water' in, say, French: 'water' *veut dire l'eau*. But if you're worried about what the intention of 'water' could be – what it could be that picks out H<sub>2</sub>O in the context Earth and XYZ in the context Twin-Earth – I don't suppose you will find this consideration comforting.)

I think this is all rather puzzling, but notice that it's a sort of puzzle that one has to face regardless of one's views about the Twin-Earth problem. So, for example, suppose you think that Twin-Earth shows that intention doesn't determine extension and/or that intention doesn't supervene upon neurophysiology. Still – and quite aside from all that – there's the following question: When my Twin says 'water is wet' what does what he says mean?



## Banish disContent

Not, presumably, that water is wet since, on the present assumptions, there is no reason to identify the intention of 'water<sub>TW</sub>' (viz., of the Tw-English vocable that is pronounced like our word 'water') with the intention of 'water'. And not that XYZ is wet since my Twin will presumably take 'water is XYZ' to say something informative; something, indeed, which he might wish to deny. And not, for sure, H<sub>2</sub>O is wet since there isn't any H<sub>2</sub>O on Twin-Earth, and my Twin has never so much as heard of the stuff. It looks like either 'water<sub>TW</sub>' means water (i.e. has the same intention that 'water' does) or its meaning is inexpressible in English. The moral seems to be: there is going to be a problem about saying what 'water' and 'water<sub>TW</sub>' mean whatever view you take of the Twin-Earth problem. Or, to put it another way: the real problem is to figure out what 'water' means. You tell me what 'water' means and I'll bet that 'water<sub>TW</sub>' means that too.

*Second worry:* Since intentional identity supervenes upon neurophysical identity, and since – given a context – there is no difference of extension without a difference of intention, it follows that the corresponding thoughts of neurophysically identical intentional systems in the same context are coextensive. Good. But, as I remarked above (see fn 3) you can have intentional (hence extensional) identity of thoughts without neurophysical identity of their thinkers. This suggests the following question: how much (and what kinds of) similarity between thinkers does the intentional identity of their thoughts require? This is, notice, a question one had better be able to answer if there is going to be a scientifically interesting propositional attitude psychology. For, generally speaking, the creatures that the generalizations of such a theory would have to subsume would not be anything like neurophysically identical; at a minimum they would differ as much as you and I differ from each other (at a maximum, they might differ as much as you and I differ from a silicon chip).

'What sorts of differences does intentional identity tolerate?' is the general form of the so-called 'collateral information' problem; and whereas I think that the Twin-Earth problem is, for the reasons I've set forth, something of a red herring, the collateral information problem is really very nasty and I don't know how to fix it; not, at least, in any detail. The next section is devoted to poking at the collateral information problem in the hope that something will come loose.

### 3. PUZZLES ABOUT CONTENT: COLLATERAL INFORMATION

Suppose Psmith (a neighbour, not a Twin) shares with me the belief that water is wet. And suppose we share all our other beliefs about water as well, except that he believes, and I do not, that cats like to drink water. This

Cambridge University Press  
 978-0-521-10844-7 - Language, Mind and Logic  
 Edited by Jeremy Butterfield  
 Excerpt  
[More information](#)

JERRY FODOR

supposition looks to be OK; I mean, if it is possible that two people should have all their beliefs about water in common, it is also possible that two people should have all their beliefs about water in common except for the belief that water is something that cats like. On the other hand, the following is perhaps not OK: Psmith and I share all our beliefs about water except that he believes that water is animate and I do not. Here one feels inclined to say that if somebody really believes that water is animate, then he has a different concept of water and hence cannot, *strictu dictu*, be said to share one's belief that water is wet. In the first case, one feels inclined to say, having the belief that cats like(/don't like) water is *merely collateral* to having the concept WATER; whereas in the second case, one feels inclined to say, the belief that water is (in)animate *partially determines* the identity of the WATER concept that one has; sharing the concept WATER doesn't survive disagreement about things like that.

The philosophical point, however, is that it doesn't make the slightest bit of difference what one feels inclined to say about the examples. What is wanted is a principled way of drawing the distinction. Unless we have that, we have no way of determining identity and difference of intentional content in any case that falls short of a Twin case; in any case that falls short of neurophysical identity, that is.

This, of course, is just Quine psychologized. Quine puts it that you can't distinguish between theory and language. Putnam, drawing out the implications for computational psychology, puts it that you can't distinguish between concept and collateral information. In either case, the putative moral is the same: individual symbols (including mental representations) 'are meaningful in the sense of making a systematic contribution to the functioning of the whole language; they don't have "meanings", in the form of isolable objects, properties or processes, which are associated with them individually and which determine individual assertability conditions' (Putnam 1984: 1). But if mental representations don't have meanings, beliefs don't have contents. And if beliefs don't have contents, there are no beliefs.

*Preliminary comment:* The present problem arises whether or not you think there is a defensible notion of *narrow* mental content; that is, whether or not it is possible to preserve the intuition that my water-thoughts are somehow intentionally identical to my Twin's. You might, for example, take the view: 'save the truth conditions and cognitive science be damned; there is no kind of belief content that supervenes on the neurophysical identity of believers'. Even so, one must decide whether someone who believes that water is animate can have beliefs about water. If, as Quine and Putnam suppose,