

Fodor's New Theory of Content and Computation

Andrew Brook and Robert J. Stainton

Interdisciplinary Studies and Philosophy

Carleton University

Ottawa, ON, Canada K1S 5B6

{abrook, stainton}@ccs.carleton.ca

Abstract

In his new book, *The Elm and the Expert*, Fodor attempts to reconcile the computational model of human cognition with information-theoretic semantics, the view that semantic content consists of nothing more than causal or nomic relationships between words and the world, and intentional content of nothing more than causal or nomic relationships between brain states and the world. We do not challenge the project, not in this paper. Nor do we show that Fodor has failed to carry it out. Instead, we urge that his analysis, when made explicit, turns out rather differently than he thinks. In particular, where he sees problems, he sometimes shows that there is no problem. And while he says two conceptions of information come to much the same thing, his analysis shows that they are very different.

1. Introduction

Two old friends show up early in Fodor's new book *The Elm and the Expert* (hereafter *E&E*). First old friend: psychology must employ intentional concepts such as belief and desire. Second old friend: cognitive processes consist of computations and "computational processes are ones defined over syntactically structured objects" (Fodor, 1994, p. 8). What's new is Fodor's view of intentional content. Narrow content is out; information theoretic semantics is in.¹

As we read it, *E&E* is about one central problem brought about by this change: whether the old picture of cognition as a computational process can be made to jibe with the new view that content is information, that content consists in a brain-world relationship. Could computational processes correctly 'track' content thus understood? We do not want to challenge the project. Nor do we aim to show that Fodor has failed to carry it out. Instead, we urge that his analysis, when made explicit, turns out rather differently than he thinks. First we will try to get the problem a little clearer, then we will consider peculiarities in Fodor's proposed resolution of it.

2. The Problem: The Argument for Incompatibility

By *computationalism* we mean the view that psychological states and processes are implemented computationally, where a computation is an operation over syntactic objects: a mapping from symbols to symbols, such that the transformations pay attention only to form, never to content. Let

intentionalism be the view that psychological states are ineliminably content-bearing. Intentionalism can be many different things, depending on how 'intentional' is read. In particular: combining intentionalism with the view that content is information yields what might be called *informational intentionalism*. To avoid such an ugly label, we will speak of 'info-intentionalism'. It's the view that psychological states and processes are ineliminably intentional, *in the sense of being information-bearing*.

With these rough and ready definitions in hand, let us now pose The Question: are computationalism and info-intentionalism compatible? Here's an argument, reconstructed from *E&E*, that they are not. We call it the Argument for Incompatibility.

Premise One: If psychological states and processes are ineliminably intentional and psychological states and processes are implemented computationally then there must be computationally sufficient conditions for the instantiation of intentional properties.

Premise Two: Content, being information, is relational.

Premise Three: If content is relational then there aren't any computationally sufficient conditions for the instantiation of intentional properties.

Now, the antecedent of Premise One follows directly from intentionalism, and therefore from info-intentionalism, when conjoined with computationalism. So, computationalism plus info-intentionalism plus Premise One entail that there *must* be computationally sufficient conditions for the instantiation of intentional properties. Premise Two *also* follows from info-intentionalism. And from Premises Two and Three it follows that there *aren't any* computationally sufficient conditions for the instantiation of intentional properties. Evidently, these two conclusions are inconsistent. It begins to look, then, as though one must give a negative answer to The Question: info-intentionalism and computationalism are not consistent.

¹ Incidentally, Fodor underplays the attractions of narrow content. Far from being some unwelcome stepchild to be embraced only in theoretic extremis, narrow content is an intuitively plausible notion. First, it seems that our contents stay with us no matter what causal environment we find ourselves in. Second, we are aware of many of our mental contents in a way that seems inconsistent with mental content being broad. There seems to be an interesting asymmetry here between mental content and semantic content, however. It is easier to think of the latter as relational than the former. Among other reasons, we often do not know the meaning of a word but it is not easy to think of ourselves not knowing the contents of our thoughts, desires, perceptions, etc.

But there may be hope yet for the view that cognition is a matter of computations which track information. Perhaps one could reject one of the Premises. Premise Two, we repeat, is entailed by info-intentionalism. And Premise Three is motivated as follows: computational properties, being syntactic, are internal. But no internal property is such that satisfying it is sufficient for having an external relation. (This is presumably what Fodor has in mind when he remarks that, "It's as though one's having ears should somehow guarantee that one has siblings." [Fodor, 1994, p. 14]) Applying this to the case of computational and intentional states/processes, it seems that computational properties cannot possibly guarantee intentional properties—the former are internal while (*ex hypothesi*) the latter are relational. Hence Premise Three. Given the solidity of Premises Two and Three, then, Fodor goes after Premise One.

The appeal of Premise One resides in our need for what Fodor calls "a property theory" connecting intentional laws with their computational implementations. As Fodor puts it,

If the implementing mechanisms for intentional laws are computational, then we need a property theory that provides for *computationally sufficient conditions* for the instantiation of intentional properties. (Fodor, 1994, p. 12; his emphasis)

To this Fodor gives a natural reply: the demand is too strong. In fact, he urges, psychological states and processes could be ineliminably intentional *and* be implemented computationally, even if there were no computationally sufficient conditions for the instantiation of intentional properties. If so, then despite the need for the theory to which Fodor refers, Premise One is too strong; indeed, Not True. In sum: the consequent of Premise One is inconsistent with the consequent of Premise Two, and the antecedents of both are true. But Premise One is not true. So, as far as the foregoing argument shows, the answer to The Question may well be 'yes': computationalism and info-intentionalism are compatible.

3. The Revised Argument for Incompatibility

So far so good. But a still small voice is insistently asking, 'Does the need for a property theory not commit us to *anything*?' Indeed it does. However, Fodor urges, all it commits us to is:

Premise One (Revised): If psychological states and processes are ineliminably intentional and psychological states and processes are implemented computationally, then the co-instantiation of the computational implementer and intentional implemented must be *reliable*.

Here is the argument that Premise One (Revised) is strong enough to allow for intentional psychological laws. One condition sufficient for computations to track content correctly would be supervenience: all differences of content being reflected in a difference within the computational system, across all possible worlds. That something this strong would *suffice* is presumably what motivates premises like our original Premise One. But, argues Fodor, all we *require* for psychology are conditions that reliably link

intentional content to syntactic states. Exceptions are quite all right so long as they *are* (just) exceptions, not counterexamples: that is to say, so long as they are infrequent and unsystematic—particularly unsystematic. (Psychology, goes the mantra, is a *ceteris paribus* science, not a basic one.) Moreover, these conditions need be reliable only in *this* world and worlds nomologically like this one. Psychology is beholden to worlds with the psychological (and related) laws of *our* world; other nomologies need not concern it. In sum: for a property theory to be in the offing, there must be *something* that keeps computation and content in phase, such that computational states/processes track intentional states/processes most of the time—but the tie need not be perfect nor hold in all possible worlds. These conditions can be met far short of sufficient conditions as conceived in classical conceptual analysis.

Unfortunately, the revised Premise One immediately suggests a revised Argument for Incompatibility. Frege cases (such as 'the Morning Star' and 'the Evening Star', 'Mark Twain' and 'Samuel Clemens') appear to be examples of computational type-distinctions to which no content distinction corresponds; Putnam's Twin Earth and Expert ('elm'/'beech') cases appear to be examples of content type-distinctions to which no computational distinction corresponds. (Fodor also discusses what he calls 'Quine cases': 'rabbit' and 'undetached rabbit part', for example. We'll introduce Quine cases later.) Because of Frege, Twin Earth, Expert and other cases, it is tempting to think that:

Premise Three (Revised): If content is relational then the co-instantiation of the computational implementer and intentional implemented will not be reliable.

Premise Two and Premise Three (Revised) entail that co-instantiation *will not be* reliable. Info-intentionalism, computationalism and Premise One (Revised) together entail that co-instantiation *must be* reliable. The ancillary premises seem solid. So it appears, once again, that info-intentionalism and computationalism are not consistent.

In response to this variant of the Argument for Incompatibility, Fodor goes after Premise Three (Revised). Roughly speaking, Fodor argues that—Frege, Twin, Expert and other cases notwithstanding—computational states/processes and intentional states/processes *do* reliably co-instantiate, at least when it matters. So the revised Argument for Incompatibility is also unsound. We turn, at last, to Fodor's defence of this claim.

4. Complications: Tying Content to Computation

When content is construed informationally—that is to say, as a matter of syntactic forms being in relationships to external objects—there are two broad ways in which content and computation could come apart, Fodor suggests.

1. There could be computational distinctions that do not reflect differences of informational content (Frege cases), and,
2. There could be differences of informational content that are not reflected in difference of computational state.

The latter in turn might happen in two ways:

2a. The difference of content is not available to the cognitive system, as in Twin Earth (and also Expert) cases. (In Expert cases it is available to an expert: Fodor cannot tell a beech from an elm but an expert can. In Twin Earth cases, it is not available to anyone.)

2b. The difference is available to the cognitive system, but it cannot be captured in a purely informational theory, as in the case of Quine's rabbit/undetached rabbit part.

For each of these ways in which disconnection might seem possible, 1, 2a, and 2b, Fodor tells us that there is a mechanism tying content and computation together. To explain what he has in mind, he offers an analogy. Why are appearing to be a dollar bill and actually being a dollar bill tied together—not perfectly but very, very reliably? Because of the mechanism of the police stamping out counterfeiting. We would then expect Fodor to say, 'And here are the mechanisms for content and computation for each of the three kinds of case'. But that is not what he does at all. Instead, the analysis goes off in a curious direction, indeed in two curious directions.

i. Instead of identifying three police-like mechanisms, Fodor offers us in effect *a series of explanations of why we do not need one*, at least in two of the three kinds of case. (The problem in his treatment of the third is even bigger, as we will see.)

The second curiosity arises from the fact that Fodor uses at least two variations on the content-as-information theme, which he calls causal and nomic. (We'll untangle these terms shortly.) But,

ii. The explanations mentioned in (i.) go through straightforwardly only for the causal variant.

4(i). Mechanism or Explanation?

The first curiosity is: instead of mechanisms, in two of the three cases Fodor offers us explanations of why we do not need any. These are the Frege and Twin Earth cases. (As we said, the problem with his treatment of Quine cases is different.) Let us begin with Frege cases. Take Oedipus and his unfortunate affair with his mother. Fodor says that, so far as *content* is concerned, the proposition thought by Oedipus to be true—that he was marrying Jocasta—and the proposition thought by Oedipus to be false—that he was (gasp!) marrying his mother—are the same: they have the same reference so carry the same information. Nevertheless, they're computationally distinct and clearly had or would have had different effects on Oedipus. If so, content doesn't map one-to-one with computational state.

But cases like Oedipus are not a problem, says Fodor, because they are *unsystematic accidents*; if it happened regularly that we did not know that co-referring terms referred to the same object, practical reasoning would become useless. Indeed, it takes such complicated circumstances to make a story like Oedipus's plausible that such cases would have to be rare. If these cases are rare accidents, however, all they show is that content and computational state are not perfectly linked. They do not show that the two are not *reliably* linked, and this is all we need. So far so good. But now ask: where's the mechanism? If acci-

dents are allowed, we do not need to block them. So we not need a content-computation mechanism; there would nothing for such a mechanism to do. Curious.

'Surely,' it will be protested, 'you have missed Fodor's point. We *also* need something to explain how it is that Oedipus cases *are* unsystematic, that we do generally know that our co-referring terms refer to the same object. That the mechanism Fodor has in mind'. That this is what Fodor has in mind is, at the very least, not obvious. Let us follow his analysis. He starts by saying that,

'Intentional systems' invariably incorporate mechanisms which insure that they generally know the facts upon which the success of their behaviour depends (Fodor, 1994, p. 48)

Is this the mechanism we are looking for? No; for "rational agents" to "reliably make a point of knowing the facts the success of their behaviour depends upon" (Fodor, 1994, p. 46), the facts have to cooperate. We have to have a poor Oedipus lacked, information adequate to know the truth of the relevant identity-statements. As it happens, agents generally do, because "the syntactic structure of a modal presentation reliably carries information about its causal history" (Fodor, 1994, p. 54). And do they also carry information about other syntactic structures, that they refer to the same object?, we might ask. But that is not the question. The vital question is this: Does any mechanism guarantee this reliability? Indeed, does anything guarantee it? Fodor says nothing to suggest that it is more than a happy accident, though one essential to practical reason indeed probably to life itself.

If there is a mechanism producing the happy covariation that are supposed to allow us to know, most of the time, the identity of reference of our coreferring expressions, the candidate that we can think of is natural selection. In such contexts, Fodor expresses a distinct lack of enthusiasm for evolutionary arguments (1994, p. 20 for example), but may be one place where they can do some work. What it would be peculiar to call anything that could be involved here a *mechanism*, certainly if the police arresting counterfeiters is an example of a mechanism. Compare: 'Aren't there lots of elephants on Lake Ontario today? Well, you see, there's a mechanism...'

We discover the same pattern even more clearly in Fodor's treatment of Twin Earth cases. Fodor's response there is to claim that they do not occur in our world will not occur in any nomic near one (Fodor, 1994, 38-9). He may well be right; but if he is, we do not need any *mechanism* to deal with them. Once again, instead of identifying a mechanism, Fodor has shown that we do not need one.

That leaves the Quine cases. (We will not consider Expert cases.) They need a word of explanation. The problem with 'rabbit'/undetached rabbit part ('urp') cases is that information semantics Fodor-style is important. The terms 'rabbit' and 'urp' clearly have different contents. Indeed since a rabbit is more than an urp, if one of them correctly applies, then the other does not. Yet they are always instantiated. Given that, the *information* contained in 'rabbit' and 'urp' is the same, on Fodor's stringent notion

information. From this it follows that no semantics based purely on such a notion of information is going to capture the difference of content between the two terms. Fodor's suggestion is that we can capture the difference between 'rabbit' and 'urp' if we add a notion of "inference potential" (shades of conceptual role!). In particular, by checking whether an Informant (Inf) will accept or reject certain conjunction reductions, we can tell whether she uses 'rabbit' to mean *rabbit* or *urp*, 'urp' to mean *urp* or *rabbit*.

Here is how the story goes. (Fodor tells it in terms of triangles and triangle parts but we will tell it in its rabbit and urp version.) In addition to a rabbit and an urp, consider also, say, the front half of a rabbit. An appropriately located urp could be part of both a rabbit and the front half of a rabbit (an undetached ear, eye, or nose would be some examples). Now we ask, does 'rabbit' mean *rabbit* or *urp*? Does 'urp' mean *an undetached rabbit part* or *rabbit*? Being the front half of a rabbit excludes being a rabbit but being an (appropriate) urp is compatible with being part of both a rabbit and the front half of a rabbit. Thus, if 'rabbit' and 'front half of a rabbit' mean *rabbit* and *front half of a rabbit*, Inf will accept 'A is a rabbit' and 'A is a front half of a rabbit' in certain circumstances but never, it seems, 'A is a rabbit and a front half of a rabbit'. On the other hand, if 'rabbit' and 'front half of a rabbit' mean *rabbit part* and *part of the front half of a rabbit*, Inf will accept 'A is a rabbit' and 'A is a front half of a rabbit' and also 'A is a rabbit and a front half of a rabbit' in certain circumstances. So all we have to do is check and see which conjunction reductions Inf accepts and we can determine what she means by 'rabbit' and 'urp' (1994, p. 73). So far so good.

But now a problem appears: What in Inf's *computational structure* could implement the difference of content between 'rabbit' and 'urp'—as the case has been described? It cannot be anything informational because Inf's information about rabbits and about urps is the same; 'rabbit' and 'urp' covary perfectly with both rabbits and urps. So it has to be something syntactic. And the problem is this: what could anything purely syntactic have to do with content here? Notice: here Fodor is not demanding merely that syntax *covary* with content. Here syntax is to be *part* of content.

In the story about 'rabbit', 'urp', and 'front half of a rabbit' we told, most people would be inclined to think that Inf's conjunction reduction predilections are driven by information: she reduces or refuses to reduce in accord with what satisfies these terms and her beliefs about these items. But that cannot be Fodor's story; for him, there is *no difference* in Inf's information about rabbits, front halves of rabbits, and urps. So what drives Inf to reduce or not reduce has to be something purely syntactic, some computational rules for uninterpreted symbol-sets in Mentalese buried in her. This sharpens the problem. What could anything purely syntactic in this way have to do with content?

We think the most probable answer is: nothing. If what governs Inf's reduction predilections is something purely syntactic, rules for the manipulation of uninterpreted symbols, Inf could just as well reduce sentences the *content* of whose predicates are incompatible as compatible, sen-

tences whose subjects are *satisfied* by different objects as by the same one. In short, on a purely syntactic account, conjunction reduction seems to tell us *nothing* about content. And the Quine cases remain, alive and well.

Does Fodor have any escape? If conjunction reduction proclivities are to be part of content, they must at least covary with the relevant differences of content. We can think of two possibilities:

(i) Fodor could simply *define* incompatibility of satisfaction as an Inf's unwillingness to conjunction reduce. This seems a desperate expedient: what does the fact that 'rabbit' is satisfied only when 'urp' is not have to do with some rule for transforming uninterpreted symbols buried deep in the brain? But out of desperation grows an idea. What if,

(ii) we could find a *mechanism* that tied the computational reduction to the semantic difference? This should be the escape Fodor wants, given the rest of *E&E*. Can we think of one?

It does not look promising. Strangely enough, Fodor is no help here; indeed, he does not even mention a mechanism in connection with Quine cases! This gap and the difficulty of bridging it may be an instance of a wider problem. At the beginning of *E&E*, Fodor tells us that if computations are to house content, computations must track truth in the way that our reasonings do. To see the bigger problem first, suppose that Fodor is right about everything and there is no way in which symbol structures and contents might systematically come apart, not in this and near worlds at any rate. Would that be enough to ensure that computation will *track truth* reliably? It is not clear that it would; even if computational and contentful *nodes* line up correctly, the *relationships among them* could still come apart.

Now apply this bigger problem to conjunction reduction and the problem of covarying terms that are not jointly satisfiable. The problem here is to find something to ensure that a purely computational move tracks mutual nonsatisfiability correctly. It's not the same problem, but it is a closely related one. And Fodor says as much about it, namely, nothing.

This problem is far more serious than the problems we identified earlier in connection with Frege and Twin Earth examples. There Fodor promised us mechanisms and instead gave us arguments for why we don't need them—a fairly parochial failing. Here he produces no mechanism *and* no argument that we do not need one. This looks more serious.

The source of the problem seems to be as follows: To discriminate rabbits from urps, we would have to use something that comes to us from these items. Or so it would seem. But Fodor thinks that there is no such information. So he has to go shopping for something purely syntactic, purely internal. Once we see this, a solution to the Quine cases becomes readily apparent. The relationships of rabbits and our syntactic structures and urps and our syntactic structures differ, in lots of relevant ways, in fact, and ways as naturalistic as covariance; all the *causal* relationships are different, as different as the very different causal powers of rabbits and urps. These differences are just as good candidates to ground differences of information in connection with 'rabbit'

and 'urp' as co-variance, the one Fodor chooses.

Fodor cannot avail himself of this solution only because he adopts an excessively stringent notion of information. For him, the *only* way in which the information carried by two terms can vary is if there are contexts in which the terms do not co-vary; and the only way they can do that is if they refer to objects that do not always co-vary. In short, Fodor seems to think that the only link between words and the world that counts is covariance. Content is also determined by what *undergirds* the covariance, however: 'dog' has the content it has not simply because it co-varies with dogs, but because tokens of it are *caused* by dogs (see Quine, 1960, p. 30). This is entirely compatible with a naturalistic, externalist account of content. Assuming—as seems plausible—that it is the causal powers of *rabbits* that maintain the co-variance between 'rabbit' and rabbit stuff and that the causal powers of *urps* are relevantly different, 'rabbit' will mean *rabbit*, not *urp*. In which case, 'rabbit' meaning *urp* is excluded without appeal to conjunction reduction or any other kind of 'inference potential'. And the Quine cases would no longer be a problem. True enough, content so construed will go beyond covariance, but we cannot imagine why this should concern Fodor.

The introduction of 'inference potential' of any kind as an element of content also raises troubling issues all its own. For example, though Fodor argues that this move invites only a benign form of holism, it's surely very hard to help oneself to just a soupçon of conceptual role semantics. Fodor says that he can isolate the "logical syntax" of conjunction reduction from the rest of language, but his treatment of this seems a little blasé. In particular, to make the separation, Fodor would have to be able to separate conjunction-reductions based on syntax from reductions based on evidence, i.e. purely on information. Perhaps he can; but given Quine's worries about the very possibility of doing such things, we'd like to see the argument.

4(ii). Variations on the Notion of Information

Above we urged that, just when Fodor seems poised to describe the police-like 'mechanism' that keeps content and computation reliably in phase, his discussion goes off in peculiar directions. Having explored one of them—no mechanism, just explanation, and in Quine cases no explanation either—, we turn now to the other one. Fodor acknowledges two variations on the content-as-information theme: causal and nomic. Here's what's strange: the explanations Fodor gives in connection with Frege and Twin Earth cases go through straight-forwardly only for the first.

Fodor calls the two conceptions the *causal* informational and the *nomic* informational. Notice first that they are both quite different from the biographical-historical conception propounded by Dretske (1993) and others. Dretske argues that AI systems do not and could not have content, intentionality, etc., simply because they have the wrong kind of history. On this view, even if a system were to be built, to the appropriate fineness of grain, exactly like us in all relevant respects, and even if it behaved exactly like us, the difference in its history would ensure that it does not have content—even though we do. For Fodor, however, history

of the system has little to do with content.

Fodor calls the causal informational notion of content the view that, "the content of mental representations is constituted by their etiology" (Fodor, 1994, p. 82). As we've seen, Fodor goes a long way toward showing that this simple little notion is adequate (though perhaps not quite in the direction he thinks). However, there is also a nomic notion of content-as-information at work in the book.

On the causal story, a computational state comes to carry information by entering a causal relationship with some object: a token of 'dog' comes to carry the information contained in the concept DOG by being in a causal relationship with a dog. On the *nomic* informational story, in contrast, a computational state comes to carry information by satisfying certain counter-factuals: a token of 'dog' as found in me carries the information *dog*, for example, if I *would* say or otherwise token 'dog' *were I* to be in the presence of a dog. It is not necessary that my token of 'dog' was ever *actually* in a causal relationship with a dog.

On the face of it, these seem to be quite different conceptions of content. About all that immediately holds them together is that they are both developed in terms of information. Strangely enough, Fodor seems to think that they are quite similar—similar enough to be roughly substitutable one for the other. Thus, he tells us at one point that he prefers the nomic-informational notion but for the sake of simplicity of exposition has told his story in terms of the causal-informational one (Fodor, 1994, p. 54) and from then on switches back and forth between them pretty much at will, embracing the causal story on pp. 52, 82 and 86, and the nomic story on pp. 72, 90, and 116, for example. (Fodor even makes use of concepts that could only be part of the nomic story before he formally introduces the notion. For example, he speaks of content as dispositions on p. 30 and as specified by counter-factuals on p. 37.)

Contrary to Fodor and anyone else who may suppose that these notions of broad content are quite similar, the two are very different. Consider Davidson's (1987) old friend the Swampman, the molecule for molecule duplicate of Davidson who springs to life ('life?') one day in a swamp. Intuitively, Swampman seems to have content at the very start of his new life, i.e., prior to entering into causal relationships with objects. Now, this is also what the nomic story would entail, and it's the story Fodor embraces. However, on the *causal* story, we would have to deny that the Swampman had content—and a number of theorists have done so. This is enough by itself to show that the two conceptions are quite different from one another. We so conclude; and turn to the implications for Fodor.

Fodor, as we've seen, argues that—Twin, Frege and Quine cases notwithstanding—the link between content and computational implementation is reliable. However, the nomic story goes quite differently for some of these cases from the causal one. Fodor gives no indication that he sees this.

However well his story about why Frege cases are no problem works on the causal informational account, for example, it does not work at all on the nomic informational version. Here is why. If content is a matter of actual causal

connection to an object, then it is perfectly possible to locate what is common to Oedipus's two beliefs—it consists in their being linked causally to the same object. Then this becomes the content of both 'Jocasta' and 'Mom'. Because these two are syntactically distinct, we have a case of computation-content disconnection, but that's okay because it's accidental. Treat content as a matter of some kind of disposition, however, and this story utterly collapses. What could be *common* to the disposition activated in Oedipus by the woman presented as Jocasta and the one activated by the woman presented as Mom? Certainly not a disposition to get married! And not, to be more serious, the respective word-tokening dispositions either. Even more seriously, what motive could one have for splitting the dispositions into two elements in the first place, a part common to both and a part distinctive to each? In so far as there is no good answer to this question, Fodor's story about Frege cases does not work on the nomic account of information.

What about Twin Earth cases? Again a problem. Where information derives from *etiology*, the apparent difference in content between (say) Adam and Twadam is that Adam's tokens of 'water' are causally linked to H₂O, Twadam's to XYZ—a substance phenomenally indistinguishable but chemically (or something) distinct from H₂O. Here there is a clear sense in which the two tokens of 'water' have different contents, even though elicited by indistinguishable environments: one is causally linked to H₂O, the other to XYZ. Different contents, same computational state. (And, says Fodor, it doesn't matter because Twin Earth isn't 'nearby'.)

As with Frege cases, however, things turn sour when we go nomic. On the nomic story, the difference in content between Adam and Twadam has to be that Adam would token 'water' in the presence of H₂O while Twadam would token 'water' in the presence of XYZ. The trouble is, Adam's disposition would also lead him to token 'water' in the presence of XYZ and Twadam's to token 'water' in the presence of H₂O. That is to say, on the nomic theory, there is no difference of content between Adam and Twadam in the first place. Which would imply that, for the nomic version of informational content, Twin Earth cases do not pose an even apparent problem for Fodor in the first place.

For Twin Earth cases, then, things are easier for Fodor on the nomic view than the causal one. But if he wants his story of Frege cases to work for the nomic view as well as the causal one, he owes us an argument. Fodor's treatment of Quine cases seems to work equally well, or equally badly, on either view.

In sum: Fodor's strategy in *E&E* is to urge that psychological states and processes can be both ineliminably intentional (information-bearing) and implemented computationally—without there being computationally sufficient conditions for the instantiation of intentional properties. All we need is that computational states/processes reliably co-instantiate with intentional states/processes in our world and worlds nomically near to ours. Which, despite the problems with Fodor's arguments, they may do. But Fodor argues for the thesis in some strange ways and owes us some additional arguments.

Acknowledgements

Andrew Brook would like to thank his fellow participants in a discussion group on *The Elm and the Expert*, The Queen's College, Oxford, Trinity Term 1995, for many important ideas: Alex Rosenberg, Galen Strawson, Michael Lockwood, David Bakhurst, and Frank Jackson. Robert Stainton would like to thank Tracy Isaacs, Daniel Stoljar, and the students in his Mental Representation tutorial at Carleton University for helpful comments. Thanks also to the participants in our workshop on Computationalism and Intentionalism in the Philosophy of Mind, Canadian Philosophical Association, Université du Québec à Montréal, June 3, 1995.

References

- Davidson, D. (1987). Knowing One's Own Mind. In *Proceedings and Addresses of the American Philosophical Association* 61:441-458.
- Dretske, F. (1993). The Possibility of Artificial Intelligence. Presented to the American Philosophical Association, Eastern Division, December 28, 1993.
- Fodor, J.A. (1994). *The Elm and the Expert*. Cambridge, MA: The MIT Press.
- Fodor, J.A. (1987). *Psychosemantics*. Cambridge, MA: The MIT Press.
- Quine, W.V.O. (1960). *Word and Object*. Cambridge, MA: The MIT Press.