

SSF: [Still Small Voice—could it be the voice of conscience?] I do believe you’ve gone over to Steve Stich.

Have you no conscience?

Answer: There, there; don’t fret! What is emerging here is, in a certain sense, a “no content” account of narrow content; but it is nevertheless also a fully intentionalist account . . . In effect I’m prepared to give Stich everything except what he wants.

FODOR, *PSYCHOSEMANTICS*

8

NARROW CONTENT MEETS FAT SYNTAX

1 Introduction

A bit over a decade ago I published a paper in which I argued that Putnam’s much discussed Twin Earth thought experiments posed a problem for the view that a psychological theory aimed at explaining human behavior will invoke common-sense intentional concepts like belief and desire.¹ That argument relied on a pair of premises. The first, which I (perhaps infelicitously) called the *principle of psychological autonomy* maintains that any state or property properly invoked in a psychological explanation should supervene on the current, internal, physical state of the organism. Thus, a pair of Putnamian doppelgangers, being molecule for molecule replicas of one another, must share all the same explanatory psychological states and properties. The second premise was that commonsense intentional properties, properties like *believing that Eisenhower played golf* (or *having a belief with the content that Eisenhower played golf*) and *believing that water is wet* (or *having a belief with the content that water is wet*) do not supervene on a person’s current, internal, physical state. For want of a better label, I’ll call this the *Autonomy* argument.

The first premise of the Autonomy argument was one that I took to be intuitively obvious and widely shared. Thus I offered little by way of support. The second premise seemed to be a straightforward consequence of the usual intuitions about Twin-Earth

1. Stich (1978). For The details of Putnam’s thought experiment, see Putnam (1975a).

style thought experiments. According to those intuitions, the belief that my doppelganger expresses when he says “Eisenhower played golf” is not *about* Eisenhower, the man whose hand I almost got to shake during the 1956 presidential election; it is about some other statesman in a far off corner of the universe. Thus the truth conditions of my belief and my doppelganger’s are different. But it is plausible to suppose that on the conception of content implicit in commonsense psychology, belief tokens that are about different people and that have different truth conditions must have different contents. So my doppelganger and I do not both have beliefs with the content that Eisenhower played golf.

This argument was part of a larger project. Influenced by Quine, I have long been suspicious about the integrity and scientific utility of the commonsense notions of meaning and intentional content. This is not, of course, to deny that the intentional idioms of ordinary discourse have their uses, nor that these uses are important. But, like Quine, I view ordinary intentional locutions as projective, context sensitive, observer relative, and essentially dramatic.² They are not the sorts of locutions we should welcome in serious scientific discourse. For those who share this Quinean skepticism, the sudden flourishing of cognitive psychology in the 1970s posed something of a problem. On the account offered by Fodor and other observers, the cognitive psychology of that period was exploiting both the ontology and the explanatory strategy of commonsense psychology. It proposed to explain cognition and certain aspects of behavior by positing beliefs, desires, and other psychological states with intentional content, and by couching generalizations about the interactions among those states in terms of their intentional content.³ If this was right, then those of us who would banish talk of content in scientific settings would be throwing out the cognitive psychological baby with the intentional bath water. On my view, however, this account of cognitive psychology was seriously mistaken. The cognitive psychology of the 1970s and early 1980s was not positing contentful intentional states, nor was it adverting to content in its generalizations. Rather, I maintained, the cognitive psychology of the day was “really a kind of logical syntax (only psychologized).”⁴ Moreover, it seemed to me that there were good reasons why cognitive psychology not only did not but *should* not traffic in intentional states. One of these reasons was provided by the Autonomy argument.

2. Quine has urged this view of the propositional attitudes in many places. See, for example Quine (1960:219). For some elaboration on these themes, see Stich (1982); Stich (1983:chs 4–6); Gordon (1986); and Levin (1988).

3. See, for example, Fodor, (1975:ch. 1); Fodor (1980); Fodor (1981: “Introduction”); and Fodor (1987:ch. 1).

4. The quote is from Fodor (1978). For my account of the explanatory strategy of cognitive psychology, *circa* 1980, see Stich (1983:chs 7–9). Perhaps this is the place to say that when I talk of the cognitive psychology of the 1970s and early 1980s, what I have in mind is pre-connectionist cognitive psychology. The qualification is important since, on my view, neither Fodor’s account of cognitive theorizing nor my syntactic account will mesh comfortably with the connectionist paradigm. For some elaboration of this point, see Ramsey, Stich, and Garon (1990).

During the last decade, that argument and similar arguments offered by other writers have attracted a fair amount of attention, very little of it favorable.⁵ Some critics have focused on the first premise, and have argued that explanatory psychology need not, and does not, restrict itself to states and properties that organisms and their doppelgangers share.⁶ Others have focused on the second premise, with some arguing that commonsense psychology does not insist that beliefs with different truth conditions differ in content, or at least that it does not do so consistently, while others challenged the intuition that the beliefs of doppelgangers on Earth and Twin-Earth differ in truth conditions.⁷ I think each of these objections raises serious issues, and each merits a detailed reply. But in the present paper I'll say very little about them. My focus here will be on quite a different reaction to the Autonomy argument—a reaction which grants both premises of the argument. This reaction concedes that the *commonsense* notion of intentional content will not play a role in scientific psychology. But it insists that *another* notion of intentional content will be central to psychology. For this second, more technical, and less commonsensical notion of content, it is not the case that if a pair of belief tokens differ in truth conditions, or in what they are about, then they also differ in content. Thus Twin Earth cases and others of their ilk will not show that *this* sort of content does not supervene on the current internal state of the organism. Though my doppelganger and I have beliefs that are about different people (or stuff) and thus have different truth conditions, those beliefs may still have the same content, when content is construed in this new way. Since the ordinary notion of content determines truth conditions—typically conditions in the world beyond the head—while the new technical notion does not, the new notion has been dubbed *narrow content*; the old commonsense notion is often said to be *broad* or *wide*. There are various lines along which the narrow content response to the Autonomy argument can be developed.⁸ But, as is appropriate in a volume focused on Fodor, the line I propose to explore is the one that Fodor follows. In section 3.1, I'll give a quick overview of Fodor's account of narrow content.

An objection often urged against the notion of narrow content is that it is not really a species of content at all.⁹ One reason for this suspicion is that while it is generally easy to *say* what the (ordinary, broad) content of a belief is, there often seems no way at all to say what the narrow content of a belief is. Narrow content appears to be “radically inexpressible.” However, I will argue that this suspicion is mistaken. Indeed, in section 3.2 I will sketch a straightforward way in which readily available resources can be used to construct a vocabulary for attributing narrow content. Of course, this alone is not enough to show that narrow content really is a kind of content, properly so-called. And I must confess

5. For similar arguments, see Stack (unpublished); Putnam (1978); and Putnam (1983).

6. See, for example, Burge (1979); Burge (1986); Kitcher (1985); Owens (1987); Baker (1987a).

7. Loar (1987); Lycan (1988: 76–9); Dow (1991).

8. See, for example, Block (1986); Dennett (1982); Devitt (1989); Dow (1991); Loar (1987a).

9. See, for example, Owens (1987) and Baker (1987b).

that I'm not at all sure what it *would* take to show that narrow content is, or isn't, really a kind of content. So I propose to leave that question to be debated by those who think they understand it. As I see it, the major objection to narrow content, as Fodor develops the notion, is that it is very unlikely to be of any more use to psychology than the commonsense notion of broad content. If we taxonomize mental states by their narrow content, there are going to be lots of psychological generalizations that we are not going to be able to state. My argument for this claim is set out in section 3.3.

That argument presupposes a certain conception of the cognitive mind—a conception that portrays the mind as analogous to a kind of computer. Though very familiar, this picture of the mind has never been without its critics, and with the recent flowering of connectionism it has become particularly controversial. However, in the present paper I don't propose to challenge the picture. Since Fodor himself has long been one of its most eloquent advocates, I will simply accept it, if only for argument's sake. In order to launch my argument against narrow content, it will be necessary to sketch in parts of the picture with somewhat more detail than is usually provided. This is the project I'll pursue in section 2.

Before getting on to any of this, however, we would do well to get a bit clearer about the issue that is in dispute. In the article in which I first set out the Autonomy argument, and in various subsequent publications, my "official" thesis was that serious scientific psychology should not invoke commonsense intentional notions like belief and desire. The official thesis certainly does not entail that beliefs, desires, and other propositional attitudes do not exist,¹⁰ nor even that commonsense psychology is not "pretty close to being true"¹¹—though it is, of course, consistent with these claims. But it is these claims that are at the heart of Fodor's concern. On his view, "if commonsense intentional psychology really were to collapse, that would be, beyond comparison, the greatest intellectual catastrophe in the history of our species . . ." (*PS*:xii). This leaves us with a rather delicate question. Just what would it take to show that commonsense intentional psychology had collapsed? Nobody thinks that *all* of commonsense psychology is going to turn out to be correct. Indeed, Fodor cheerfully concedes that "a lot of what common sense believes about the attitudes must surely be false (a lot of what common sense believes about *anything* must surely be false)" (*PS*:15). He also concedes that "you can't make respectable science out of the attitudes as commonsensically individuated" (*PS*:30). The "identity conditions for mental states" that "we need, when doing psychology" are not going to be "those that common sense prefers" (*PS*:30). If all of this is not enough to undermine commonsense psychology and its intentional ontology, one might well wonder how much more it will take. Fortunately, Fodor tells us. He stipulates that a psychological theory will count as "endorsing" commonsense propositional attitudes "just in case it postulates states (entities, events, whatever) satisfying the following conditions:

10. On this point see Stich (1983:ch. 11 sec. 1).

11. Fodor (1987:x). Subsequent references to Fodor's *Psychosemantics* will be referred to as *PS* in the text.

- (i) They are semantically evaluable.
- (ii) They have causal powers.
- (iii) The implicit generalizations of commonsense belief/desire psychology are largely true of them.

In effect,” Fodor tells us, “I am assuming that (i)–(iii) are the essential properties of the attitudes. This seems to me intuitively plausible; if it doesn’t seem intuitively plausible to you, so be it. Squabbling about intuitions strikes me as vulgar” (*PS*: 10).

I am not at all sure whether my intuitions agree with Fodor’s here; indeed, I’m not even sure I *have* any intuitions about the essential properties of the attitudes. But no matter. This book is for Fodor; I’ll play by his rules. What I propose to argue is that most of the implicit generalizations of commonsense psychology are not likely to turn out to be true of the states posited by psychological theories that cleave to the computational paradigm. So Fodor loses on (iii). Moreover, on at least one plausible reading of what it is to be “semantically evaluable,” these states are not semantically evaluable either. Thus Fodor loses on (i) too. Whether or not we accept Fodor’s intuitions about what is essential to the attitudes, this should be enough to show that propositional attitude psychology is in trouble.

2 The Computational Paradigm

My goal in this section is to provide a brief sketch of a familiar story about the cognitive mind. Since the basic outline is so well known, I will devote most of my attention to clarifying the ontological underpinnings of this account and the taxonomic strategies it exploits. Much of what I say in this section is based on the rather more detailed account I developed in *From Folk Psychology to Cognitive Science*. Since talk of *states* and the various ways in which they get taxonomized or individuated is going to be of some importance in what follows, I’ll begin by making a few proposals about how this talk should be construed. So far as I can see, nothing in the arguments to follow depends on the details. We just need *some* systematic way of talking about states. Most any sensible proposal would do.

As I propose we view them, states are the instantiation of a *property* by an *object* during a *time interval*.¹² There are, of course, venerable disputes about what sorts of things properties are.¹³ But for present purposes I propose to be quite permissive. Near enough, I’ll count any open sentence with a reasonably clear extension as specifying a property. That raises the notorious question of when two open sentences specify the same property. Fortunately, this is not a question for which we will need any fully general answer. All

12. My account of states is modeled on Kim’s account of events. See Kim (1969) and Kim (1976).

13. See, for example, Armstrong (1978).

we'll need is the weak principle that open sentences with different extensions specify different properties.

On the view I'm recommending, states count as *particulars* with a more or less definite location in space and time. States also admit of what might be called an *essential* classification into types. A pair of states are of the same *essential type* if and only if they are instantiations of the same property. Although each state has only one essential type, states, like other particulars, can be grouped into nonessential types in an endless variety of ways. A type of state is simply a category of particulars, and we have specified a type when we have set out conditions for membership in the category. Though we are conceiving of states as particulars, it will sometimes be convenient to use the word "state" to talk about a type or category of states, or the property that members of a category have in common. When ambiguity threatens, I'll use "state token" to refer to particulars and "state type" to refer to categories or types.

So much for states. Let me turn, now, to the story about the mind that I have been calling the *computational paradigm*. The central assumption of the story is that the cognitive mind can be viewed as a particular kind of computer—that the mind is, in Fodor's phrase, "a syntax-driven machine" (*PS*: 20). On this view, each cognitive state token is a brain state token—its essential type is determined by some neurophysiological property or other. However, these neurophysiological state tokens can also be viewed as having syntactic structure in something like the same way that sentence tokens in a natural or formal language have a syntactic structure. That is, each cognitive state token can be viewed as belonging to a syntactic type (or having a "syntactic form"), just as each inscription of a sentence in English or in first order predicate calculus can be viewed as having a syntactic form. Cognitive processes consist of temporal sequences of these syntactically structured states. The reason that the cognitive mind can be thought of as a kind of computer is that the mechanism that controls these cognitive processes is "sensitive solely to syntactic properties" (*PS*:19).

This account of the cognitive mind as a computer or a "syntactic engine" has become very familiar in recent years. But, as Michael Devitt notes in a recent article,¹⁴ the account is very easy to misconstrue. Often, when offering quick sketches of the mind-as-computer story, writers will conjure the image of a "belief box" and a "desire box" inside the head in which syntactically structured sentence-like entities are stored. For vividness, it may even be suggested that the sentences be thought of as well formed formulas of some familiar formalized language.¹⁵ But, as Devitt notes, this image invites us to think of the syntactic properties of cognitive state tokens (the properties in virtue of which they fall into one or another syntactic category) as *intrinsic* or "brute physical" properties—properties that we could detect if we looked at the appropriate bits of the brain in isolation, much as we

14. Devitt (1989).

15. For a particularly vivid and influential example of the Belief-Box metaphor, see Schiffer (1981).

could see whether an inscription in a “belief box” had the shape: $(x) Fx \rightarrow Gx$. If we think of the syntactic properties of mental states in this way, then it would make perfectly good sense to suppose that in certain brains syntactically structured states might be stored in the “belief box,” though the mechanisms which control cognitive processes are *not* sensitive to the syntax. But, along with Devitt, I would urge that this is just the wrong way to conceive of things. Mental state tokens are brain state tokens. But the properties in virtue of which mental state tokens are classified into syntactic categories are not intrinsic features of those brain states; they are not features which depend exclusively on the shape or form or “brute physical” properties of the states. Rather, the syntactic properties of mental states are relational or functional properties—they are properties that certain states of the brain have in virtue of the way in which they causally interact with various other states of the system. To put the point in a slightly different way, we would have no reason to view brain states as syntactically structured unless that structure can be exploited in capturing generalizations about the workings of mind/brain’s mechanisms. Attributing syntactic structure to brain state tokens—assigning them to syntactic types—is justified only if some interesting set of causal interactions among those tokens is isomorphic to formal relations among abstract syntactic objects. Here is how I elaborated on this theme in *From Folk Psychology to Cognitive Science*:

The basic idea . . . is that the cognitive states whose interaction is (in part) responsible for behavior can be systematically mapped to abstract syntactic objects in such a way that causal interactions among cognitive states, as well as causal links with stimuli and behavioral events, can be described in terms of the syntactic properties and relations of the abstract objects to which the cognitive states are mapped. More briefly, the idea is that causal relations among cognitive states mirror formal relations among syntactic objects. If this is right, then it will be natural to view cognitive state tokens as tokens of abstract syntactic objects . . .

The theorist’s job in setting out [this sort of] cognitive theory can be viewed as having three parts. First, he must specify a class of [abstract] syntactic objects . . . and do so in a way which assigns a formal or syntactic structure to each of these objects . . .

Second, the theorist hypothesizes that for each organism covered by the theory, there exists a set of state types whose tokens are causally implicated in the production of behavior. He also hypothesizes that there is a mapping from these state types to syntactic objects in the specified class. Several observations about these hypotheses are in order. First, the theorist need say very little about the essential nature of the state tokens which are causally implicated in the production of behavior. Presumably they are physical states of the brain, and thus the properties which constitute their essential types are neurological properties . . . Second, in asserting the existence of the mapping, the order of the quantifiers is of some importance. The theorist is not claiming that the mapping is the same for each subject, but only

that for each subject there is a mapping. So in different subjects, quite different neurological states types may be mapped to a given syntactic object. These . . . two points . . . are in the spirit of functionalism, which stresses the possibility of multiple realizations of mental states . . .

The third part of [this kind of] cognitive theory . . . is a specification of the theory's generalizations. The core idea . . . is that generalizations detailing causal relations among the hypothesized neurological states are to be specified indirectly via the formal relations among the syntactic objects to which the neurological states are mapped. Similarly, generalizations specifying causal relations between stimuli and neurological states will identify the neurological states not by adverting to their essential neurological types but, rather, by adverting to the syntactic objects to which the neurological types are mapped. Ditto for generalizations specifying causal relations between neurological states and behavior.¹⁶

As Devitt rightly points out, there is a certain tension in this passage that emerges when we ask how we would go about determining whether a pair of brain state tokens in two different people (or in one person at two different times) are tokens of the same syntactic type. One criterion for the syntactic type identity of tokens would require only that the tokens' patterns of causal interactions *with other tokens* be pretty much the same, so both patterns could be captured by the same formal relations among the appropriate system of syntactic objects. A more stringent criterion would require not only that the tokens' patterns of causal interactions with each other be the same, but also that their patterns of causal interaction with *stimuli* and *behavior* be pretty much the same as well. Since the terms "broad" and "narrow" have been appropriated for distinguishing kinds of content, I will call these two standards for determining the syntactic type of a hypothesized brain state token *skinny* and *fat* respectively. Though my writing has sometimes been less than clear on the point, it has always been my intention to invoke *fat syntax* in typing mental state tokens. When Fodor describes the mind as a "syntax-driven machine" it is not clear whether the standard of syntactic type individuation he has in mind is fat or skinny. In what follows, I'll assume that the syntactic types exploited in computational theories of the mind are fat, not skinny, though most of my argument will work either way.

3 Narrow Content

So much for the computational paradigm. Let's now return to the Autonomy argument, and Fodor's strategy for dealing with it. Since the notion of "narrow" content plays a central role in that strategy, I'll start with a sketch of how Fodor proposes to construct the notion. Once that's been done, I'll set out a pair of reasons for doubting that Fodor's

16. Stich (1983:149–51).

notion of narrow content will do what he wants. One of these, I'll argue, is pretty easy to handle. The other is not.

3.1

Mental states, we are supposing, are states of the brain. And, while their essential type is neurophysiological, they can also be classified into all sorts of other categories. One such categorization, provided by commonsense psychology, is to type mental state tokens by their content. The problem posed by the Autonomy argument is that the taxonomy imposed by ordinary, "broad" content does not supervene on a person's current, internal, physical properties. So while those states in Fodor's brain which count as beliefs and those in Twin-Fodor's brain which count as beliefs are neurophysiologically the same, they may well differ in content. This difference in content, Fodor notes, must be due to differences in the world around them and their relations to that world.

Presumably . . . there's something about the relation between Twin-Earth and Twin-Me in virtue of which his "water"-thoughts are about XYZ even though my water-thoughts are not. Call this condition that's satisfied by (Twin-Me, Twin-Earth) condition C (because it determines the *Context* of his "water"-thoughts). (*PS*:48)

Fodor's proposal for constructing a notion of narrow content is to start with the taxonomy provided by the ordinary, broad, truth-condition determining notion of content, and subtract out the contribution of the contextual conditions, like condition C, that "anchor" it.¹⁷ One way of thinking of the narrow content of a thought is that it is what remains of the broad content when we "take away the anchoring conditions" (*PS*: 51). But Fodor cautions against taking this subtraction picture too literally. A better way of thinking of narrow content, he suggests, is to view the narrow content of a thought as a function (in the mathematical sense—a mapping) from contexts to broad contents. Since broad contents determine truth conditions, narrow contents will determine mappings from contexts to truth conditions. "Two [narrow] thought contents are identical only if they effect the same mapping of thoughts and contexts onto truth conditions" (*PS*:48). Thus the thought tokens that lead both Fodor and Twin-Fodor to say "Water is wet" have the same narrow content, since they would have the same broad content if they were embedded in the same context.

[S]hort of a miracle the following counterfactual must be true: Given the neurological identity between us, in a world where I am in my Twin's context, my

17. "I learned 'anchors' at Stanford," Fodor tells us. "[I]t is a very useful term despite—or maybe because of—not being very well defined" (*PS*:49).

“water”-thoughts are about XYZ iff his are. (And, of course, vice versa: in a world in which my Twin is in my context . . . it must be that his water-thoughts are about H₂O iff mine are.) (PS:48)

3.2

One complaint about this notion of narrow content, the one that Fodor suspects “*really* bugs people” (PS:50), is that it seems impossible to say what the narrow content of a thought is. Fodor and Twin-Fodor have thought tokens with the same narrow content. But what is it that they both think? What is the narrow content of those thoughts? It can’t be *that water is wet*, since Twin-Fodor doesn’t think that. Nor can it be *that XYZ is wet*, since Fodor doesn’t think that. It seems that “narrow content is radically inexpressible” (PC:50). If this is right, however, it is hard to see how narrow content could serve the purpose for which it is intended. Recall that narrow content was supposed to provide a species of content-based taxonomy that would be useful in scientific psychology. If we insist, as Fodor does, that the states and properties invoked in scientific psychology must supervene on physiological states and properties, then psychological generalizations cannot invoke broad content. An alternative strategy is to couch those generalizations in terms of narrow content. But if narrow content is “radically inexpressible” it would appear that psychology’s generalizations could never be stated.¹⁸

Fodor’s response to this problem is to suggest that while we can’t *express* the narrow content of the thought that he and his Twin share, we can “sneak up on the shared content by *mentioning*” an appropriate English expression—in this case presumably the sentence: ‘Water is wet.’ But in offering this response I think Fodor seriously understates the case to be made for his notion of narrow content. We can do more than “sneak up” on the narrow content of a mental state; we can explicitly introduce a way of talking about it. The central idea is very simple. Expressions of the form: “—believes that p” are predicates whose extension in any possible world is the class of people who believe that p in that world. Given these predicates along with the notion of a doppelganger, we can introduce expressions of the form “—believes that [p]” (think of it as “bracketed” belief) whose extensions in any possible world include everyone in that world who believes that p, along with all of their doppelgangers. Similarly, expressions of the form “—has the (broad) content that p” are predicates whose extension in any possible world included the class of brain state tokens whose broad content is p. Here we can introduce expressions of the form “—has the (narrow) content that [p]” whose extension in any possible world includes the class of brain state tokens whose (broad) content is p, along with the

18. I owe this way of making the point to Warren Dow. See Dow (1991). A similar point is made by Baker (1987a).

physically identical tokens in all doppelgangers of people who harbor tokens whose broad content is *p*.¹⁹ These “bracketed” predicates are no less clear and no less systematic than the broad-content predicates on which they are based.

This strategy for talking about narrow content has what might at first seem to be a curious feature. In some cases the extension of “—has the (narrow) content that [*p*]” and the extension of “—has the (narrow) content that [*q*]” are going to be the same even though ‘*p*’ and ‘*q*’ are replaced by sentences that differ in reference and truth value. Consider, for example, a version of Putnam’s aluminum/molybdenum story. In the southern province of a certain English-speaking country, pots are typically made of aluminum, and this fact is known to a southerner, (Southern)Sam, who knows very little else about aluminum. In the northern province, pots are typically made of molybdenum. But in the north, molybdenum is called “aluminum.” (Northern)Sam, who is (Southern)Sam’s doppelganger, has a belief which he expresses with the words “Pots are typically made of aluminum.” Though of course given the standard intuitions in these cases, the belief token he is expressing has the (broad) content that pots are typically made of molybdenum. Now what about the narrow content of the belief (Northern)Sam expresses? Since that belief has the (broad) content that pots are typically made of molybdenum, it has the (narrow) content that [pots are typically made of molybdenum]. But since it is neurophysiologically identical to (Southern)Sam’s belief whose (broad) content is that pots are typically made of aluminum, it also has the (narrow) content that [pots are typically made of aluminum]. Similarly, the belief token that (Southern)Sam expresses when he says “Pots are typically made of aluminum” has both narrow (or bracketed) contents. There is nothing particularly surprising about any of this. The device we’ve introduced for attributing narrow contents exploits the expressions we would use in attributing broad contents and expands their extensions in a systematic way. It is to be expected that in some cases two of these enlarged extensions will coincide.²⁰

19. Something rather like this was suggested very briefly in Stich (1983:192, fn.). More recently, similar ideas have been developed by Valerie Hardcastle (1990) and Michael Devitt (1989). Perhaps I should add that I do not take my suggestion to be in competition with Fodor’s strategy for “sneaking up on” narrow content; mine is just a bit more explicit. Indeed, were I to develop my definition more carefully, and without riding roughshod over the fine distinction between use and mention, it would be obvious that my story, like Fodor’s, enables us to talk about narrow content by *mentioning* sentences.

20. This note is for aficionados only. I have argued that (Northern)Sam’s belief falls within the extension of both

- (i) “—has the (narrow) content that [pots are typically made of molybdenum],” and
- (ii) “—has the (narrow) content that [pots are typically made of aluminum].”

But it does not follow that (i) and (ii) are co-extensive. For consider the case of an expert in the North, someone who knows a great deal about how to distinguish aluminum from molybdenum and who also (broadly) believes that pots are typically made of molybdenum. Plainly, his belief is in the extension of (i). Is it also in the extension of (ii)? Not unless he has a doppelganger whose belief has the (broad) content that pots are typically made of aluminum. But if he has a doppelganger in the South, it is not at all clear that his doppelganger

The conclusion I would draw here is that the putative “radical inexpressibility” of narrow content is not a problem that Fodor need worry much about. It is easy enough to devise locutions for attributing narrow content to cognitive states, and these locutions can be used to state psychological generalizations in much the same way that locutions attributing broad content can.

3.3

As I see it, the real problem with narrow content does not derive from our inability to talk about it, and thus state generalizations in terms of it. Rather, the problem is that if the computational paradigm sketched in section 2 is on the right track, then many of the true generalizations—many of those that actually describe mental processes—are not going to be statable in terms of narrow content. The taxonomy of mental states imposed by narrow content is going to be both too coarse and too ill behaved to exploit in a serious scientific psychology. Perhaps the best way to see why a narrow content taxonomy is too coarse is to compare three taxonomic schemes: the one imposed by fat syntax, the one imposed by broad content, and the one imposed by narrow content.

Each mental state token is a brain state token; its “essential” type will be specified neurophysiologically. But a pair of brain state tokens in a pair of people may be very different neurophysiologically, and still count as tokens of the same fat syntactic type, provided that they have basically the same pattern of causal connections with stimuli, with behavior, and with other appropriate brain states. It’s also worth noting that if there is a pair of neurophysiologically identical states embedded in a pair of neurophysiologically identical organisms, and if one of these states is in a fat syntactic category, the other will always be in the same fat syntactic category. Fat syntax supervenes on physiology. Our commonsense intuitions about broad content provide another scheme for classifying brain state tokens. The lesson to be learned from Twin-Earth, and from Burge’s thought experiments, is that classification by broad content turns on physical, historical and linguistic *context*. Thus broad content does not supervene on physiology, and in this respect its taxonomic categories slice too finely; it sometimes puts an organism and its

would (broadly) believe that pots are typically made of aluminum. More likely, the relevant mental state of the expert’s Southern doppelganger would be so anomalous that it would have no broad content at all. For unless the story is told in a pretty strange way, you *can’t*

- (a) be the doppelganger of an expert on aluminum and molybdenum who broadly believes that pots are typically made of molybdenum,
- (b) live in a world in which pots are typically made of aluminum,
- (c) (broadly) believe that pots are typically made of aluminum.

To see the point, imagine that the Northern expert can distinguish the two metals by touch and sight, and ask what his Southern doppelganger would say when confronting the aluminum pots that are typical in his environment.

doppelgänger in different categories. Narrow content provides a third strategy for classifying brain state tokens, one which starts with broad content but ignores context. Thus, despite the terminological oddness, the categories of narrow content are larger than those imposed by a broad content taxonomy. Moreover, like fat syntax, narrow content supervenes on physiology.

All of this might lead one to suppose that the taxonomies imposed by narrow content and fat syntax *coincide*. That is, it might lead one to think that a pair of brain state tokens in a pair of individuals will be of the same fat syntactic type if and only if they have the same narrow content. However, this is all but certain to be a mistake. If we ignore the vagueness of the narrow content taxonomy, a theme to which I'll return shortly, then it may be the case that sameness of fat syntax guarantees sameness of narrow content ("plus or minus a bit," as Fodor might say).²¹ But on almost any plausible reading, the categories imposed by a narrow content taxonomy are much larger than those imposed by fat syntax. Thus sameness of narrow content does not guarantee sameness of fat syntax.

The literature is full of examples that illustrate this mismatch. Perhaps the most obvious examples involve people with unusual or defective perceptual systems. To take an extreme case, consider Helen Keller. If Ms Keller were to be told by a trusted informant that there is a fat cat in the room, she would come to believe that there is a fat cat in the room. That is, she would acquire a brain state which functions like a belief and which has the (broad) content that there is a fat cat in the room. Similarly, if I were told by a trusted informant that there is a fat cat in the room, I would acquire a brain state which functions like a belief and which has the (broad) content that there is a fat cat in the room. Thus both Ms Keller's brain state and mine would have the (narrow) content that [there is a fat cat in the room]. But surely those two states differ radically in their fat syntax. There are all sorts of perceptual stimuli (both visual and auditory) that would cause me, but not Ms Keller, to acquire the belief that [there is a fat cat in the room]. And states whose patterns of causal interaction with stimuli differ substantially do not share the same fat syntax. Much the same point could be made, though perhaps less dramatically, with examples of people with other perceptual anomalies, both real, like color blindness, and imagined.²²

21. Fodor (1980:240).

22. For detailed examples along these lines, see Stich (1983:66–8) and Stich (1982:185–8).

Kenneth Taylor has suggested that the objection I am urging against Fodor dissolves if we focus more steadfastly on Fodor's "official" account of narrow content which takes the narrow content of a thought to be *a function* from contexts to broad contents. On my account of narrow content, any two thoughts with the same broad content must have the same narrow content. But, Taylor urges, if we view the narrow content of a thought as a function from contexts to broad contents, then it is entirely possible that Ms Keller's thought and mine do not have the same narrow content. For there might be some contexts in which Ms Keller's thought and mine did not have the same broad content, and if this is possible, then on the function account of narrow content our thoughts do not have the same narrow content.

I am inclined to think that the function account of narrow content is more than a bit obscure. For I am not at all clear about what a *context* is; nor am I sure how we are supposed to play the game of imagining peo-

In the Helen Keller example, differences in fat syntax are due to differences in the way stimuli affect mental states. But there are also cases in which differences in syntactic type are engendered by differences in the way mental states interact with *each other*. Some people are logically acute; it is plausible to suppose that the mechanism underlying their reasoning makes many valid inferences and few invalid ones. Other people are significantly less acute; their mental mechanism makes many fewer valid inferences and many more invalid ones. On a syntactic taxonomy—even a *skinny syntactic taxonomy*—the states being manipulated by these mechanisms are of different syntactic types. But in many such cases the intuitive commonsense taxonomy of broad content classifies the states being manipulated as having the same (broad) content. And, of course, states with the same broad content have the same narrow content. In addition to these normal interpersonal differences in inferential capacities, there are also lots of pathological cases, some real and some imagined, in which people reason in ways very different from the way I reason, but where commonsense psychology is still comfortable in attributing the same broad content.²³ Here too, syntax and narrow content will diverge.

What I have been arguing is that there are major differences between a taxonomy based on narrow content and one based on fat syntax (or skinny syntax, for that matter). In many cases the syntactic taxonomy will be substantially more fine grained, and will draw substantially more distinctions, than the narrow content taxonomy. There are lots of examples in which a pair of belief state tokens will differ in their fat syntax though not in their narrow content. The reason this is important is that, along with Fodor, I have been assuming that the cognitive mind is a “syntactic engine” and that the mechanism controlling cognitive processes is “sensitive solely to syntactic properties.” But if this is right, then the generalizations that describe cognitive processes will be storable in syntactic terms, and these will typically be more fine grained than generalizations storable in terms of narrow content. The generalizations of a computational theory will describe different patterns of causal interaction for cognitive states with different fat syntax, even though in many cases those states will have the same narrow content. So if the

ple and their thoughts embedded in other contexts. Consider the example in the text. Is Ms Keller’s context different from mine? If so, what would it be for me to be in her context? Would I have to have her handicaps? Would I have to have had the same biography? The mind boggles.

But even if we suppose these questions can be answered in some coherent and principled way, I doubt the answers will do Fodor much good. To avoid the objection I am urging, it will have to be the case that the taxonomy generated by the function account of narrow content coincides with the taxonomy generated by fat syntax. And I see no reason to think this will be the case. Certainly, Fodor has offered no argument for this claim. Moreover, if as Taylor suggests, Ms Keller’s belief and mine have different narrow contents on the function account, it is hard to see why the same will not be true of the beliefs of other people who broadly believe that there is a fat cat in the room, but who differ from me less radically than Ms Keller does. However, if this is the case, then the function account of narrow content runs the risk of individuating much too finely. Only doppelgangers will have thoughts with the same narrow content.

23. For some examples, see Stich (1983:68–72). For examples of a rather different sort, see Cherniak (1986). For another example, see Dennett (1981b:54–5).

computational paradigm is the right one, then many of the generalizations that describe the mind's workings are simply not going to be statable in terms of narrow content.

Throughout this section I have been writing as though the broad content taxonomy provided by commonsense psychology is reasonably clear and stable, and thus that predicates of the form “—has the (broad) content that p” have a reasonably well defined extension. However, there is good reason to doubt that this is so. Following Quine's lead, a number of writers have assembled cases which seem to show that commonsense intuitions about the extensions of such predicates are highly context sensitive. Whether or not a state can be comfortably classified as having the content that p depends, to a significant degree, on the context in which the question arises.²⁴ I have developed an account of the tacit principle underlying commonsense content attribution which views them as a sort of similarity judgment. This account explains their context sensitivity, and various other phenomena as well. But whether or not my explanation of the phenomena is correct, I am inclined to think that the data speak for themselves. By varying the context in which the question is asked, we can get competent users of commonsense psychology to judge that a particular cognitive state token clearly has the content that p, or that it clearly does not. If this is right, it provides yet another reason for thinking that the generalizations of a serious scientific psychology will not be statable in the taxonomic categories provided by narrow content. For the categories of a narrow content taxonomy are simply the categories of a broad content taxonomy extended to meet the demands of the principle of autonomy. But the broad content taxonomy of commonsense psychology is too vague, too context-sensitive and too unstable to use in a serious scientific theory. *Narrow* content inherits all of these deficits.

4 Keeping Score

Toward the end of section 1, I quoted the three conditions that, on Fodor's view, would have to be met by the states a psychological theory postulates, if that theory is to count as “endorsing” the propositional attitudes, and thus avoiding the “catastrophe” that would ensue “if commonsense intentional psychology really were to collapse.” It's time to ask which of those conditions are likely to be met. Along with Fodor, I'll assume, as I have been all along, that the computational paradigm is correct, and that the mind is “a syntax driven machine” whose operations are “sensitive solely to syntactic properties.”

The third condition on Fodor's list is that “the implicit generalizations of commonsense belief/desire psychology” must be “largely true” of the states postulated by the psychological theory in question. Presumably, Fodor's hope went something like this:

24. See, for example, Stich (1982:180–203), where I describe the phenomenon as the “pragmatic sensitivity” of belief attributions. See also Stich (1983:90–110). Much the same moral can be drawn from Dennett's examples of the use of intentional notions to describe trees and his example of the young child who asserts that Daddy is a doctor. For the first, see Dennett (1981a:22); for the second, see Dennett (1969:183).

The generalizations of commonsense psychology are couched in terms of (broad) content. But the Twin Earth examples show that “you can’t make respectable science out of the attitudes as commonsensically individuated” (*PS*:30). Very well, then, we’ll move to narrow content, since, unlike broad content, narrow content supervenes on physiology. Given any commonsense generalization about tokens of the belief that *p*, there will be a parallel narrow generalization—a generalization about the tokens of the belief that [*p*]. And that latter generalization will be scientifically respectable.

To satisfy Fodor’s third condition, however, it is not sufficient that the narrow analogues of broad content generalizations be scientifically respectable. Most of them must also be true. Now if the mind really is a syntax driven machine, and if syntactic categories can be matched up, near enough, with the categories of narrow content, then it looks like we’re home free. But the burden of my argument in 3.3 was that syntactic and narrow content taxonomies will not match up, because the latter is both too coarse and too ill-behaved. If the computational paradigm is correct, I argued, then many of the generalizations that describe the mind’s workings are not going to be storable in terms of narrow content. If that’s right, then Fodor’s third condition will not be satisfied.

Let’s turn, now, to Fodor’s first condition; that the states a psychological theory postulates must be “semantically evaluable.” How well do we fare on this one if the computational paradigm is correct? I am inclined to think that here again the ill behaved context sensitivity of semantic taxonomies poses real problems. If it is indeed the case that by varying the context of the question we can get competent users of commonsense psychology to judge that a particular cognitive state token clearly has the content that *p*, or that it clearly does not, then it’s hard to see how even the tokens, let alone the types posited by a serious, computational, scientific psychology will be “semantically evaluable.”

One final point. Suppose I am wrong about the mismatch between syntactic and narrow content taxonomies; suppose that the generalizations of a scientifically solid psychology really can be stated in terms of narrow content. Would it then follow that the states postulated by such a theory are “semantically evaluable”? Fodor himself seems ambivalent. Consider the following:

[I]f you mean by content what can be semantically evaluated, then what my water-thoughts share with Twin “water”-thoughts *isn’t* content . . . We can’t say . . . what Twin thoughts have in common. This is because what can be said is ipso facto semantically evaluable; and what Twin-thoughts have in common is ipso facto not. (*PS*:50; the emphasis is Fodor’s.)

But, of course, what Fodor’s water-thoughts share with Twin ‘water’-thoughts, what “Twin-thoughts have in common,” *is* narrow content. So in this passage Fodor seems to

admit—indeed insist—that narrow content is *not* semantically evaluable. Elsewhere he is even more explicit:

You can have narrow content without functional-role semantics because *narrow contents aren't semantically evaluable*; only wide contents have conditions of satisfaction. (*PS*:83; the emphasis is Fodor's.)

Still, perhaps this is just a debater's point. For in several other passages Fodor notes that narrow content "is semantically evaluable relative to a context" (*PS*:51). And perhaps this is all that is required to satisfy his first condition. There's no need to decide the point since, as I see it, the real problem with narrow content is not that it fails to be "semantically evaluable" (whatever that might come to) but that it fails to match up with the syntactic taxonomy of a computational psychology.

The remaining item on Fodor's list of conditions is that the states posited by a psychological theory must "have causal powers." On this one Fodor wins easily. If the computational paradigm is on the right track, then the syntactically taxonomized states posited by a correct computational theory are sure to have causal powers.

By my count, the score against Fodor—and against intentional psychology—is two to one.²⁵

REFERENCES

- Armstrong, D. (1978). *A Theory of Universals*, Cambridge: Cambridge University Press.
- Baker, L. (1987a). *Saving Belief: A Critique of Physicalism*, Princeton: Princeton University Press.
- Baker, L. (1987b). "Content by courtesy," *Journal of Philosophy*, 84, 197–213.
- Block, N. (1986). "Advertisement for a semantics for psychology," in P. French, T. Uehling & H. Wettstein (eds.), *Midwest Studies in Philosophy*, Vol. 10, *Studies in the Philosophy of Mind*, Minneapolis: University of Minnesota Press.
- Burge, T. (1979). "Individualism and the mental," *Midwest Studies in Philosophy*, 4, 73–121.
- Burge, T. (1986). "Individualism and psychology," *Philosophical Review*, 95, 1, 3–46.
- Cherniak, C. (1986). *Minimal Rationality*, Cambridge, MA: MIT Press.
- Dennett, D. (1969). *Content and Consciousness*, London: Routledge and Kegan Paul.
- Dennett, D. (1981a). "True believers: the intentional stance and why it works," in A. Heath (ed.), *Scientific Explanation*, Oxford: Oxford University Press.
- Dennett, D. (1981b). "Three kinds of intentional psychology," in R. Healy (ed.), *Reduction, Time and Reality*, Cambridge: Cambridge University Press.
- Dennett, D. (1982). "Beyond belief," in A. Woodfield (ed), *Thought and Object*, Oxford: Oxford University Press.

25. I am indebted to many people for much useful conversation on these matters. Those I can recall are Daniel Dennett, Michael Devitt, Warren Dow, Jerry Fodor, Gary Hardcastle, Patricia Kitcher, Kenneth Taylor and Valerie Hardcastle. I hope the others will accept my thanks anonymously. Special thanks are due to Warren Dow for his helpful comments on an earlier version of this paper.

- Devitt, M. (1989). "A narrow representational theory of the mind," in S. Silvers (ed.), *Representations: Readings in the Philosophy of Mental Representation*, Dordrecht: Kluwer.
- Dow, W. (1991). *Content and Psychology*, Ph.D. thesis, Department of Philosophy, University of California–San Diego.
- Fodor, J. (1975). *The Language of Thought*, Cambridge, MA: Harvard University Press.
- Fodor, J. (1978). "Tom Swift and his procedural grandmother," *Cognition*, 6, 229–247.
- Fodor, J. (1980). "Methodological solipsism considered as a research strategy in cognitive science," *Behavioral and Brain Sciences*, 3, 1, 63–109.
- Fodor, J. (1981). *RePresentations: Philosophical Essays on the Foundations of Cognitive Science*, Cambridge, MA: MIT Press.
- Fodor, J. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, MA: MIT Press.
- Gordon, R. (1986). "Folk psychology as simulation," *Mind and Language*, 1, 158–171.
- Hardcastle, V. (1990). "In defense of a different taxonomy: a reply to Owens," *Philosophical Review*, 99(3), 425, 431.
- Kim, J. (1969). "Events and their descriptions: Some considerations," in N. Rescher et al. (eds.), *Essays in Honor of C. G. Hempel*, Dordrecht: Reidel.
- Kim, J. (1976). "Events as property exemplifications," in M. Brand & D. Walton (eds.), *Action Theory*, Dordrecht: Reidel.
- Kitcher, P. (1985). "Narrow taxonomy and wide functionalism," *Philosophy of Science*, 52, 1, 78–97.
- Levin, J. (1988). "Must reasons be rational?" *Philosophy of Science*, 55.
- Loar, B. (1987). "Social content and psychological content," in R. Grimm & D. Merrill (eds.), *Contents of Thought: Proceedings of the 1985 Oberlin Colloquium in Philosophy*, Tucson: University of Arizona Press.
- Lycan, W. (1988). *Judgment and Justification*, Cambridge: Cambridge University Press.
- Owens, J. (1987). "In defense of a different doppelganger," *Philosophical Review*, 96, 521–554.
- Putnam, H. (1975a). "The meaning of 'meaning'" in Putnam (1975b).
- Putnam, H. (1975b). *Mind, Language and Reality*, Cambridge: Cambridge University Press.
- Putnam, H. (1978). *Meaning and the Moral Sciences*, London: Routledge & Kegan Paul.
- Putnam, H. (1983). "Computational psychology and interpretation theory," in *Realism and Reason*, Philosophical Papers (Vol. 3), Cambridge: Cambridge University Press.
- Quine, W. V. (1960). *Word and Object*, Cambridge, MA: MIT Press.
- Ramsey, W., Stich, S. and Garon, J. (1990). "Connectionism, eliminativism and the future of folk psychology," *Philosophical Perspectives*, 4: *Action Theory and Philosophy of Mind*, 499–533.
- Schiffer, S. (1981). "Truth and the theory of content," in H. Parrett & P. Bouverese (eds.), *Meaning and Understanding*, Berlin & New York: de Gruyter.
- Stack, M. (unpublished). "Why I don't believe in beliefs and you shouldn't either," paper delivered at the annual meeting of the Society for Philosophy and Psychology, 1980.
- Stich, S. (1982). "On the ascription of content," in A. Woodfield (ed.), *Thought and Object*, Oxford: Oxford University Press.
- Stich, S. (1983). *From Folk Psychology to Cognitive Science: The Case Against Belief*, Cambridge, MA: MIT Press/Bradford Books.