

1993

Belief States and Narrow Content

Curtis Brown

Trinity University, cbrown@trinity.edu

Follow this and additional works at: http://digitalcommons.trinity.edu/phil_faculty

 Part of the [Philosophy Commons](#)

Repository Citation

Curtis Brown, "Belief States and Narrow Content," *Mind and Language* 8 (1993): 343-67.

This Pre-Print is brought to you for free and open access by the Philosophy Department at Digital Commons @ Trinity. It has been accepted for inclusion in Philosophy Faculty Research by an authorized administrator of Digital Commons @ Trinity. For more information, please contact jcostanz@trinity.edu.

(*Mind and Language* 8 (1993): 343-367.
There may be small differences between
this version and the published one.)

Belief States and Narrow Content

Curtis Brown

1. Introduction: Four Incompatible Theses About Belief

Each of the following four theses is a plausible view about belief and other propositional attitudes, but the four theses are jointly incompatible, and so cannot all be true.

The first thesis is that beliefs play a role in explaining behavior. This is reasonably uncontroversial, though it has been controverted. Why did I raise my arm? Because I wanted to emphasize a point, and believed that I could do so by raising my arm. The belief that I could emphasize a point by raising my arm is central to the most natural explanation of my action.

The second thesis is that those features of mental states which play a genuine role in explaining behavior must be intrinsic features of the individual: they must supervene on the current, internal physical states of the agent, which is to say they must *not* depend on the agent's environment, situation, or context (except insofar as the environment causally affects the agent's intrinsic states). The rough idea here is that nothing outside my skin can cause my behavior except by causing changes inside my skin. Of course, not all features of our mental states play a role in explaining our behavior. Consider knowledge, a mental state that sometimes helps explain what one does. On a particular occasion, I may go to the store because I want food and know that there will be food there. Now, part of what makes my mental state an instance of knowledge is the truth of the proposition that there is food at the store. That this proposition is true is certainly not a fact about matters inside my skin. But, although my knowledge explains my behavior, and although the truth of the proposition known is a feature of my knowledge that is not an intrinsic feature of me, it does not follow that non-intrinsic features of my mental states play a role in explaining my behavior, since the truth of the proposition known does *not* play a role in explaining my behavior. If the second thesis is correct, then any feature of my mental states that, like the truth of known propositions,

does not supervene on intrinsic facts about me will not be a feature that explains what I do.

The third thesis is that *content* is a feature of beliefs that plays a role in explaining behavior. Beliefs explain actions in part by making them reasonable; it is hard to see how anything but the contents of the beliefs could do that. Again, my beliefs at a particular time help explain my later beliefs, through something like induction or inference to the best explanation, and one of the notable failures of positivism was its attempt to "syntactify" such inference. If induction plays a role in belief change, and induction depends on the semantic properties of the beliefs involved, then content is central to the explanation of changes in belief.

The fourth and final thesis is that the content of most beliefs is not intrinsic. This is perhaps the only thesis of the four that seems to run counter to common sense. But although the fourth thesis is not a bit of common sense, it has become widely accepted among philosophers, as a result of powerful and persuasive examples introduced by Hilary Putnam (Putnam 1975b, 1983) and Tyler Burge (Burge 1979).

Clearly, at least one of the four theses must be false. Each has its opponents: eliminative materialists such as Paul Churchland (1985) and Patricia Churchland (1986, esp. Chapters 7 and 9) reject the idea that beliefs play a role in explaining behavior; Tyler Burge (1979, 1982a, 1982b, 1986, 1990) and Robert Stalnaker (1989), among others (e.g. Baker 1987; Owens 1987, 1990; Jackson and Pettit, 1988), reject the idea that explanatory states must be intrinsic; Stephen Stich (1983) rejects the idea that the explanatory role of beliefs depends on their content, arguing that for purposes of psychological explanation beliefs, if they play a role at all, should be individuated syntactically rather than semantically; and Jerry Fodor (1987, 1991) and others (White 1982, 1991; McGinn 1982; Loar 1987, 1988; Devitt 1990) have argued that, although it is true that the ordinary or "broad" content of mental states depends on the situation of the agent, there is also another sort of content, "narrow content," which depends only on factors internal to the agent. It is this last move that concerns me here.

I will begin by suggesting two criteria an account of narrow content should meet. I then discuss a device offered in three recent publications by Valerie Walker (1990) and Stephen Stich (Stich 1991; Jones, Mullaie, and Stich 1991) as a means of expressing the narrow content of mental states or of actions. Walker offers the device as a response to Joseph Owens' critique of the notion of narrow content (Owens 1987), while Stich offers it as an explication of Fodor's characterization of narrow content. I argue that the device is fatally

flawed. It will not provide a useful psychological typology, since it will count people with radically different beliefs as sharing the same narrow contents. And it does not capture the notion Fodor had in mind, since it does not determine a function from contexts to broad contents. I then suggest that a related device recognizably in the same spirit as Walker's and Stich's provides a better explication of Fodor's notion. Pleasantly enough, this device is one I advanced in a previous publication (Brown 1986), though I did not use the term "narrow content" and indeed did not claim that the device captured a notion of content at all. Finally, with this device in hand, I argue that current conceptions of narrow content should be revised in three respects: Fodor's notion is better thought of as describing belief states rather than narrow contents, and as leaving it an open question whether belief states are or are not characterizable in terms of content; particular beliefs very likely do not have narrow contents taken individually; and narrow content should not be regarded as a kind of content shared by mental states and English sentences.

2. Criteria for a Successful Account of Narrow Content

The point of a theory of narrow content is to provide a characterization of mental states that captures precisely those properties that play a role in explaining action. Our discussion so far already points to one reason why ordinary or broad content risks missing important explanatory generalizations, namely that two people with exactly the same intrinsic properties might have different broad contents, provided that their contexts were sufficiently different. Whatever the explanatorily relevant features of agents are, it seems that they should be the same for agents who are intrinsically exactly similar, since the environment can contribute to the explanation of behavior only by way of its effects on the intrinsic features of the agent.

One criterion of success for an account of narrow content, then, is that it should group together token mental states that have the same effects on behavior. Since two people who are intrinsically exactly alike will, if placed in the same circumstances, behave alike, it follows that an account of narrow content should, in particular, group together the mental states of people who are intrinsically exactly alike -- of doubles or doppelgängers, for short. The examples of Putnam and Burge show persuasively that ordinary attitude ascriptions do not satisfy this criterion. Narrow content, if there is such a thing, ought to say what my Twin Earth counterpart and I share by virtue of which I believe that lakes contain water and he believes that lakes contain XYZ. Although the *broad* contents of my Twin's and my mental states are different, it seems that they ought to have the same *narrow* content.

A second criterion of success for a theory of narrow content, however, is less often remarked on. This is that such a theory should discriminate between mental states that have different effects on behavior. If two token mental states would, under the same circumstances, lead subjects to act differently, then they should be regarded as tokens of different types -- as having different narrow contents. Just as Putnam's and Burge's examples show that ordinary attitude ascriptions do not satisfy the first condition, so Kripke's famous puzzle about belief (Kripke 1979) shows that they do not satisfy the second.

In Kripke's familiar story Pierre, a native Frenchman initially innocent of English, forms the belief that London is pretty, a belief he expresses by the French 'Londres est jolie.' He then moves to London, learns English, and comes to believe that London is not pretty, a belief he expresses by the English sentence 'London is not pretty', without realizing that the city he calls 'Londres' is the same city as the one he calls 'London.' It seems natural to say that Pierre is related to the same proposition, that London is pretty, by way of two different mental states; that the mental states are different is shown by the fact that he believes the proposition when it is taken in one way and disbelieves it when it is taken in another. Now change the example slightly. Suppose that on moving to London Pierre came to believe, not that London is *not* pretty but, again, that London *is* pretty. Then he believes the same proposition twice over. His two beliefs that London is pretty have very different causal consequences; for example, the belief he acquired in England helps explain why he does not move away from London when he gets job offers elsewhere, while the belief he acquired in France does not. Since the two beliefs, despite having the same broad content, have different effects on behavior, an account of narrow content ought to discriminate between them.

The two sorts of examples just mentioned make related but distinct points. Putnam's and Burge's examples show that such broad contents as the proposition that lakes contain water do not characterize one's intrinsic mental state *essentially*, since my Twin Earth counterpart and I are in the same intrinsic state, but one of us believes this proposition and one of us does not. Our modified version of Kripke's example shows that such broad contents as the proposition that London is pretty do not characterize one's intrinsic mental state *completely*, since there are different ways to believe this proposition. In Twin Earth examples we have the same mental state, but different propositions believed; in examples like that of Pierre, we have the same proposition believed, but different mental states. Unfortunately, the second sort of example has, in *this* context, been relatively neglected.

3. Stich's Device for Expressing Narrow Content

Fodor (1987) develops his notion of narrow content primarily as a response to Twin-Earth examples. What do I share with my Twin when I believe that lakes contain water and he believes that lakes contain XYZ? Fodor describes narrow contents as "functions from contexts and thoughts onto truth conditions" (47), where truth conditions themselves are understood to be functions from worlds to truth values. The idea is that the narrow content of my thought when I believe that lakes contain water is a function that takes my actual context into the proposition that lakes contain water, and takes my Twin's context into the proposition that lakes contain XYZ. "The 'broad content' of a thought . . . [is] what you get when you specify a narrow content *and fix a context*" (48). Fodor suggests that "you can't, in absolute strictness, express narrow content," although "there are ways of sneaking up on it" (51).

Stich responds that "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." I would now like to consider the apparatus Stich offers for doing so. He presents it as follows:

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 doppelgänger, 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 doppelgängers. Similarly, expressions of the form '___ has the (broad) content that p' are predicates whose extension in any possible world includes 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 physically identical tokens in all doppelgängers of people who harbor tokens whose broad content is p (Stich 1991; Jones, Mullaire, and Stich 1991, pp. 65-66).

It is unclear from this explanation whether cross-world doppelgängers are allowed: that is, it is unclear whether the extension of '___ believes that [p]' at a world *w* includes only individuals in *w* who are doppelgängers of people who believe that p *in w*, or whether it includes those in *w* who are doppelgängers of people who believe that p *in other worlds as well*. Either way, the proposal has unacceptable results.

4. The Intra-World Interpretation

It may be useful to set out the two possible interpretations of Stich's device in a quasi-formal way. For x to be a doppelgänger of y is for x and y to have what we might call the same *total intrinsic state*, that is, to have all the same intrinsic properties (excluding haecceities), though their extrinsic or relational properties may differ. Clearly, the notion of a doppelgänger needs to be relativized to time: if Twin Earth diverges from Earth just after t , then we will want to say that my Twin is my doppelgänger *at* t , but not *after* t . However, the three-place relation, x is a doppelgänger of y at t , still does not adequately capture the notion. Being a doppelgänger is a *cross-temporal relation*. If Twin Earth is exactly like earth, but fifty years behind, then we will want to say that my Twin as he is at t is a doppelgänger of me as I am at t minus 50 years. If we symbolize the total intrinsic state of a thing x at a time t as $TS(x, t)$, then we can say that x at t is a doppelgänger of y at t' just in case $TS(x, t) = TS(y, t')$. This apparatus enables us to give a formulation of the first possible reading of Stich's explanation of his bracketing device, a reading we may call the *intra-world interpretation*. On the intra-world reading, x believes that [p] at t just in case *either* x believes that p at t , *or* $(\exists y)(\exists t')(y$ believes that p at t' and $TS(y, t') = TS(x, t)$). (Although the first disjunct helps to see how the proposal works, it could be eliminated without loss of content, since it merely captures the degenerate case of the second disjunct in which $x = y$ and $t = t'$.) On this reading, to find the people who actually believe that [p], we take the people who actually believe that p and add all their actual doppelgängers; to find the people who believe that [p] in some other possible world w , we take the people who believe that p in w and add those in w who are doppelgängers of people who believe that p in w .

On this reading, Stich's proposal clearly will not do. There are three main problems with Stich's account on the intra-world interpretation, two of which will carry over to the second interpretation as well. First, the account is too narrow, in the sense that it fails to group together mental states that have the same effects on behavior. It thus fails to satisfy our first criterion for the success of an account of narrow content.

Two examples will illustrate the point. First, consider the case of *the waterless world*. Pretend that our world, w_1 , contains an actual Twin Earth. Then in our world, those who believe that [lakes contain water] include not only those (like me) who believe that lakes contain water, but also my Twin, who believes instead that lakes contain XYZ. My Twin is included by virtue of being a doppelgänger of someone, namely me, who believes that lakes contain water. So far, all seems well. Now consider another world, w_2 , which contains Twin Earth just as it is in w_1 , but does not contain Earth at all, and in which there is no water, only XYZ. Supposing that one cannot have beliefs about a substance without either being acquainted with it or knowing its essential nature, and that

you cannot be acquainted with what does not exist, and stipulating that no one in w_2 speculates about the possible but unrealized substance H₂O, we get the result that no one in w_2 has beliefs about water, so in particular no one believes that lakes contain water. Now, to determine the extension of ‘believes that [p]’ in w_2 , we take those in w_2 who believe that lakes contain water, and then add their doppelgängers. So no one in w_2 believes that [lakes contain water], since no one in w_2 believes that lakes contain water! This cannot be the right result, since w_2 contains an exact replica of my Twin, and hence an exact replica of me, and the motivation behind the notion of narrow content was to provide a sort of content all my replicas would share with me, no matter how differently they were situated.

Nor can we simply say that it doesn’t matter what narrow contents we ascribe to otherworldly beings, so long as we get the right results for those in the actual world. First, if narrow content is to be of any psychological use, then statements about narrow content should support counterfactuals: for example, that my beliefs would have had the same narrow content they do have no matter how different my environment was. On the proposal we are considering, however, given the usual semantics for counterfactuals, it will turn out instead that, if the world had never contained water, my mental state would not have had the narrow content [lakes contain water]. And second, the proposal will also fail to get the actual narrow contents of actual people right. Consider a second example, the case of *chocolate versus vanilla*. Suppose Twin Earth is *slightly* different from Earth: my near-Twin there has a narrow mental state *almost* exactly like mine, but unlike me he believes that chocolate milkshakes are better than vanilla. Since this difference has no effect on the behavioral consequences of his beliefs about XYZ, he should still count as sharing my belief that [lakes contain water]. Since he is no longer a doppelgänger of someone who believes that lakes contain water, however, the proposal under consideration will not ascribe to him the narrow belief that [lakes contain water].

We have seen that Stich’s proposal, on the intra-world interpretation, gives too narrow an account of narrow content: it fails to group together mental states that should count as of the same type, and thus violates the first condition for a successful account. A second problem with the intra-world version of the proposal is that it is too broad. It straightforwardly violates the second condition for a successful account of narrow content, since it counts anyone who believes that p as being in the same narrow mental state even though, as Kripke’s example shows, there may be many different ways of believing p with quite different consequences for one’s behavior.

Perhaps the second condition seems too restrictive. Beliefs may share much of interest even if their consequences for behavior are not *quite* the same. Why not allow the notion of narrow content to have the feature that beliefs with somewhat different causal powers nevertheless share a narrow content? This objection misunderstands the problem. My insistence that beliefs with different behavioral effects have different narrow contents in no way prevents me from recognizing that such beliefs may also share interesting content-related features. Compare: The broad content of the belief that Beethoven was a composer is different from the broad content of the belief that Mozart was a composer. Recognizing this in no way prevents us from recognizing that the beliefs have an interesting content-related feature in common, namely that both are beliefs that someone is a composer. Similarly, insisting that Pierre's two beliefs that London is pretty have different narrow contents does not prevent us from seeing similarities between the two. Conversely, recognizing the similarities should not prevent us from seeing that there are important differences between the two, differences that seem to be differences of content, not merely differences of syntax, or neuronal firings, or something of the sort. If there are differences of content between the two beliefs that London is pretty, one would hope that an account of narrow content would enable us to say what they are (or at the very least acknowledge that they are present).

A third problem with the intra-world version of the account is that it does not determine a function from context to broad content. Since Fodor's notion is defined precisely as such a function, Stich's proposal cannot be what Fodor had in mind. Narrow content is supposed to be something which, together with my context, yields the belief that lakes contain water, and which, together with my Twin's context, yields the belief that lakes contain XYZ. But now consider a slightly modified version of the Twin Earth story. Earth is more or less as we know it, but in a remote and isolated laboratory, a scientist has synthesized a small quantity of XYZ. He shows his laboratory assistant samples of both water and XYZ, and explains the chemical structure of each. Together, they do the sophisticated experiments that alone can reveal the differences between the apparently similar substances. Then the scientist tells his assistant that water exists only in tiny quantities on Earth, while the abundant substance which falls from the skies, flows from the taps, and fills the lakes is XYZ!

The gullible lab technician believes broadly that lakes contain XYZ. He therefore also believes that [lakes contain XYZ]. I believe broadly that lakes contain water but, since my Twin Earth counterpart believes that lakes contain XYZ, I share with the lab technician the narrow belief that [lakes contain XYZ]. My belief that [lakes contain XYZ], together with my present context, yields the broad belief that lakes contain water. Now consider what the lab technician's belief that [lakes contain XYZ], together with *my* present context,

yields. (Since the context of a person at a time includes the person's history, superimposing the lab technician's mental state on my context will result in some discontinuity. This, although strange, does not affect the argument.) The lab technician's belief, even in my context, will surely yield the broad belief that lakes contain XYZ rather than the belief that lakes contain water! After all, unlike me, the lab tech knows the difference between the two, and thinks that, of the two, it is XYZ that fills the lakes. So the belief that [lakes contain XYZ] takes us from the same context to two different beliefs; therefore it is not a function.

This problem is of course closely related to the problem of broadness. It is precisely because the intra-world version of Stich's proposal is too broad that his "narrow contents" fail to determine a function from contexts to broad contents: the same narrow content groups together very different mental states, states so different that they yield different broad beliefs when plugged into the same contexts.

5. The Cross-World Interpretation

We have seen that the intra-world version of Stich's proposal faces serious difficulties. Let us consider a second, alternative version of the proposal and see how it fares. The second, *cross-world* version construes being a doppelgänger not just as a cross-temporal relation but also as a cross-world relation, so that a person, at a time, in a world, can be a doppelgänger of another person, at another time, in a different possible world. Now we can say that, for people x and y , times t and t' , and worlds w and w' , x at t in w is a doppelgänger of y at t' in w' just in case $TS(x,t,w) = TS(y,t',w')$. This gives us the necessary resources to offer a second interpretation of Stich's device. On this second interpretation, x believes that [p] at t in w just in case *either* x believes that p at t in w , *or* $(\exists y)(\exists t')(\exists w')(y$ believes that p at t' in w' and $TS(y,t',w') = TS(x,t,w)$). (As with the first interpretation, we could eliminate the first disjunct.) We may begin to feel some qualms about the proposal on this interpretation when we realize that it amounts to saying that I believe [p] if it is possible that a doppelgänger of mine should believe p; I will soon argue that such qualms are amply justified.

Indirect evidence that the cross-world interpretation must be what Stich intends comes from Valerie Walker's development of virtually the same idea in a piece cited by Stich (Walker 1990). The only substantial difference is that Walker develops the device as a mechanism for narrowly describing behavior rather than mental states. Walker unambiguously intends to include cross-world doppelgängers. She begins with "any common-sense intentional description of behavior D" (429). Such a description, she says, "has as its extension a class of

actual and possible actions, viz., those actions that comprise all the actual and possible referents of D" (430). (This is an unusual use of the term 'extension'; the class in question is what C. I. Lewis called the 'comprehension' of a term (Lewis 1946, p. 39), and is closely related to what a number of more recent writers mean by the 'intension' of a term, namely, a function from worlds to extensions.) Next, we add "Doppel(D)," "the set of all actual and possible instances of behavior that are physically identical to some member of D but are performed by a Doppelgänger of the agent of that action." Finally, we add the class determined by D to Doppel(D) to get [D], the set of actual and possible behaviors with the same narrow description.

How does the cross-world version of Stich's proposal fare with respect to the problems I have noted with the intra-world version? It avoids the difficulties of narrowness. First, the extension of 'believes that [p]' in the waterless world does include my Twin, since he is a doppelgänger of me-in-the-actual-world. Second, my actual near-Twin who, unlike me, prefers chocolate to vanilla, will now count as sharing my belief that [lakes contain water], since he presumably is a doppelgänger of some merely possible person who believes (broadly) that lakes contain water. (It is worth noting that if cross-world doppelgängers were not allowed, then the people who actually share a narrow content would be all and only the people who share the broad content in terms of which it is defined, since there are no actual doppelgängers.)

How does the cross-world version of the proposal fare with respect to the problem of broadness, the problem of grouping together very different mental states as having the same narrow content? This problem is actually exacerbated by the cross-world interpretation. The purported narrow descriptions of behavior '[D]' and of belief 'believes that [p]' are going to be far too inclusive to be of any serious psychological interest.

We may begin to see how inclusive the narrow descriptions will be by considering a point Stich himself makes (in Stich 1991). Stich introduces a different example to make the point, but I will simply adapt the Twin Earth example we have been considering. My Twin believes that lakes contain XYZ by virtue of being in the very same mental state that makes it the case that I believe that lakes contain water. So my Twin's mental state has the narrow content that [lakes contain XYZ], since on Stich's account having the broad content that p guarantees having the narrow content that [p], and also has the narrow content that [lakes contain water], since my doppelgängers' mental states have the same narrow content as mine. By the same token, of course, I also believe both that [lakes contain water] and that [lakes contain XYZ].

Does it follow that believing that [lakes contain water] and believing that [lakes contain XYZ] are just the same thing, i.e., that 'believes that [lakes contain water]' and 'believes that [lakes contain XYZ]' are cointensional? Stich suggests that it does not. Suppose that, known only to a few earthly experts, XYZ exists in small quantities on Earth, and that, known only to a few Twin-Earthly experts, water exists in small quantities on Twin Earth. Now consider an Earthly expert who can distinguish the two, who has studied the substance in lakes, and who believes that lakes contain water. The expert thereby believes that [lakes contain water]. But does he believe that [lakes contain XYZ]? Since he does not broadly believe that lakes contain XYZ, he will not believe that [lakes contain XYZ] unless he has a doppelgänger whose belief has the (broad) content that lakes contain XYZ. But a scientist on Twin Earth who knows how to distinguish water from XYZ and who has studied Twin-Earthly lakes will either come up with different results than the Earthly expert, namely results which show that the lakes contain XYZ rather than water, in which case he is not a doppelgänger of the earthly expert, or, if he comes up with the same results as the earthly expert, will conclude that the lakes contain water rather than XYZ! Either way the earthly expert does not have a doppelgänger who believes broadly that lakes contain XYZ.

Thus Stich. But wait! Once we include cross-world doppelgängers, the situation changes. For unless the characteristics the expert uses to distinguish water from XYZ are the very characteristics that make the two substances water and XYZ, there will be a possible world in which the characteristics the expert associates with water and uses to distinguish it from XYZ are possessed instead by XYZ (and vice versa), and in which a doppelgänger of our Earthly expert believes that lakes contain XYZ (which he of course calls 'water'). Since someone who believes that lakes contain XYZ, albeit in a different possible world, is a doppelgänger of the earthly expert, the earthly expert believes after all that [lakes contain XYZ].

It still does not follow that believing that [lakes contain water] is the same thing as believing that [lakes contain XYZ]. However, they will be different only if someone could distinguish water from XYZ directly by means of their essential properties rather than by properties that are only contingently related to their essences and which therefore are "switchable" in other possible worlds. I would argue that we always make only indirect use of the chemical structure of water in identifying it; chemical structure is usable by us only insofar as it produces observational experiences, and it seems that we could always imagine a sufficiently weird world in which the very observational experiences produced in the actual world by H₂O are instead produced by XYZ. At any rate, even if 'believing that [lakes contain water]' is not quite cointensional with 'believing

that [lakes contain XYZ]', it should now be clear that they overlap to a significantly greater extent than Stich suggests.

I make this point in order to illustrate that, on the cross-world interpretation, Stich's device has the consequence that narrow content ascriptions will include a much broader range of mental states than one might have thought. It now appears that an expert on the difference between water and XYZ who believes that lakes contain water, and an expert who believes lakes contain XYZ, have beliefs with the same narrow content. (To be more precise, they share *a* narrow content, since on Stich's view the same belief may have many different narrow contents.) For such an expert, the difference between believing that lakes contain water and believing that they contain XYZ will be of great significance; a notion of narrow content which groups them together would seem to be of dubious psychological significance. But worse is yet to come.

Other possible worlds can be very strange indeed, with the result that Stich's proposal, on our second reading, will group together states which have *very* little in common. For example, suppose Putnam is correct in thinking that brains in vats with experiences exactly like ours would have beliefs not about water, trees, and the like, but about "electronic impulses and program features" (Putnam 1981, p. 14). I suppose I cannot be a doppelgänger of a brain in a vat. But if what Putnam says is true of brains in vats, it should equally be true of bodies in vats. In that case, when I believe lakes contain water, I have a doppelgänger in a vat who believes that circuit 25 is open (or something of the sort). So my belief has the narrow content [circuit 25 is open]. But actual Earthly computer scientists may also believe that circuit 25 is open, and hence believe that [circuit 25 is open]. So my belief that lakes contain water has the same narrow content as my computer science colleague's belief that circuit 25 is open! This does not look promising. Similarly, there may be worlds in which cats look the way dogs actually look; in some such worlds cats will be called 'dogs'; some of these worlds in turn will contain people who think that there is a cat in front of them who are doppelgängers of actual Earthlings who think that there is a dog in front of them; so my thought that there is a cat in front of me has the same narrow content as my neighbor's thought that there is a dog in front of him.

Such examples are even easier to come by in the case of descriptions of action, which Walker's purportedly narrow descriptions concern. The actions described by most ordinary action descriptions may be accomplished in a wide variety of ways, all of which will be grouped together into the same "narrow" description; moreover, the same actions *could* be accomplished in an even wider variety of ways, thus extending still further the range of quite different behaviors grouped into the same narrow kind. Consider the action description

"greeting your neighbor," which includes a wide variety of utterances, hand gestures, and so on. If this were not bad enough, there are possible worlds in which one greets one's neighbor by doing something indistinguishable from actually feeding the dog, or picking up the newspaper, or even swearing at one's neighbor. But then people who are actually feeding the dog, or picking up the newspaper, or swearing at their neighbors will have doppelgängers in other worlds who are greeting their neighbors. So someone actually feeding his dog will be [greeting his neighbor], by virtue of having a cross-world doppelgänger whose behavior is correctly describable (broadly) as "greeting his neighbor," and his behavior will thus fall into the same "narrow" type as that of someone who is actually greeting his neighbor.

6. A Device for Characterizing Narrow Content

Stich's and Walker's device, then, does not look promising as a useful notion for narrow psychology. Stich recognizes this (though he apparently does not realize *how* unpromising it is). But Stich infers that Fodor's notion of narrow content, which his device is intended to explicate, will not be a useful psychological notion. We have seen that the inference is faulty, since the notion explicated by Stich's device is clearly not what Fodor had in mind by "narrow content."

As it happens, I think that something a *bit* like Stich's apparatus will give a reasonably accurate picture of the sort of thing Fodor has in mind by "narrow content." However, the device I will now discuss does not provide a means of expressing narrow contents, as Stich's device is intended to do. Nothing like Stich's device will provide such a means of expression *unless* there is a function from the embedded sentences in belief ascriptions to the narrow contents of the beliefs they ascribe. On my view, the connection between narrow contents and 'that'-clauses is not close enough to provide such a function. Still, I think the apparatus I will now discuss has some interest independent of the aim of expressing narrow contents.

Stich begins with the insight that any doppelgänger of someone whose belief has a certain narrow content should have the same narrow content. This is surely correct. He also recognizes that it will not do to consider only exact replicas; if the notion of narrow content is to have any use, then it must be possible for two people to have beliefs with the same narrow content even though they are not exact doubles. The trouble with Stich's proposal is a result of the means he uses of determining who, other than my doubles, has beliefs with the same narrow content as mine. We have seen that, depending on how it is interpreted, his answer lets in either too much or too little *and* too much.

A better answer would take more seriously Fodor's description of narrow content as a function from contexts to broad contents. Consider my current mental state. Place a brain just like mine in my situation, and you get someone who believes that lakes contain water. Place a brain just like mine in my Twin's situation (where his situation includes not just his current environment but his past history), and you get someone who believes that lakes contain XYZ. And so on, for a variety of imaginable further situations. The brain state my doppelgängers and I share, then, determines the relevant function from contexts to contents, and so should count as having the same narrow content. What other brain states should count as instances of the same narrow content? *Those which yield the same results as my current actual brain state when plugged into the various relevant situations.* So not just any old belief that lakes contain water has the same narrow content as *my* belief that lakes contain water; rather, only those which, in the Twin Earth situation, yield the belief that lakes contain XYZ. Thus the state of the expert on the differences between water and XYZ who believes that lakes contain water will *not* be included as having the same narrow content as my state, since, plugged into the Twin Earth context, he will not believe that lakes contain XYZ.

In "What is a Belief State?" I developed and defended the account I have just briefly described. The account was not offered as an account of narrow content, however. I was trying to determine which actual and possible individuals should count as being in the same "belief state," where a belief state was intended to be, to employ Daniel Dennett's useful term (Dennett 1982, p. 19), the "organismic contribution to belief." The idea was that one's belief state, in this sense, should depend only on one's intrinsic properties and not on facts about one's environment; and anyone whose intrinsic properties made the same contribution to belief should count as being in the same belief state. Fairly clearly, if the notion of narrow content is to have any purchase, it should characterize belief states in this sense: so the set of actual and possible individuals who share the same belief state should be the same as the set of actual and possible individuals whose beliefs have the same narrow content. Now, Stich does not define narrow content directly; instead, he defines the set of actual and possible individuals whose beliefs have the *same* narrow content. So my attempt to characterize the individuals who share a belief state would seem to be an attempt to characterize precisely the same set of individuals Stich's device is intended to capture. The only difference in the goals of the two devices is that Stich wants a notion of narrow content for particular beliefs, whereas I tried to characterize only what we might call one's total belief state: the organismic contribution to all my beliefs taken together rather than to any particular belief.

I attempted to characterize one's total belief state as follows. Using ts_1 to symbolize my total intrinsic state, I wrote:

I am in the same belief state, no matter what situation I am in, provided that I am in ts_1 . But . . . there are also other total states that I (or others) could be in and be in the same belief state: namely, all those total states that make no difference to my beliefs. A total state would make a difference to my beliefs if there were some situation in which it led to different propositions believed. So those total states that make no difference to what I believe are those for which there is no situation in which they produce different propositions believed than ts_1 (Brown 1986, p. 361).

Identifying a belief state with the set of actual and possible total intrinsic states in which one would be in that belief state, and letting 'P(TS,S)' symbolize the set of propositions determined by plugging a particular total intrinsic state into a particular situation, I added: "The belief state of individual x , at time t , and in world w -- $BS(x,t,w)$ -- is the set of all those TS such that: for all [situations] S , if $P(TS(x,t,w),S) = Q$, then $P(TS,S) = Q$. In something closer to English, x 's belief state (at time t and in world w) is the set of total states that determine, with respect to any situation, the same propositions x 's total state (at t and in w) determines with respect to that situation."

An analogy I used in that earlier paper will help both to motivate this account and to show what is wrong with Stich's proposal. *Weight* is a property which varies from one situation to another: if I were on Mars, I would weigh less than I do; if I were on Venus, I would weigh more. But we can consider the intrinsic contribution to my weight--my "weight state," as it were. Clearly, any doppelgänger of mine is in the same weight state I am in. Just as clearly, things which are not doppelgängers of mine may nevertheless be in the same weight state as me. Which non-doppelgängers have the same weight state I do? The analog of Stich's procedure would be to take everything that has my weight (regardless of its context) and then to add all the doppelgängers of such objects. But this would give us a hopeless jumble of objects. They would include, for example, Earthly doppelgängers of things that weigh on Venus what I weigh here, and Earthly doppelgängers of things that weigh on Mars what I weigh here: so my weight state would be shared by things substantially more and substantially less massive than I am. (If we include cross-world doppelgängers, this procedure will make *everything* share my weight state, since any actual object could in imaginable circumstances share my weight.) By contrast, the analog of the procedure I favor would count as having the same weight state only objects which, placed on Earth, weigh what I weigh here; placed on Mars, weigh what I *would* weigh there; placed on Venus, weigh what I would weigh there; and so on. In short, the analog of the procedure I favor will yield the set

of objects with my *mass* as the objects which share my intrinsic contribution to my weight -- which is just the result we should hope for.

7. Differences Between the Present Account and Fodor's Narrow Content

Three principal differences between my explication of the notion of a belief state and Fodor's explication of narrow content are: that my notion is not intended to be a notion of content; that it does not assign separate belief states for each distinct belief; and that it does not purport to say anything about the linguistic meaning of the embedded sentences in belief ascriptions. I now argue that each of these differences is an advantage of my account.

(1) *Fodor's "narrow contents" characterize belief states, not contents.* It is better to regard Fodor's notion, as well as Stich's and my related notions, as identifying belief states rather than narrow contents. Stich's notion gives us a class of people who, or of mental states which, share a narrow content. But it does not say what content they share, and indeed leaves it an open possibility that there is no content they share. It is better to regard the members of the class as sharing a belief state, and leaving it as a separate question whether it is possible to provide a semantic characterization of belief states. Fodor's notion does not, like Stich's, describe which people or states have a given narrow content; instead it says something about what narrow content *does*, namely, provide a function from contexts to broad contents. But there is no guarantee that such a function will be semantically interesting enough to count as a notion of content. Most discussions of narrow content, including Fodor's, have modeled narrow content on David Kaplan's notion of *character*. Kaplan's characters are semantically very interesting, and do determine a function from contexts to broad contents (see Kaplan 1989a, pp. 505ff). But it hardly follows that any such function is determined by something semantically interesting.

What makes character "semantically interesting?" At least the following facts. First, there are independently identifiable linguistic units (the word 'I', for example) which *have* characters. Second, we can identify the character of a particular expression with a storable *rule* which determines a function from contexts to contents (e.g. the referent of 'I' with respect to a given context is the agent of that context). Third, competent users of the linguistic expression know the rule and intend to use the expression in accordance with it; that is *why* expressions have the character they do. And fourth, character fits neatly into a compositional semantics; the character of a compound expression is a function of the characters of its components. How many of these features are true of Fodor's narrow contents remains an open question. Are there independently identifiable mental units which have narrow contents? Perhaps, if there is a "language of thought." But it is not clear that there is. Certainly no one

can *now* identify the bearers of narrow content. Can narrow contents be stated as rules? If so no one now knows how to do so; that is the reason for Stich's roundabout attempt to characterize narrow content in terms of broad contents and doppelgänger. We certainly cannot explain why mental states have the narrow contents they do in terms of the intentions of agents. This by itself is not a damaging point, of course; it cannot be a requirement on notions of content that they be explicable in terms of the intentions of agents, since intentions themselves have contents. But the more general point is this: we have an account of how linguistic expressions acquire their characters, and we can then use character plus context to explain content. Fodor and company have offered no similar account of what determines narrow content, and as a result we cannot at present use narrow content to explain broad content, but must work backward from broad content to speculate about the possible existence of narrow contents. Finally, it is clear that we are nowhere near an account of narrow content which would yield anything like a compositional semantics of mental states.

I do not mean to imply that there is no semantically interesting notion of narrow content. (I suspect, in fact, that there is.) My point is that we should not regard the notion of a function from contexts to broad contents *as* a notion of narrow content. There *must* be such a function, but it is an open question whether the function is determined by any sort of content, or indeed by any storable rule at all. We should regard Fodor's notion as saying something about what belief states do -- determine what the people in them believe in various possible situations -- while leaving it open whether they do it by virtue of having a certain kind of content.

(2) *Particular beliefs do not have narrow contents.* Both Fodor and Stich want a notion of narrow content which applies to particular beliefs, such as the belief that grass is green or that lakes contain water, rather than to one's entire collection of beliefs taken as a whole. This seems a reasonable thing to want, and the view I have defended so far does not provide it, since I provided only an account of one's belief state as a whole, not of what we might call "partial belief states." However, I doubt that either partial belief states or narrow contents can be parceled out belief by belief.

How would one provide an account of partial belief states, along the lines I have suggested for a total belief state? Perhaps the most natural approach is the following. Suppose we want to know what partial belief state I am in by virtue of which I believe that lakes contain water. Take my actual total intrinsic state, and see what beliefs it yields when plugged into whatever possible situations are of interest. Now, make the minimal change to my total state required to delete my belief that lakes contain water from my total stock of beliefs. Plug

the revised total state into all those possible situations, and see what beliefs are lost. (For example, presumably on Twin Earth what is lost will be the belief that lakes contain XYZ.) Now, consider the set of all total states which determine, in my actual context, the belief that lakes contain water, and which determine, in each other context, the beliefs which were lost in that context as a result of the initial revision of my total state. That set seems to be a reasonable candidate for the partial belief state by virtue of which I believe that lakes contain water, and hence also to be a reasonable candidate for the set of states which share the same narrow content as my belief that lakes contain water.

But in fact I am very pessimistic about whether this procedure will specify a unique partial belief state. First, it isn't clear that there is a unique minimal revision of my total belief state which results in my losing the belief that lakes contain water. The fact that I believe that lakes contain water is a result of an extremely varied and complex set of facts about my mental state; there may be a wide variety of equally minimal but quite different ways of deleting that belief, and each of the different ways might have different effects on what I would believe in other possible situations. Second, you cannot subtract the belief that lakes contain water without subtracting a lot of other related beliefs (though perhaps--the first problem again--no *determinate* set of them), e.g. the belief that you can water ski in lakes, the belief that lake-dwelling fish extract oxygen from the water by means of their gills, and the belief that if you stay below the surface of a lake for too long, you will drown. So it is not clear that the present proposal will yield the belief state by virtue of which I believe that lakes contain water, period, as opposed to the belief state by virtue of which I believe that lakes contain water *and* a variety of other propositions. And so, if there is a narrow content which characterizes the partial belief state in question, it is not clear that it will be the narrow content of the belief that lakes contain water, period, as opposed to the narrow content of a collection of related beliefs. (Notice that you do not need to adopt the implausible view that the content of every belief depends on the content of every other to be holistic enough to see problems with the notion of a partial belief state.)

I have offered an account of what a total belief state is, and said that we should regard it as an open question whether belief states are characterizable in semantic terms -- i.e. whether they have narrow contents. Let us suppose optimistically that there are narrow contents. Then the skepticism I have just expressed about the existence of *partial* belief states raises a pressing question: What use are narrow contents if particular beliefs do not have them? The point of developing a notion of narrow content is to find the right way of formulating psychological laws that appeal to the content of our beliefs; presumably such laws will involve relations between various beliefs. But how can we formulate such laws if we can't specify the narrow content of each belief independently?

How can we see how various beliefs interact unless we can isolate them as individuals?

This is a less serious worry than it may seem. It will seem serious only if our model of psychological processes involves discrete units bumping into each other and pushing each other around like billiard balls on a table. But laws need not be causal, in this simple sense, to be of interest. Suppose that the narrow content of someone's total belief state is something like P & Q & R & S, but that we cannot find plausible discrete entities to play the roles of the beliefs that P, Q, R, and S taken separately. Such generalizations as that people who believe P & Q will usually come to believe T may still be true and interesting. We can regard the contents P, Q, etc. as properties of one's total belief state, and can make generalizations about the relations between such properties, without being able to associate each such property with a distinct partial belief state. (Compare: we can talk about the relation between the temperature and the pressure of a gas without being able to identify them with separate parts of the gas or even regarding them as entirely distinct.) It may be that this is one of the lessons of work on parallel distributed processing, though I don't want to enter *that* tangled controversy here.

(3) *Narrow content is not a kind of linguistic meaning.* Stich, Fodor, others seem to assume that not only mental states but also English sentences will have narrow contents, and moreover that the narrow content of the mental state of someone who believes p will be the narrow content of the English sentence 'p'. For example, in *Psychosemantics* Fodor puts the following question in the mouth of an objector: "Well, since on your view your water-thoughts are content-identical to your Twin's, I suppose we may infer that the English word 'water' has the same intension as its Tw-English homonym . . .?" And Fodor responds, "We may" (49). Similarly, Michael Devitt takes the same notion of narrow content to apply to both mental and linguistic representations. And Joseph Owens, a critic of narrow content, takes defenders of the notion to hold that each sentence is assigned a narrow content "as one of its semantic values," and that, where *N* is the narrow content of a sentence *P*, '*S* believes that *P*' "is deemed true if and only if *S* bears the believing relation to the narrow content *N* and, in addition, the context satisfies some specific conditions *C*" (Owens 1990, p. 171).

I suspect that it is Kaplan's notion of "character," on which, as I have already noted, notions of narrow content have been modeled, which has led writers to think that narrow content should be a different and more basic level of linguistic meaning of the that-clause that ascribes the broad content of a belief. In the case of indexicals, for which after all the notion of character was developed, this may be all right: perhaps something like the character of the

English word 'I' is part of the narrow content of my mental state when I think thoughts whose ascriptions contain the word 'I' in their that-clauses. But in other cases this surely is not the right model. Consider my mental state when I believe a singular proposition (for example, the proposition that London is pretty). As Kripke's example shows, there are different ways of believing this proposition. Pierre believes that London is pretty when the proposition is taken in one way and disbelieves it when it is taken in another. This vividly makes the point that the mental states involved in the two ways of believing that London is pretty are different. Surely an account of the narrow content of our mental states ought to provide a way of capturing this difference in mental state. But the notion of character is the wrong model for capturing the difference. Granted, in Kripke's example and our modification of it, the difference between Pierre's two relations to the proposition that London is pretty has something to do with his differing attitudes to the names 'London' and 'Londres'. But these names do not have different characters, in Kaplan's sense. There might be some point to claiming that the names have different narrow contents *for Pierre*, but a notion of content that must be relativized to particular speakers is not, properly speaking, a notion of *linguistic* content at all.

A wide variety of mental states will qualify subjects as competent users of the name 'London'. I don't think there is any central core which all these states must share and which we could identify as the narrow meaning of 'London'. We might say that the narrow meaning of 'London' *for Pierre* is something like "city I read about in this and that travel book, have longed to visit for years, and which contains Big Ben and Westminster Abbey." And we might make use of this "narrow content" in explaining the state of mind by virtue of which he believes that London is pretty. But we should be clear about the fact that narrow content in this sense is not a semantic feature of 'London' but a psychological feature of Pierre. And recognizing this may free us to recognize that there is no reason the "narrow contents" which explain differences in belief state should be contents *of words* at all. One might have beliefs about London -- believe singular propositions about London -- without having *any* name for London; we should be able to distinguish between the narrow contents of different people who believe singular propositions about London without having names for it; and when we do so we will clearly not look to the narrow content *of a name*.

8. Conclusion

Having criticized Stich's device for expressing narrow content, and having offered my own related device for describing the nature of belief states, I should discuss how the account I have just offered fares with respect to the

problems I have identified with Stich's account. This exercise is a little tricky, since in offering counterexamples to Stich's account I have for the sake of argument presupposed views which Stich accepts but I in fact reject. Nevertheless my criticisms of Stich's account may help to highlight some differences between his account and my own. The first problem with Stich's account was that, in its intra-world version, it individuated narrow content too narrowly, so that my cross-world doppelgänger in the waterless world, and my this-worldly near-doppelgänger who prefers chocolate to vanilla, do not share my belief that [lakes contain water]. My own account is not offered as an account of content *per se*, but as an account of the nature of belief states. So the relevant question for my own account would seem to be whether, on my account, my waterless-world doppelgänger and my chocolate-loving near-doppelgänger share my belief state. In the example of the waterless world, the difference between the two accounts is clear: on my account, my doppelgänger in the waterless world will clearly share my total belief state, since any doppelgänger of mine will, in any situation, have the same broad beliefs I would have in that situation. The case of the chocolate-lover is more complex. Since he is not only not a complete doppelgänger of mine, but in fact has different broad beliefs than I do, he will not share the same total belief state as me. That of course is as it should be; the question about my own account which parallels my criticism of Stich's account is whether he shares the partial belief state by virtue of which I believe that lakes contain water. But since I have rejected the idea that there is a partial belief state for each broad content believed, this question cannot properly be raised about my account. The point about Stich's account is that, if you want an account of narrow content which parcels out narrow contents belief by belief, you should want something the intra-world version of his account cannot provide, namely that a difference in beliefs *prima facie* very remote from the belief that lakes contain water will not result in a difference in the narrow content of the belief that lakes contain water. On my view, by contrast, the expression "the narrow content of the belief that lakes contain water" has no referent; neither belief states nor narrow contents are parceled out in that way.

My second criticism of Stich's account, which applied to both versions but was especially egregious on the cross-world version, was that it individuates narrow contents too broadly. For starters, it counts anyone who believes broadly that, say, London is pretty as sharing a certain narrow content, namely [London is pretty], with all other individuals who have the same broad belief. Again, comparing Stich's account with my own account of belief states is complicated by our different presuppositions. But it should be clear that, on my account, sameness of broad content believed does not guarantee sameness of belief state. Consider the belief state by virtue of which Pierre believed that London was pretty before he moved to London. Now consider a cross-world doppelgänger

of Pierre as he was before moving to London, in a world in which Birmingham is called 'Londres' by the French, while London is called something else, and in which the information Pierre has gleaned from French travel books about the city they call 'Londres' has therefore been about Birmingham rather than about London. Such a doppelgänger of Pierre will believe that Birmingham is pretty rather than that London is pretty. On the other hand, the belief that London is pretty which Pierre acquired while living in London is immune to this particular change in context. Since my account counts belief states as different if they would in some possible situation yield different broad beliefs, it counts someone with the sort of belief that London is pretty Pierre first acquired as being in a different belief state than someone with the sort of belief that London is pretty Pierre later acquired.

The third difficulty with Stich's account was that it did not yield a function from contexts to broad contents. But, like Fodor's notion, my account of belief states is defined precisely in terms of such a function. So my account does not face the sorts of difficulties I have noted with Stich's account.

I have been discussing, as my title indicates, what narrow content is *not*. I would like to conclude with some brief remarks about what narrow content *is*. A better model for the relation between narrow and broad content than that between character and content is the relation between belief *de dicto* and *de re* - or, a bit more accurately, the closely connected relation between beliefs in singular propositions and the beliefs in general propositions on which they depend. In such a case, the sort of content to which we need to appeal is not a different kind of content possessed by the sentence which forms the embedded that-clause of a belief attribution, but rather a content of the same kind but one which is not a content of that sentence. We believe singular propositions not "directly," as it were, but indirectly, by virtue of believing related general propositions. The general beliefs by virtue of which we believe singular propositions provide a better characterization of our belief states than the singular propositions do, and hence come closer than the singular propositions to capturing the narrow content of those belief states. Moreover, the general propositions, together with principles about the conditions under which *de dicto* beliefs give rise to *de re* ones, will yield a function from contexts to broad contents, and so will do what, on Fodor's account, narrow contents are supposed to do -- *without* being a kind of special content analogous to Kaplan's character.