

Ruritania Revisited¹

Ned Block
MIT

Perhaps you are wondering what I mean by ‘holism’. After all, everyone seems to use the term in a different sense. Even if we restrict ourselves to holism of meaning and content, we have many different holisms. Some take holism about meaning to be the doctrine that if you’ve got one meaning, you’ve got lots of them.² On other views, to say meaning is holistic is to say that the meaning of each term depends on the meanings of all or most other terms.³ Others take meaning holism to be the doctrine that there is no real distinction between language and theory or between the “dictionary” and the “encyclopedia”.⁴

Although everyone seems to favor a different definition of meaning holism, there is widespread agreement that some versions of meaning holism are *extremely implausible*, and for a specific reason, namely that they lead to the following: Meaning depends on belief; if any of my W-beliefs change (that is, if any of the beliefs that I would standardly express using the word W change), then my word W changes meaning. And if any of my W-beliefs differ from your W-beliefs, then what you mean by W ≠ what I mean by W. Now suppose I accept and you reject “The Pantheon’s lead was used

¹This paper is very similar to my “An Argument for Holism”, forthcoming in *The Proceedings of the Aristotelian Society*. It appears here with the kind permission of The Aristotelian Society. The argument of the paper is descended from Hilary Putnam’s famous Ruritania argument in “Computational Psychology and Interpretation Theory”, in *Realism and Reason, Philosophical Papers Vol 3*, (Cambridge:Cambridge U. Press, 1983.) I published a very short version of the revised Putnam argument in “What Narrow Contents Are Not” in B. Loewer and G. Rey, *Fodor and His Critics*,(Oxford: Basil Blackwell, 1991). I discovered later that Hartry Field had come up with a similar revision of Putnam in an unpublished paper. I am very much indebted to Field’s paper and to conversations with Field; the paper could reasonably have both of our names on it. Field rejects the conclusion, arguing that instead one of the premises (what I am calling Field’s Principle) should be rejected. (Maybe it should be Field’s Anti-Principle to make it clear that he rejects it.) Versions of this paper were given at a conference in honor of Tyler Burge in Vancouver, in the Fall of 1993; and at the following conferences and meetings in the summer and spring of 1993: the NEH Summer Institute at Rutgers, the meeting of the Sociedad Filosofica Ibero-Americana in Tenerife and at a conference at the University of Maryland. I am grateful to audiences at those occasions. I hope to write a longer version to be included in a volume of papers from the conference in honor of Burge.

²Jerry Fodor and Ernest LePore, *Holism: A Shoppers’ Guide*,(Cambridge: MIT, 1992)

³Michael Devitt, “A Critique of the Case for Semantic Holism”, in *Philosophical Perspectives 7, Language and Logic*, 1993. (Atascadero: Ridgeview: 1993). Devitt gives some objections to the line of argument I give here which are keyed to views that do not appear in this paper. I will discuss some of Devitt’s points in the longer version; in my view, the immunity of the points in this paper to Devitt’s criticisms show that the criticisms do not get at the heart of the argument. See also Devitt’s *Coming to Our Senses: A Naturalistic Program for Semantic Localism*, forthcoming.

⁴Quine is the leading exponent of these views; sometimes he is taken to say that the unit of meaning is the whole language rather than words or sentences, but this form of words has a variety of interpretations, some of which are the ones just mentioned in the text. Putnam has used ‘meaning holism’ to mean that meanings have an identity through time, but no essence. See for example *Representation and Reality*, (Cambridge: MIT, 1988), but this notion seems very different from the family of ideas mentioned in the text.

to make a canopy.” It follows that we don’t share meanings of any of the words in this sentence.⁵ So the meaning of the sentence that I accept isn’t the same as the meaning of the sentence that you reject. So how can people ever disagree? Moreover, we may both accept “Lead is heavy”, but since we don’t mean the same by “lead”, agreement is problematic too. Further, if I accept a sentence and later reject it, then the meaning of what I accept is not the same as the meaning of what I later reject, so how can I ever change my mind?

I propose to avoid the issue of whether one or another version of holism really has this horrible consequence by simply *defining* ‘holism’ as a version of the horrible consequence. Let holism be the doctrine that any *substantial* in difference in W-beliefs, whether between two people or between one person at two times, requires a difference in the meaning or content of W. I propose to argue *for* the horrible consequence, that is, for holism in this sense.⁶

Of course, the interest of the conclusion hinges on the interpretation of ‘substantial’. The argument itself will fill in what I have in mind. I hope you will agree that I am not putting the kind of restrictions on ‘substantial’ that would make the thesis uninteresting.

So far I have not said what I mean by ‘meaning’ or ‘content’. I will use ‘meaning’ and ‘content’ more or less interchangeably.⁷ Meaning and content in this paper are *narrow* meaning and content, if such there be. I will not assume that there is such a thing. Rather, I am arguing for a conditional: if there is such a thing as narrow content, it is holistic in the sense described. One person’s *modus ponens* is another’s *modus tollens*, so the upshot for some readers may be that narrow content is holistic, and for others that there is no such thing as narrow content. .

At this point it would be nice if I could tell you what narrow content is. All I will say is this: Narrow content is internal content, content “inside the head”. Arguably, some beliefs supervene on the non-relational, physical properties of the body. These are beliefs that I necessarily share with any doppelganger, any molecular duplicate of myself, no matter how different the duplicate’s environment or language community. For example, perhaps the belief that $2+5=7$ is one I must share with any doppelganger. If so, then the belief is narrow and its content is a narrow content. Some proponents of narrow content suppose that every belief has both wide and narrow content, the narrow content

⁵This needs qualification. What follows is that your ‘lead’ differs in meaning from my ‘lead’, your ‘canopy’ differs from mine, etc. But your ‘lead’ could still mean the same as my ‘canopy’ and conversely. Then it would be possible that we didn’t disagree at all.

⁶Jane Heal argues that this doctrine (as well as a version of the third one mentioned above) can be ascribed to Fodor and LePore. See her “Semantic Holism, Still a Good Buy” in *Proceedings of the Aristotelian Society* XCIV, 1994, 325-340.

⁷I tend to use ‘meaning’ in connection with words and ‘content’ with sentences.

being what is needed for explanation of behavior; others suppose that beliefs have only narrow content.⁸

Perhaps you are disappointed with this meager account. I don't say more for two reasons. First, I don't have an account of narrow content. Second, even if I did, I'd be reluctant to offer it because my argument does not depend on any specific theory of what narrow content is, but rather only that it is narrow and explanatory (and other conditions to be spelled out below). To give an account would be to encourage the idea that my argument depends on it.

If you believe that content *is* narrow content or that there is both narrow and wide content, then presumably you will regard the conclusion that I will argue for as interesting (though it would be rational to wait to see what the "substantial" qualification comes to). But even if you are agnostic about narrow content or feel at sea about what it might be, still the conclusion ought to be of interest to you. If we can establish for sure that any content that is narrow is holistic, that may help us to think about whether wide content is holistic or why it has seemed to be holistic. There has long been a gulf between holists and anti-holists. Perhaps it will turn out that they have been talking past each other: holists are right about one thing, anti-holists about another.

THE PREMISES

The assumptions made so far are relatively uncontroversial and will remain in the background. Here are the assumptions that will be in the foreground:

NARROWNESS

Narrow content supervenes on non-relational physical features of the body. Or in slogan form, narrow content is narrow. This is just a definition.

DIFFERENCE

If at one time, a person has substantially different beliefs associated with term t_1 and t_2 , then t_1 differs in narrow content from t_2 for that person at that time. So for any normal person, words like 'cat' and 'dog' and 'panda' have different narrow contents.

EXPLANATION

Narrow content's main purpose is its role in psychological explanation.

INCOMPLETE UNDERSTANDING

Incomplete understanding and full mastery of a concept are completely compatible. This idea should be familiar from the work of Tyler Burge. I can have full mastery of the concept of arthritis, so that it is correct to ascribe to me beliefs such as that arthritis is a disease that I expect but don't want, even if my understanding of the concept is incomplete. I will also assume that where there is a wide concept there is a narrow concept.

INTER/INTRA: FIELD'S PREMISE⁹

⁸ Jerry Fodor and David Lewis argue for both narrow and wide content, but they give rather different accounts of why one should accept narrow content. See Fodor's *Psychosemantics: The Problem of Meaning in Philosophy of Mind*, (Cambridge: MIT Press: 1987) and Lewis, "Reduction of Mind" in Samuel Guttenplan (ed), *A Companion to Philosophy of Mind*, (Oxford: Basil Blackwell: 1994). Fodor now rejects narrow content. See his *The Elm and the Expert*, (Cambridge: MIT: 1994).

The relation of *same narrow content* that holds between people is the same relation of *same narrow content* that holds within a single person. So if my word X has the same narrow content as my word Y, and my word Y has the same narrow content as your word Z, it is legitimate to conclude that X is the same in narrow content as Z.

These principles are not independent. For example, one of the roles of EXPLANATION is to bolster DIFFERENCE. If a theorist holds that a single person's 'dog' and 'cat' have the same narrow content, one should wonder what this theorist thinks narrow contents are *for*. Any narrow contents that are usable for psychological explanation will have to be more fine-grained than that.

THE EXAMPLE

The argument will be based on a version of Putnam's Ruritania example. There are two parts of Ruritania, B and W. Bruce lives in B, Walter, his twin, lives in W. The dialects of the two parts are exactly the same save for the fact that in B, they use 'grug' instead of 'beer', whereas in W, they use 'grug' instead of 'whiskey'. The B dialect lacks 'beer'; the W dialect lacks 'whiskey'. Ruritanian = English except for the use of 'grug' in the two dialects. So in B, 'whiskey' means whiskey and in W, 'beer' means beer. At age 10, Bruce and Walter alike in every relevant respect. We could suppose that they are molecular duplicates. In particular, the difference between the different substances known as 'grug' in their communities has not impinged on them at all. They share all beliefs, images, recognitional dispositions and the like having to do with the referent of 'grug' in their communities. If asked about grug, the following is as much as could be squeezed out of them:

- "Grug is a brownish liquid."
- "Drinking grug makes grownups act funny."
- "Grownups like to drink grug at social occasions."
- "Grug is bought in "liquor" stores."
- "Grug is often served before dinner."
- "It would be peculiar to drink grug with breakfast."

Two years pass in which Bruce and Walter become more integrated into their societies, learning more about the different items called 'grug' by their language communities and the connection between those terms and English. Here is what they know at age 12: Bruce could give voice to the following:

1. "'Grug' translates in English to 'beer'."
2. "Grug comes in small cans (in both parts of Ruritania)."
3. "More than six cans makes people very drunk."
4. "Grug is relatively cheap."

⁹Field points out that this premise is assumed in the argument and argues that we should reject it rather than accept the conclusion

5. “‘Grug’ in the other dialect is characterized by 1-4 below.”

Walter could give voice to the following:

1. “‘Grug’ translates in English to ‘whiskey’.
2. “‘Grug’ comes in liter bottles (in both parts of Ruritania).”
3. “One glass knocks you out.”
4. “‘Grug’ is expensive.”
5. “‘Grug’ in the other dialect is characterized by 1-4 above.”

Both twins are bilingual, and they have all the same beliefs--at least they would utter all the same sentences--except for indexical ‘grug’ beliefs. In contexts in which B-terminology is appropriate, both say “Grug is cheap and comes in 6-packs.” In contexts in which W-terminology is appropriate, both say “Grug is expensive and comes in liter bottles.”. In contexts in which neither dialect is singled out, they use indexicals. Walter says “‘Grug’ is *our* word for whiskey,” Bruce says “‘Grug’ is *their* word for whiskey.”

THE ARGUMENT

The basic idea of the argument is simple. At age 10, Bruce’s ‘grug’ has the same narrow content as Walter’s ‘grug’. But at age 12, the narrow contents of their native or home ‘grug’s differ. So at least one must have changed, and symmetry requires that both have changed. Why is any further argument needed? The first premise (the identity at 10) depends only on the definition of ‘narrow’. But what is the justification of the second premise, that the narrow contents are different at age 12? If we have wide content in mind, the second premise will seem just obvious--it is our practice to take reference and beliefs into account in deciding about translation, and in this case both the reference and the beliefs are different, and further, there is a competitor: each twin’s foreign ‘grug’ seems the right translation of the other twin’s home ‘grug’. But if it is wide content we have in mind, the first premise is false.

Concentrating on narrow content, it may still seem obvious that the twins’ home ‘grug’s differ in narrow content at age 12. The twins’ home ‘grug’ beliefs are very different. E.g., using their home ‘grug’s, one says “Grug is expensive,” and the other says “Grug is cheap.” As I said above, in regard to a different matter, if we are to regard these as the same in narrow content, what would narrow content be for? Any narrow contents usable for psychological explanation will have to be more fine-grained than that. But now it looks as if there is very little difference between the premises of the argument and the conclusion. After all, ‘grug’ at age 10 and ‘grug’ at age 12 differ in associated beliefs too. If belief differences between people count for differences of narrow content, why shouldn’t belief differences between two stages of one person count as well? And then we could dispense with the argument altogether!

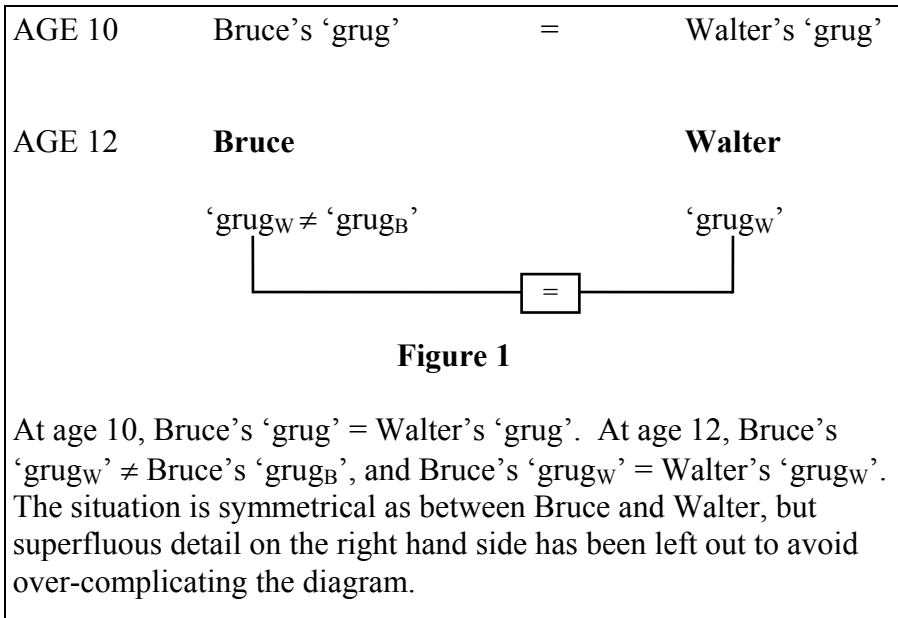
So the purpose of the added complexity in the argument to follow is to justify the claim that at age 12 their home ‘grug’s differ in narrow content. The idea is to justify this inter-personal claim by appeal to an *intra-personal* claim, the claim that each twin’s

two ‘grug’s differ in narrow content from each other, a consequence of the DIFFERENCE principle.¹⁰ But isn’t this just a matter of a difference in beliefs too? There is an important difference, but I’ll wait until later to say why.

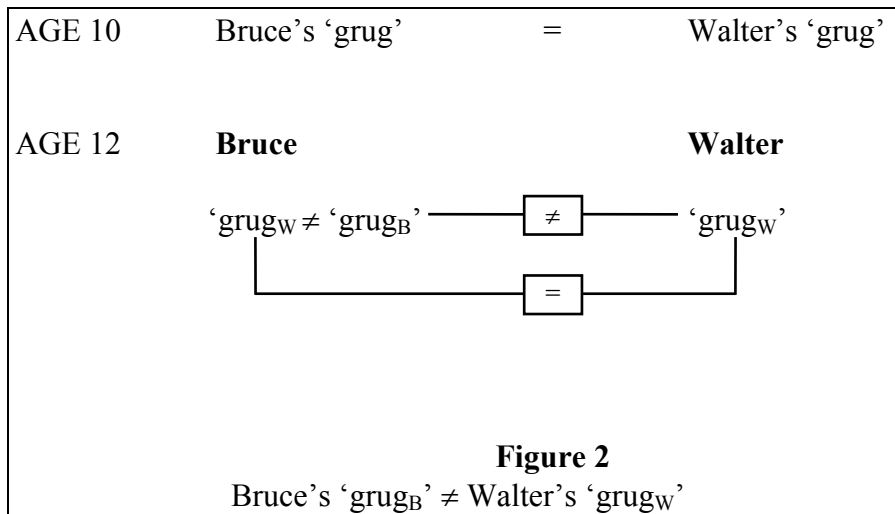
One final preliminary: the argument is complicated, so to make it easier to follow, I will adopt a simplification. I will speak of sameness and difference of narrow meaning and content as identity and difference, simpliciter, using ‘=’ and ‘≠’ as symbols. ‘Word₁ = Word₂’ is to be understood as saying that the two words have the same narrow content.

- I. At age 10, Bruce’s ‘grug’ has the same narrow content as Walter’s ‘grug’. At age 10, Bruce and Walter are molecular doppelgangers, so all their words have the same narrow content, by the principle that narrow contents are narrow (NARROWNESS). But do the boys’ ‘grug’s have narrow content at all? Here is where the INCOMPLETE UNDERSTANDING principle comes in--to justify the claim that they have narrow contents despite a large measure of ignorance.
- II. At age 12, both boys understand the home ‘grug’ and the foreign ‘grug’, and they attach different narrow contents to them. Bruce’s ‘grug_B’, that is, his word ‘grug’ in his home dialect used to mean ‘beer’, has a different narrow content from his ‘grug_W’, the word that as a bilingual he uses to mean what ‘grug’ means in the other dialect, namely, whiskey. As explained above, I will put this by saying that for Bruce, ‘grug_B’ ≠ ‘grug_W’. (See Figure 1) Here I appeal to the DIFFERENCE principle, the principle that says that substantial differences in belief associated with two terms at one time make for different narrow contents, as with ‘cat’ and ‘dog’ for a normal speaker. For Bruce, ‘grug_B’ and ‘grug_W’ are as different in narrow content as your ‘whiskey’ and ‘beer’. For example, he knows that ‘grug_B’ translates in English to ‘beer’, and ‘grug_W’ translates in English to ‘whiskey’.
At age 12, one boy’s foreign ‘grug’ is the same as the other’s home ‘grug’. Bruce’s ‘grug_W’ = Walter’s ‘grug_W’, that is, Bruce’s ‘grug_W’ has the same narrow content as Walter’s ‘grug_W’. (See Figure 1.) Recall that Bruce and Walter are doppelgangers at age 12 except for indexical ‘grug’ beliefs having to do with *whose* ‘grug’ is in question. Except for the indexicals, they are the same with respect to the word ‘grug’ used to mean whiskey. I appeal to the NARROWNESS principle, and to the EXPLANATION principle to justify the claim that the indexical difference doesn’t make a difference to narrow contents. (More on this later.) This is the step in the argument that requires the example’s two ‘grug’s. If not for the need to make Bruce and Walter doppelgangers, I could have run the argument with different words for the two ‘grug’s.

¹⁰My 1991 rendition of the argument (mentioned in footnote 1) involves an unhappy compromise between the simple and the complex versions, in which I attempted to justify the second premise by appeal to the intra-subjective difference between ‘beer’ and ‘whiskey’

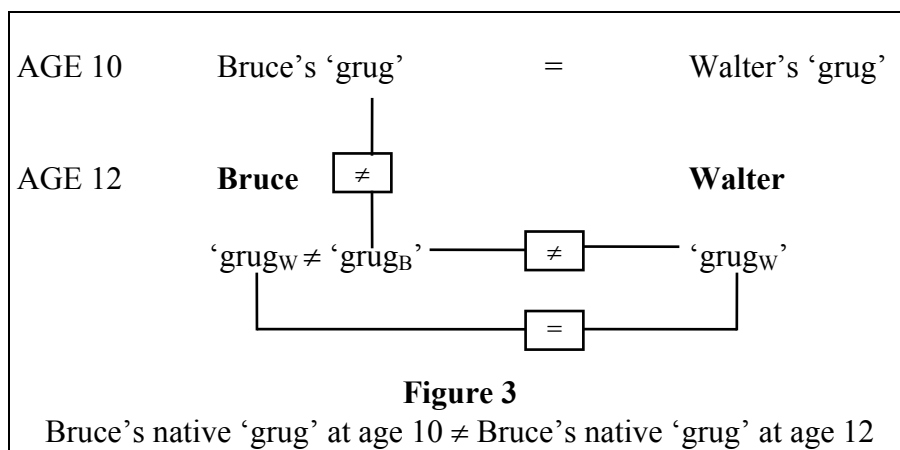


III. It follows that at age 12, one boy's home 'grug' ≠ the other boy's home 'grug'. Since Bruce's 'grug_W' ≠ Bruce's 'grug_B', and since Bruce's 'grug_W' = Walter's 'grug_W', it follows by the logic of identity (or rather identity of narrow meaning or content) that Bruce's 'grug_B' ≠ Walter's 'grug_W'. (A representation that's not synonymous with one of a pair of synonymous representations can't be synonymous with the other either, and the same holds for narrow synonymy.) See Figure 2.



IV. At age 10, Bruce's 'grug' = Walter's 'grug'. At age 12, Bruce's 'grug_B' ≠ Walter's 'grug_W'. It follows by the logic of identity that either Bruce's 'grug' at age 10 ≠

Bruce's 'grug_B' at age 12 or Walter's 'grug' at age 10 ≠ Walter's 'grug_W' at age 12. And since there is no asymmetry in the details of the case and therefore no reason to treat one child differently from the other, *both* Bruce's 'grug' at age 10 ≠ Bruce's 'grug_B' at age 12 *and* Walter's 'grug' at age 10 ≠ Walter's 'grug_W' at age 12. See Figure 3. But Bruce's 'grug_B' at age 12 just is the word 'grug' in his dialect; 'grug_B' is Bruce's home 'grug'. So 'grug' in Bruce's dialect changed narrow meaning between age 10 and age 12. And the same for Walter. So a substantial change in Bruce's 'grug' beliefs results in a change in narrow content of his 'grug'. Further, there is nothing unusual about this change. It involves just the sort of change in belief that we often undergo. Reading the New York Times can induce this sort of change in a single sitting.



Note that the argument assumed no particular *theory* about what narrow content is. Some say that narrow content is functional role, others that it is observational content, others that it is specified by a function from contexts of language acquisition to the contents acquired in those contexts and others that it is the same as wide content. But I have not assumed the truth or falsity of any of these views. Of course, the list of assumptions given above does put some restrictions on what narrow content could be, but I hope these restrictions will not be controversial.

ESSENTIAL AND ANALYTIC PROPERTIES

In arguing against Putnam on Ruritania in *Psychosemantics*¹¹, Fodor suggests that holism can be avoided. He says "Learning what anything really is changes one's narrow concept of that thing", but learning other sorts of features do not change narrow concepts. The suggestion is that we can avoid holism by distinguishing between two classes of truths--those that attribute some sort of essential property (Fodor prescind from declaring allegiance to any particular form of linguistic or metaphysical essentialism) and those that don't. Learning the former features do change narrow concept, but that is unexceptionable, and learning the latter features do not. But this suggestion is powerless to avoid my argument, for in my example Bruce and Walter do not learn anything that

¹¹*Psychosemantics*, op. cit., pp. 94-95.

might count as learning what “grug” really is. Here’s what beer really is: a fermented alcoholic beverage, brewed from malt and flavored with hops that is less than 20% alcohol. Here’s what whiskey really is: an alcoholic beverage distilled from fermented mash of grain (corn, barley, rye), aged in wood, and roughly 40%-50% alcohol. Note that none of the things that Bruce or Walter learn are very closely connected to these facts.¹²

INDEXICAL OBJECTION

In III, I said that Bruce’s ‘grug_w’ = Walter’s ‘grug_w’. I noted that the twins are just alike with respect to ‘grug_w’ except for indexical beliefs, appealing to NARROWNESS, the principle that says that narrow content is supervenient on the body. Though Bruce and Walter are no longer perfect duplicates by age 12, they are the same with respect to ‘grug_w’. Well, almost the same. There is the indexical difference I mentioned. But does the indexical difference make a difference to psychological explanation? If not, EXPLANATION, the idea that the main purpose of narrow content is psychological explanation, dictates that the indexical difference doesn’t make a difference. But now I’m in trouble, for indexical differences are famous for making an explanatory difference. Famously, if I think my pants are on fire, I jump into the pool, but if I think your pants are on fire I push you in.

True, and there will be important behavioral differences between the twins that hinge on the indexical difference. If they are both told that the part of the country in which ‘grug’ is used for whiskey has been invaded, Walter, but not Bruce, will be worried about the fate of his house and family. But the difference comes from the indexicals, not from the narrow content of ‘grug’ itself. Bruce says “‘Grug’ is their word for whiskey,” whereas Walter says “‘Grug’ is our word for whiskey.” It is the indexical difference that makes the difference.

SPLITTING OBJECTION

Does the argument commit a well-known fallacy about the notorious “splitting” cases made famous in recent discussions of personal identity? Suppose that next year you split into two persons more or less just like you are now. (Each successor has half of your cells combined with duplicates of the other half.) Call the two new people ‘A’ and ‘B’. $A \neq B$, since A and B occupy distinct places at one time. But then you can’t be identical to *both* A and B, since one thing can’t be identical to non-identical things. And since there is no relevant difference between A and B, $you \neq A$ and $you \neq B$. But no one should conclude that the mere possibility of a split shows that I am not the same person who wrote the first word in this sentence. Even if an actual split undermines identity over time, that does not show that the mere possibility of a split does so. That would be a fallacy. But is my argument an instance of the same fallacy? How can the mere possibility of a split of ‘grug’ into ‘grug_B’ and ‘grug_w’ show that cases of substantial changes in belief without any such split are changes of narrow meaning or content?

¹²They do learn the translation of ‘grug’ into English, but learning that isn’t learning an essential property unless they know an essential property under the English description. I make use of their knowing the translation only in making them perfect twins at age 12 aside from indexical beliefs. This helps to motivate the idea that one boy’s foreign ‘grug’ has the same narrow content as the other’s home ‘grug’.

But the ‘grug’ example, spelled out in the way I would fix on in response to this objection, is not a case of ‘grug’ splitting into ‘grug_B’ and ‘grug_W’. Bruce’s word ‘grug’ at age 10 is his native or home ‘grug’. I have called his home ‘grug’ at age 12 ‘grug_B’, but this terminology should not mislead us. These are the same words at different times. Words can maintain an identity over time just as people can. (Note that I have temporarily abandoned my terminology of talking about identity of narrow content in terms of identity of words; ‘grug’ at age 10 is the same word as ‘grug_B’ at age 12, but they have different narrow contents.) To use a popular metaphor, he has a ‘grug’ “file” at age 10. As he learns more about beer, he puts more information in his ‘grug’ file. At age 12, he has quite a bit of information in his ‘grug’ file--“Grug is cheap,” “Grug comes in 6-packs”, etc. The ‘grug’ file at 10 is the same file as the ‘grug_B’ file at age 12. When he learns about whiskey, he opens a new file, the ‘whiskey’ file. I hereby stipulate that this all happens before he learns about the other part of Ruritania or their dialect. When he learns about the W-Ruritanian dialect, and the fact that they use ‘grug’ to mean whiskey, that’s just another bit of information in his ‘whiskey’ file. Of course, I’m filling in details that were left vague in the original description so as to fit this case, but that is perfectly legitimate because it does nothing to weaken the original argument and if I am right, it allows a firm rejection of this objection.

Suppose you find out that the term ‘beer’ is used in Outer Mongolian to mean whiskey. Does that make your ‘beer’ “split” in any way that raises one of these identity over time cases? I assume not, and the same is true for our twins. I could have run the example without Bruce and Walter having learned the word in the other dialect of Ruritanian at all. My purpose in running the example the way I did was to make it easier to justify the idea that Bruce’s ‘grug_W’ = Walter’s ‘grug_W’. I appealed to the idea that they are exactly alike except in indexical beliefs. But even if the “exactly alike” had to be weakened a bit, I think that premise would not be very much weakened.

SMALL DISTANCE OBJECTION

I mentioned earlier that there is a simple version of the argument that suffers from an excessively small difference between premises and conclusion. The simple version is: the twins’ native dialect ‘grug’s are the same in narrow content at age 10 and different at age 12. So at least one must have changed, and given symmetry considerations, both changed. The problem is: What’s the justification of the claim that the twins’ home ‘grug’s are different at age 12? Well, its true that they have very different beliefs connected with them (e.g. “expensive” vs “cheap”) But if differences in belief make for a difference in narrow content, why not just appeal to that directly to support the difference between 10 and 12 and dispense with the argument altogether? Further, the DIFFERENCE principle just appeals to a difference in belief, and so doesn’t it just beg the question in the same way?

Let me reply by mentioning a justification for the DIFFERENCE principle. So far, all I have said is: what good would narrow contents be if a person’s ‘cat’ and ‘dog’ (at

one time) had the same narrow content? But doesn't that apply between people and within a person over time too? We can see a difference between the other cases and the intra-personal case at a time by considering what we might call incoherent cognition. By this phrase, I mean the kind of confused thinking and acting that a person engages in if he has contradictory beliefs. If I accept "Pandas are cute" and "Pandas are not cute", that could cause and causally explain incoherent cognition of a sort that would not arise in someone who accepted instead "Pandas are cute" and "Cats are not cute". 'Panda' has the same narrow content in both occurrences, so "Pandas are cute" and "Pandas are not cute" are incompatible contents. 'Panda' and 'cat' have different narrow contents, so "Pandas are cute" and "Cats are not cute" are not incompatible contents. The file metaphor may help: "is cute" and "is not cute" are fine in different files, problematic in the same file. In the same file, they lead to incoherent cognition. So here's why the DIFFERENCE principle is right: different beliefs require different files and different files involve different narrow contents. Note that no such reasoning will apply in justifying an inter-personal version of the DIFFERENCE principle or one that applies to two stages of a single person. Incoherent cognition only operates in the intra-personal synchronic case.

But wait! The file explanation doesn't depend on there being different *beliefs* in the two files. Suppose I have two different files headed 'panda'. Both contain "Furry" and "Found in Asia" and "Not identical to the other animal called "panda". Both files have the same beliefs.¹³ In fact, this does describe my epistemic situation some years ago. (Now I know more--that one is the great panda and the other is the lesser panda.) All that we need for different narrow contents in a single person are different *files*--the contents can be the same.¹⁴ So a line of thought which seemed to back up the DIFFERENCE principle actually appears to argue for something much stronger, something like this: if at one time a person thinks Xs are distinct from Ys, then 'X' differs from 'Y' in narrow content for that person at that time. Differences in beliefs are only relevant because it would be hard for one to have 'X' beliefs that are different from one's 'Y' beliefs without thinking that Xs are distinct from Ys.

The upshot is that I could have motivated the difference in narrow content between the two 'grug's within each twin at age 12 by appealing just to the fact that each twin thought the two words picked out different things. There was no real need to appeal to a difference in beliefs. So if the DIFFERENCE principle is reformulated as just indicated, the premise that says that a twin's two 'grug's at age 12 are different in narrow content does not depend on a difference in beliefs. Hence the step in the argument that

¹³Why should we say two files that are the same instead of one file that is written down redundantly? Even if there is no functional difference between the two files there could still be a functional difference between two files and one file.

¹⁴See Putnam's 'elm/'beech' point, Kripke's 'Paderowski' example, and Loar's 'chat' example. Putnam, H., "The Meaning of 'Meaning' in *Mind, Language and Reality*, Philosophical Papers Vol 2, (Cambridge: Cambridge University Press, 1975). S. Kripke, "A Puzzle About Belief" in *Meaning and Use*, A. Margalit (ed), (Dordrecht: D. Reidel: 1979). B. Loar, "Social Content and Psychological Content" in *Contents of Thought: Proceedings of the 1985 Oberlin Colloquium in Philosophy*, R. Grimm and D. Merrill (eds), (Tucson: University of Arizona Press, 1987).

says that the twins' native 'grug's are different in narrow content at age 12 is suitably distant from the conclusion about change over time.

FIELD'S PRINCIPLE

Thus far the role of Field's principle (that is, the principle Field points out is required for the argument but, according to Field, should be rejected), INTER/INTRA has been mainly in the background. This principle says that the intra-personal relation of *same narrow content* is the same as the inter-personal relation of *same narrow content*. One of the places in which the principle was used was in the last stage of the reasoning. At age 10, the native 'grug's are the same. At age 12, the native 'grug's are different. So one or both must have changed between 10 and 12. The content identity in the premises is interpersonal, whereas the content identity in the conclusion is intrapersonal. Are these the same identity relations? Field says no, but I disagree.

Here is one possible justification for the claim that intra-personal content identity is the same relation as interpersonal content identity. Any relation between representations relevant to psychological theory that can obtain between representations of 2 people can also obtain between a representation of or in a person at one time and a representation of the same person at another time. Consider, for example, a difference, D, between representations of Oscar and Elmer that explains why a given stimulus causes Oscar to do one thing and Elmer to do another. D can also explain in just the same way why Oscar does one thing at one time and another thing at another time. From the point of view of psychological theory, any explanatory difference or similarity that can be found operating between representations in two people can also be found operating between two stages of a single person.¹⁵

This justification is fine as far as it goes. But it doesn't go far enough because it leaves out an important case. There are three relevant content identity relations:¹⁶

- Intra-personal + Diachronic
- Inter-personal + Synchronic
- Intra-personal + Synchronic
-

The reasoning in step V just described includes the first two of these. Bruce and Walter are relevantly the same at one time, different at another time, and a conclusion is reached about a difference within each over time. I believe that the justification mentioned can show that this reasoning is OK. But there is no point in going into the matter, because the reasoning in steps II, III and IV includes the latter two relations, and the justification given does nothing to show that these are the same.

In II, I said that Bruce's two 'grug's had different narrow contents. This is intra-personal and synchronic. In III, I said that one's foreign 'grug' = the other's home 'grug':

¹⁵The converse is not required for the argument. I do not claim the converse, because I think there are psychologically interesting principles of representational change over time. So what I hold is that interpersonal representational relations are a subset of intra-personal relations over time.

¹⁶I am indebted to discussions with Field and Brian Loar.

inter-personal and synchronic. And in IV I concluded that at age 12, their home ‘grug’s are different: again inter-personal and synchronic. Both the premises and the conclusion are synchronic.

Is the last of the three content identity relations listed above the same as the first two? Well, there certainly is a large difference. As I pointed out, no difference in beliefs is required for intra-personal synchronic narrow content difference. All we need are two files, even if the same things are written in both of them. Beliefs seem much more relevant to the other types of comparisons. Consider, for example, the rationale for claiming that Bruce’s ‘grug_W’ has the same narrow content as Walter’s ‘grug_W’ rather than Walter’s ‘grug_B’. The two ‘grug_W’s have all the same non-indexical beliefs associated with them, whereas one twin’s ‘grug_W’ is the same as the other’s ‘grug_B’ only in indexical beliefs.

So are the narrow content identity relations the same or not? If there is a well-defined intra-personal narrow content identity relation and a well defined inter-personal relation, trivially we can compose these relations can get a well-defined narrow content identity relation that is both inter and intra-personal. But that doesn’t answer the question. Consider the done-by-one-person relation, $S_{XY} \leftrightarrow X$ and Y are done by the same person. Consider the done-by-cousins relation $C_{XY} \leftrightarrow X$ and Y are done by cousins. The first is internal and the second is external. We can compose them to get the a relation that holds of a pair of actions if and only if they are done by the same person or by cousins. But this trivial exercise sheds no light on whether C and S belong to the same kind. To say that X and Y are the same in narrow content is to say that each has a narrow content, and that they are identical. So the question of one versus two (or three) identity relations is really a question of whether there is more than one *kind* of narrow content. If there is more than one identity relation, then the twins’ ‘grug’s have two or more kinds of narrow content, one of which governs intra-personal relations at a time, another of which governs inter-personal relations.

If the kind of narrow content that comes into intra-personal synchronic comparisons is different from the kind that comes into the other comparisons, the sting goes out of the problems of holism that I started with. We were worried about the idea that you can’t change your mind and that two people never disagree. But if there are two or more different kinds of narrow content, each appropriate to different comparisons, we have departed so much from common sense that we can hardly expect common sense ideas to apply.

So I draw a disjunctive conclusion: If there is one kind of narrow content, it is holistic. If there are more, then maybe holism is true but it loses its sting.¹⁷

¹⁷I am grateful to Paul Boghossian and Paul Horwich for comments on a previous draft, to Tyler Burge, Martin Davies, Michael Devitt, Jerry Fodor, Brian Loar and Hartry Field for helpful comments at meetings at which earlier versions were delivered, and to comments from the audience in Robert Stalnaker’s and my graduate class, especially to Alex Byrne, Ned Hall, Diana Raffman, Robert Stalnaker and Daniel Stoljar.