

The Case Against Holism Reconsidered

Jonathan Cohen*

15 April 1996

I want to be very clear about this. I am not now, nor have I ever been, a member of the Communist Party, and I do not believe in meaning holism. In fact, I think there are good reasons to refuse to accept meaning holism. However, there are also bad reasons to refuse to accept meaning holism. The trouble is, I'm afraid some of the more popular motivations for rejecting holism fall into the latter camp.¹ As a responsible anti-holist, I think the best strategy is to own up to the failures of some of these motivations before holists embarrass us by noticing them first. This will be my project in the present essay.

Here's what I'm not going to do. I'm not going to offer a plausible version of holism. In fact, I don't think there are any plausible versions of holism. But I hope to convince you that some of the intuitions which are sometimes offered against the view are less decisive than has been supposed.

The motivations for rejecting meaning holism which shall be my focus here concern interpersonal communication, intrapersonal endurance of concepts, and the rational justification of intrapersonal inference. Meaning holism is the view that the meaning of a concept is determined by its relations to all other concepts or beliefs.² If this view is correct, then a change in a single concept or belief will have far-reaching effects on other concepts. This conclusion has been thought by many philosophers to have extremely counter-intuitive consequences with respect to the possibility of communication, intrapersonal concept endurance, and the justification of our inferential behavior. In what follows I want to indicate that these fears are misplaced: the holist has plausible responses to offer. I shall conclude that worries relating to communication, intrapersonal endurance, and inference justification are not the best reasons for rejecting meaning holism.

*Department of Philosophy, Rutgers University, joncohen@ruccs.rutgers.edu

¹I call these motivations, and not arguments, because I think the real arguments against holism are negative demonstrations that actual holist theories of meaning are unsuccessful, not *a priori* arguments that no possible holist theory could be successful. Since their views are so tendentious, holists end up shouldering the burden of proof; hence, atomists spend most of their time trying to refute the positive arguments for holism. (This kind of an anti-holist project is what's going on in [Fodor and Lepore, 1992].) For this reason, atomist challenges to holism don't take such a central role in such discussions.

²One finds both formulations in the literature. Because beliefs are beliefs about *xs*, it looks as though the two formulations will come to the same thing on most versions of holism, although a little work might be needed to tidy up the details. In what follows, I'll switch freely between the two formulations.

1 Communication

We have said that, for a holist, the meaning of a single concept C depends on the entire system of other beliefs and concepts to which it is attached. However, it seems unlikely, on the face of it, that any two individuals ever share all (or even very many) beliefs.³ Hence, it would appear equally unlikely that any two individuals could share a single concept. On this view, therefore, concepts are inherently non-sharable, non-public entities.

Some philosophers have held that this consequence would preclude any understanding of linguistic communication. On this line, which I'll call the atomist's argument concerning linguistic communication, or just the atomist's argument, you and I attach different concepts to the word "aardvark" since our networks of beliefs and concepts are distinct. So when you say "aardvark," intending to express your concept A_1 , I hear you and bring forth my concept A_2 , where $A_1 \neq A_2$. Because we attach different concepts to the word when we talk, we cannot be said truly to be communicating. Surely, though, politicians and telephone marketers notwithstanding, we often communicate successfully by using words of public languages. Thus, if holism really does threaten the notion of linguistic communication, holism must be jettisoned. Or so the standard story is supposed to go.

Unfortunately, I should like to suggest, it is not at all clear that holism does threaten the notion of linguistic communication in the way purveyors of the atomist's argument would have us believe. There are at least two ways in which I think the holist might want to challenge the putative *reductio* by disputing the conception of communication it assumes. I propose to examine and evaluate each.

First, the holist might appeal to a notion of concept similarity to recapture an explanation of communication.⁴ This suggestion is sometimes thought to originate in Frege's claim that "the task of our vernacular languages is essentially fulfilled if people engaged in communication with one another connect the same thought, *or approximately the same thought*, with the same proposition." ([Frege, 1980], 115, emphasis added)⁵ On this response, the holist could admit that you and I associate the word "aardvark" with different concepts, but insist

³One might hope to defuse the present worry by attempting to argue that individuals can share the beliefs which matter for concept determination even if they can't share all beliefs, but it would appear that the price for making this move is the acceptance of a principled analytic/synthetic distinction, a price widely held to be exorbitant. I cannot treat this matter here for reasons of space. See ([Fodor and Lepore, 1992], 22ff), [Fodor and Lepore, 1993], [Boghossian, 1993], and [Block, 1993] for a discussion of the issue.

⁴This answer is ubiquitous in the literature. For particularly strong recent statements, see ([Harman, 1993], 169–170), ([Smith et al., 1984], 268), and ([Churchland, 1989], 102).

⁵Frege's word "proposition" is best understood in this context as something like "sentence", while his term "thought" is what we are calling "content". Of course, Frege was talking about the connection between whole sentences and contents, not between single words (e. g., "aardvark") and concepts; the latter problem, as we are pursuing it in the present essay, would have been rejected by Frege as arising from a violation of the context principle. I mention Frege only to point out an early form of something like the move we are considering, not as a claim that Frege's concerns and ours should be identified.

that this does not prevent us from communicating with that word so long as our distinct concepts are sufficiently similar. If this is right, then a holist account of concept individuation needn't threaten linguistic communication, since our ability to communicate wouldn't depend on the identity of our concepts, but only their similarity.

The difficulty with this response is that it doesn't work. As pointed out in [Fodor and Lepore, 1992] (and also in [Fodor, 1998]), we must take care to distinguish between *sort-of-believing that P* on the one hand, and *believing that sort-of-P* on the other. It should be clear that the notion of similarity needed by the holist is the latter, not the former, and that this notion is (to put it mildly) obscure. In any case, it would appear that the availability of this sort of content similarity is seriously in question for a theorist who repudiates content identity. Roughly, to state what makes our "aardvark" beliefs similar, one needs to claim that there are at least some nodes shared by my AARDVARK₁ web and your AARDVARK₂ web, as it might be. Of course, this simply pushes the problem of giving identity conditions for beliefs back a step, since we now need to say what it means for the two webs to share a node:

The colloquial senses of "similar belief" *presuppose* some way of *counting* beliefs, so they presuppose some notion of belief *identity*. If you have most of the beliefs that I have, then, a fortiori, there are (one or more) beliefs that we both have. And if there is a proposition that you sort of believe and that I sort of believe strongly, then, a fortiori, there is a proposition that is the object of both of our beliefs. But precisely because these colloquial senses of belief similarity *presuppose* a notion of belief identity, they don't allow us to *dispense with* a notion of belief identity *in favor of* a notion of belief similarity. ([Fodor and Lepore, 1992], 18–19, emphasis in original)

Essentially the same points are made in chapter 2 of [Fodor, 1998]. I take it that, pending some unforeseen clarification of the notion of concept similarity (which doesn't depend on a notion of concept identity), this objection is a rather decisive one against the holist reply we have been considering.

Now, this would all be rather bad news for a holist who had only the previous response up her sleeve. But alas, holists tend to have many, and notoriously long, sleeves. I'm about to suggest that one of them contains a more plausible defense against the accusation that meaning holism precludes linguistic communication.

According to this second line of response, the holist can explain linguistic communication by denying, once again, that such communication constrains the relation between the concepts tokened by interlocutors as strictly as the atomist's argument requires. But unlike the last response, this suggestion doesn't just appeal to an unclarified notion of similarity. Instead, according to this proposal, we can communicate if our (distinct) concepts are necessarily co-instantiated, i.e., instantiated in all the same possible and actual situations.⁶ Suppose, for example, that when you utter the word "aardvark", I assign to the

⁶I first considered this idea seriously after reading a related proposal in [Prinz, 1995].

word the concept AARDVARK-TIME-SLICE, while you assign to the word the concept NON-DETACHED-AARDVARK-PART. No doubt these two concepts are very different. On the other hand, we will both use the word “aardvark” in the same situations, our predictions couched in “aardvark” vocabulary will be confirmed (denied) by the same instances, and so forth. If you said, “Lo, an aardvark!”, I’d know just what to expect, and would appeal to the same external evidence that you would to determine whether your observation were correct, all because our concepts are necessarily coinstantiated.⁷ In short, there seems to be a quite robust sense in which we can communicate even if we don’t associate the same concept with a given word. Therefore, if a holist is willing to understand linguistic communication in this (admittedly unconventional) way, there’s no reason for her to worry that her holism prevents her from explaining communication.

A few comments are in order.

The first thing to notice is that this notion of communication does not bring an encroaching behaviorism in its wake, for it demands what behaviorism repudiates, viz., real, mental constructs called concepts. It’s just that the question of whether communication has occurred can’t be answered simply by comparing the concepts in the heads of discussants. To put the present point in other words, accepting realism about mental constructs (such as concepts) does not commit one to holding that communication requires strict sharing of those concepts.⁸

The second thing to notice is that this notion of communication is actually quite close to the externalist semantics typically favored by atomists in at least one sense. Fodor and other atomists hold that interpersonal concept identity is determined by relations to external factors: to a zeroth approximation, you and I share a concept if we each have a causally efficacious mental particular which covaries appropriately with the external objects from the same class. But to say that you and I share a concept does not require that we share internal mechanisms: the internal mechanisms which aligns AARDVARKS in me with aardvarks in the world might be quite distinct from the mechanism which aligns AARDVARKS in you with aardvarks in the world. Notwithstanding these differences in our internal mechanisms, we manage to communicate.

This idea should permit us to see that the holist understanding of communication we have sketched is similar to that propounded by at least one sort of atomist. The externalist atomist holds that my relation to external factors fixes my concepts, while the holist denies this; this makes for the semantic differ-

⁷When I say that two predicates are coinstantiated, I mean that one is instantiated just in case the other is. I don’t mean that they are instantiated by the same individuals. [Ray, 1997] proposes that what I’m calling coinstantiation should be called simul-instantiation to keep this distinction clear. Instead, I’ll just use the term “coinstantiation” consistently in the sense defined above.

Also, notice that you need necessary coinstantiation (and not simply coinstantiation) to get the counterfactuals to work out right.

⁸Of course, holists admit that concepts are never (or almost never) literally shared interpersonally, and therefore hold that concepts lack the sort of publicity demanded of them in, for example, [Fodor, 1998]. What the holist we are imagining denies is just that this lack of publicity has any important consequences for the possibility of linguistic communication.

ence between the two theorists. But the externalist atomist and the holist have something interesting in common. Namely, both hold that communication is ensured when the parties bear a certain relation to all the same external objects and events, although they'll disagree on the question of whether this relation is mediated by the parties' having identical concepts. So in spite of the differences between the two theorists, they can largely agree in their understanding of linguistic communication. No doubt there are many other instructive dimensions along which to compare the two programs, but I'll spare you the details. The important point here is that the notion of communication which enables a holist to respond to the atomist's argument is not only coherent, but also in some ways plausible for the atomist himself to endorse.

This comparison, however, raises a certain worry. If, as we've admitted, the property of *being an aardvark time slice* and that of *being an undetached aardvark part* (and that of *being an aardvark*) are necessarily coinstantiated, and if meaning, for an externalist, is a matter of covariation, you might wonder why an externalist isn't committed to the position that nothing could determinately mean the former and not the latter. Fodor has argued on behalf of externalism (chapter 3 of [Fodor, 1994]) that such Quinean worries are answerable. I can't address the success of these arguments in the present essay, but I do want to forestall one objection. If Fodor is right that even necessarily coinstantiated properties can be shown to have determinately distinct reference in some situations, then one might worry that the sort of communication I am proposing wouldn't work out for the two properties in question: there would be some situations in which the reference relations of the two concepts would come apart. But this doesn't affect my proposal. Fodor was concerned to argue that even if concepts are necessarily coinstantiated, they can still differ in meaning. But he admits that there can be distinct concepts which are necessarily coinstantiated, and this is all that is required by the notion of communication I am imagining.

The moral of this part of the story, then, is that the holist has what appears to be a plausible explanation of linguistic communication which is not threatened by her holism. Thus, if we are trying to find reasons to avoid holism, we had better find one more compelling than the atomist's argument concerning linguistic communication.

2 Intrapersonal Endurance

The threat to meaning holism based on intrapersonal endurance of concepts is roughly analogous to the concern based on communication. Because my beliefs change over time, it is unlikely that I hold the same network of beliefs at two distinct moments. Hence, if my network of beliefs is constitutive of my concept AARDVARK, then it is unlikely that I hold that very concept at two distinct moments. On the other hand, it seems intuitively plausible that we retain concepts over time. Some considerations which support this view are (i) it seems that my beliefs about aardvarks and my ability to recognize them remain relatively stable, and (ii) it apparently makes sense to say that

I augment my AARDVARK-concept, say by learning a new fact, rather than simply acquiring a new, unrelated concept.⁹ Many philosophers have thought that these considerations are incompatible with meaning holism.

Once again, I think there are two general strategies available to the holist in fending off this line of attack. First, the meaning holist might suggest that she can accommodate the relevant intuitions concerning concept endurance by appealing to concept similarity, rather than concept identity. However, I have already argued that this maneuver is unsuccessful because the notion of concept similarity presupposes that of concept identity. For this reason, we should not permit the holist's appeal to concept similarity in the context of the present concern any more than we did when considering threats to holism based on linguistic communication.

However, as before, I think the holist has a second response to this worry which is substantially more plausible. The holist's semantics prevent her from holding that a thinker can entertain the same concept over time. But she should deny that this fact prevents her from accommodating the intuitions marshaled above. As before, a holist can explain the apparent persistence of my aardvark beliefs as a consequence of the necessary coinstantiatedness of certain distinct concepts.

To use our prior example, the distinct concepts AARDVARK-TIME-SLICE and NON-DETACHED-AARDVARK-PART could underlie a capacity to discriminate the members of the same class of physical objects (namely aardvarks) on separate occasions. Similarly, what might appear to me to be a persisting AARDVARK belief can be explained by a holist to be a collection of, as it might be, an AARDVARK-TIME-SLICE belief at time t_1 , a NON-DETACHED-AARDVARK-PART belief at time t_2 , and so on. Because these beliefs would be related to the same sorts of external evidence at the same times, I might well come to think that the two distinct beliefs were really a single, persisting belief. In a related way, the holist can hold that although I think I augment my AARDVARK concept, I am really exchanging an AARDVARK-TIME-SLICE concept for a NON-DETACHED-AARDVARK-PART concept. I am none the wiser, of course, since the two concepts are necessarily coinstantiated; hence there is no possible or actual situation which could sway me from my impression that I have augmented a single, persisting concept.

Such impressions are all that the intuitions concerning the intrapersonal persistence of concepts require; it must merely *seem* to us that our AARDVARK beliefs are relatively stable, and that we can add to our existing concepts by learning new facts. A holist, *qua* holist, denies that these intuitions have any purchase on the structure of our concepts, but she needn't deny their existence. Indeed, since she can explain why we have such intuitions without claiming that our concepts do in fact persist over time, the holist can claim that they do not endanger her holism.

⁹I'll consider separately another intuition which supports the view in section 3.

3 Intrapersonal Inference

In the last two sections I treated concerns about interpersonal and intrapersonal concept identity in roughly the same way. I claimed that the holist could respond to anti-holist demands for interpersonal and intrapersonal concept iteration (which the holist can't supply) by denying that intuition makes these requirements on the structure of our concepts. However, there's one respect in which the holist's inability to supply a robust notion of intrapersonal concept iteration might be thought to be more serious than her failures with interpersonal concept iteration. In this section I want to consider the suggestion that our inferential practice provides another reason for thinking that intrapersonal concept identity must be possible.¹⁰

This inferential worry begins with the observation that, typically, we are willing to infer from the premise p to the conclusion $p \vee q$. Now, a reasonable version of a rational justification for this inference pattern involves the claim that the pattern is truth-preserving. However, the pattern is only truth-preserving if there's no equivocation between the two occurrences of the symbol ' p ': the inference from p_1 to $p_2 \vee q$ is *not* truth-preserving. This suggests that if such inferences are carried out on mental representations whose constituents are concepts (as per the usual computationalist stories), then the rational justification for our inference pattern is not a rational justification unless there is strict intrapersonal concept identity between the occurrences of the concepts implicated in the tokening of p .¹¹ But, as we have seen, the holist claims that there is never (or almost never) strict intrapersonal concept identity over time. Therefore, we might think, the holist is forced to admit that the rational justification offered for our inferences (which take place over measurable intervals of time) is not, in fact, a rational justification.

I want to suggest that the holist can respond to this worry in much the same way that we have imagined she might respond to the concerns considered in earlier sections, but that here the move begins to lose some of whatever plausibility it might have had before. We had suggested above that the holist could answer worries about concept iteration by appealing to a relation of necessary coinstantiation between two concepts, which, on her way of counting, are distinct. The same move can be made in the case of the present worry about justifying inferences. Let the concept p_1 be an AARDVARK-TIME-SLICE concept, and let the concept p_2 be a NON-DETACHED-AARDVARK-PART concept; as before, these concepts are distinct, given a holist's method of individuating concepts, but they will be necessarily coinstantiated (or so we have been assuming). Suppose q is a CARBURETOR concept. Then the inference from p_1 to $p_2 \vee q$ is truth-preserving, because the consequent is true (i. e., instantiated) whenever the antecedent is true (i. e., instantiated). Therefore, a holist can give a rational

¹⁰This concern was pointed out to me by Jerry Fodor.

¹¹In particular, and as I shall be assuming for the sake of simplicity in exposition, if p is a primitive concept, the concepts implicated in the tokening of p collapse onto just p itself. The details required for the non-primitive case are messy but not particularly instructive to work out.

justification for the inference in terms of truth-preservation, which is just what the anti-holist claimed was impossible.

However, this move comes at a cost. Although it does enable the holist to give the same sort of rational justification for certain psychological inferential patterns as that offered by the non-holist, it requires that such justifications make use of further premises. In particular, the inference which the holist understands as an inference from p_1 to $p_2 \vee q$ only goes through because of an appeal to the fact that p_1 and p_2 are necessarily coinstantiated. It is important to bear in mind that this is not an isolated case; *every* inference whose justification appeals to truth-preservation will have to trade on an enthymemic metaphysical premise linking what the holist insists are distinct concepts, on pain of an equivocation which would make the inference non-truth-preserving. In contrast, the non-holist can explain the truth-preservingness of such inferences simply by pointing out that they have syntactic forms which we all learned are truth-preserving when we learned first-order logic. Viewed in these terms, I think it's fair to say that the explanation I have mounted on behalf of the holist looks substantially more tortured than that of the non-holist. Thus, even if I am right that the holist strategy I have imagined can be made to work, I think we would still have good reason to prefer the non-holist alternative.

4 Closing Remarks

I have argued that holists have responses to the common worries that their holism is incompatible with intuitions concerning linguistic communication, intrapersonal concept endurance, and justification of our inferences. However, several problems remain.

First, it's useful to notice that the responses I have given on behalf of the holist depend on appealing to relations between the extensions of concepts. And concept extensions are relatively punctate things (this is true even for theorists who think that meanings are holistic; see, for example, [Block, 1986]). As was noticed as early as [Lepore and Loewer, 1987], the kind of holist semantic theories we have been considering really involve two different semantic sub-theories (in this case, one for reference and another for concept individuation), only one of which is holistic. It's worth observing that the holist solutions we have examined gather whatever success they can only by ignoring the holistic aspects which caused the troubles in the first place. So, for example, we answered threats posed by a holistic theory of concept individuation to a notion of communication by pointing to relations between the extensions of AARDVARK-TIME-SLICE concepts and NON-DETACHED-AARDVARK-PART concepts. Thus, if this proposal can be made to work at all, it succeeds only by abstracting away from the holistic aspect of concepts in favor of their non-holistic referential factor. One might wonder, if the holistic factor is causing such a host of problems, and if it must be ignored in order to solve those problems, why we shouldn't just abandon it entirely.¹²

¹²A further, and familiar objection to such two factor theories involves the question of why

Second, it should be obvious that the alternatives I have sketched should give only cold comfort to a dyed in the wool holist, because I've only given her slightly more than an atomist has. The atomist says that interpersonal communication, intrapersonal concept endurance, and inference justification require distinct tokens of the very same concept type. I've allowed that the holist can do the work that needs to be done by requiring only distinct tokens of concepts which are necessarily coinstantiated. Whether this can be worked out successfully or not, we should notice that the relationship I'm imagining, while not as strict as identity, is still a pretty special one for two concepts to have: it's not a constraint which will be met by most pairs of concepts, construed holistically or otherwise. So I'm not at all sure that a holist will find much to get excited about in what I've offered her.

Third, it is highly unclear whether the defense I have mounted is sufficient in a more general setting. Holism is, of course, a theory about the individuation of all (primitive) concepts. All three of the responses I mounted on behalf of the holist worked on the assumption that for any concept C we could find a second concept C' which a holist would count as distinct from C , and so that C and C' would be necessarily coinstantiated. Whatever one thinks of the classic Quinean cases, this assumption is certainly tendentious as a general claim about all primitive concepts. That is to say, the holism which emerges from the responses I have sketched may be ultimately implausible. Although this may be so, I hope I have shown that, at the very least, the considerations against holism I have examined are less clear cut than they might first appear.

Fourth, there are other reasons for thinking holism conflicts with our intuitions which are much stronger. One such reason is the apparent encapsulation of our concepts: it seems that some concepts and beliefs are irrelevant to others. For example, it's hard to see why learning that aardvarks are nocturnal should affect the identity of my CARBURETOR concept. Because such encapsulation is precisely what holism denies, holism clashes in a more or less direct way with this intuition. Most significantly, as pointed out in [Fodor and Lepore, 1992], meaning holism both presents a *prima facie* threat to an attractive metaphysical realism in the philosophy of science and appears to jeopardize the possibility of a scientific intentional psychology. At any rate, this cluster strikes me as a more obvious place where holism and intuition part ways, and therefore as a more important set of considerations in motivating doubts about the plausibility of holism.

References

- [Block, 1986] Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10. Reprinted in [Stich and Warfield, 1994], 81–141.

the two separate factors should be expected to hang together. These and related matters are discussed at length in chapter 6 of [Fodor and Lepore, 1992].

- [Block, 1993] Block, N. (1993). Holism, hyper-analyticity and hyper-compositionality. *Mind and Language*, 8(1):1–26.
- [Boghossian, 1993] Boghossian, P. A. (1993). Does an inferential role semantics rest upon a mistake? *Mind and Language*, 8(1):27–40.
- [Churchland, 1989] Churchland, P. M. (1989). *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. MIT Press, Cambridge, Massachusetts.
- [Fodor, 1994] Fodor, J. A. (1994). *The Elm and the Expert: Mentalese and its Semantics*. MIT Press, Cambridge, Massachusetts.
- [Fodor, 1998] Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press, New York.
- [Fodor and Lepore, 1992] Fodor, J. A. and Lepore, E. (1992). *Holism: A Shopper’s Guide*. Blackwell Publishers, Cambridge, Massachusetts.
- [Fodor and Lepore, 1993] Fodor, J. A. and Lepore, E. (1993). Reply to Block and Boghossian. *Mind and Language*, 8(1):41–48.
- [Frege, 1980] Frege, G. (1980). Letter to Peano, 29 September 1896. In Gabriel, G., Hermes, H., Kambartel, F., Thiel, C., and Veraart, A., editors, *Gottlob Frege: Philosophical and Mathematical Correspondence*, pages 112–118. University of Chicago Press, Chicago.
- [Harman, 1993] Harman, G. (1993). Meaning holism defended. In Fodor, J. A. and Lepore, E., editors, *Holism: A Consumer’s Update*, volume 46, pages 163–171. Rodopi, Amsterdam.
- [Lepore and Loewer, 1987] Lepore, E. and Loewer, B. (1987). Dual aspect semantics. In Lepore, E., editor, *New Directions in Semantics*, pages 83–112. Academic Press, London.
- [Prinz, 1995] Prinz, J. J. (1995). How to defang a dogma without removing its teeth. Ms., University of Chicago.
- [Ray, 1997] Ray, G. (1997). Fodor and the inscrutability problem. *Mind and Language*, 12(3/4):475–489.
- [Smith et al., 1984] Smith, E. E., Medin, D. L., and Rips, L. J. (1984). A psychological approach to concepts: Comments on Rey’s “Concepts and Stereotypes”. *Cognition*, 17:265–274.
- [Stich and Warfield, 1994] Stich, S. P. and Warfield, T. A. (1994). *Mental Representation: A Reader*. Blackwell, Oxford.