

*Psychosemantic analyticity**

RICHARD HORSEY

(September 2001)

Abstract

It is widely agreed that the content of a logical concept such as AND is constituted by the inferences it enters into. I argue that it is impossible to draw a principled distinction between logical and non-logical concepts, and hence that the content of non-logical concepts can also be constituted by certain of their inferential relations. The traditional problem with such a view has been that, given Quine's arguments against the analytic-synthetic distinction, there does not seem to be any way to distinguish between those inferences that are content constitutive and those that are not. I propose that such a distinction can be drawn by appealing to a notion of 'psychosemantic analyticity'. This approach is immune to Quine's arguments, since psychosemantic analyticity is a psychological property, and it is thus an empirical question which inferences have this property.

I

It seems that certain inferences constitute (at least in part) the content of their constituent concepts. Consider, for example, the concept KILL. For Smith to kill Jones, then Jones must die. And it seems that this inference from 'x KILL y' to 'y DIE' is part of the content of the concept KILL. So, for example, we would not attribute the concept KILL to someone if we were not also prepared to attribute the concept DIE to them. Similarly, if someone were not disposed to accept the inference from 'x KILL y' to 'y DIE', we would not be inclined to attribute the concept KILL to that person.

There are some well-known problems with this view. If we allow that certain inferences are content constitutive, then we clearly have to say *which ones*. The only other option would be to allow that *all* the inferences that a concept enters into are constitutive of its content, and many people are uncomfortable with the holistic implications of such a position. The question is, then, how to distinguish those inferences that are content constitutive from those that are not.

* I would like to thank Deirdre Wilson for several interesting discussions on the substance of this paper, and Jerry Fodor for some very useful comments on an earlier draft.

Traditionally, this question has been answered by reference to some notion of analyticity—the claim being that all and only the analytic inferences are content constitutive. The problem is that most philosophers follow Quine in rejecting the possibility of drawing a principled analytic/synthetic distinction. Without such a distinction there would seem to be no way to distinguish content constitutive inferences from the rest.¹

These arguments suggest that perhaps there aren't any content constitutive inferences after all.

On reflection, however, it doesn't seem possible to give up *all* content-constitutive inferences. First, we still have to account for the strong intuitions we have about cases such as KILL. And there is a more serious problem. How are we going to account for the meaning of logical constants such as 'and'? Almost everyone (including atomists such as Fodor) agrees that we need inference rules to account for the meaning of the logical constants—for it would seem that the meaning of words such as 'and' *just is* their inferential role.

Fodor doesn't regard this as too much of a problem (see Fodor 1994, chapter 3; 1987, p. 78). He accepts that inference rules are necessary to account for the meaning of the 'logico-syntactic' apparatus, but he claims that the same is not true of the non-logical vocabulary. In other words, he considers that a clear distinction can be drawn between the logical and non-logical vocabularies ("there is no reason at all to suppose that the logico-syntactic vocabulary is itself interdefined with the *non-logical* vocabulary" [Fodor 1994: 76, original emphasis]). His position is thus that there are no content constitutive inferences *for the non-logical words*.

This immediately raises the question of how the logical vocabulary is to be specified. Fodor is not explicit about *which* words he takes to make up the logical vocabulary, but he seems to have in mind the standard propositional connectives ('and', 'or', 'not', and so on). This raises a number of questions. Would he want to claim that the quantifiers of predicate calculus are part of the non-logical vocabulary? Would the same be true of non-standard quantifiers such as 'many' or 'most'? How about modals, tense and aspect, for which formal logics have been developed? Other words also support entailments: how are we to account for the entailment relation between 'either'/'both' and 'two'? This raises the question of

¹ As Fodor & Lepore (1991) discuss, there is in any case a real danger of circularity with the proposal that all and only analytic inferences are content constitutive: as Quine (1953) made clear, one cannot fall back on analyticity to distinguish those inferences that are constitutive of the meaning of a concept, since analyticity is itself explicated in terms of meaning. Rather, one would have to employ a notion of analyticity explicated in terms of some non-semantic property. See also Boghossian (1993) for discussion of this point.

whether (some) numerals are included in the logical vocabulary. Fodor gives us no idea how to decide these questions.²

In any case, there is no reason to think that a neat line *can* be drawn between the logical and non-logical vocabularies. This makes Fodor's position untenable. For if a neat line cannot be drawn between logical and non-logical words, and if the content of logical words is constituted by the inferences they enter into, then it would no longer be possible to rule out content constitutive inferences for the non-logical words either.

There are a number of arguments which suggest that the logical/non-logical distinction that Fodor wants to draw cannot in fact be drawn (some of these arguments were made by Fodor himself). Fodor et al. (1980) note that logic is commonly thought of as providing a reconstruction of our intuitions about the validity of arguments such as that in (1), but not of arguments such as that in (2).

- (1) John left and Mary wept, therefore Mary wept
 (2) John is a bachelor, therefore John is unmarried
 (Fodor et al. 1980: 269)

The traditional view in logic has been that this is because argument (1) is valid in virtue of the meaning of 'and', which is part of the logical vocabulary, whereas argument (2) is valid in virtue of the meaning of 'bachelor', which is part of the non-logical vocabulary. But why, given that we have the intuition that both (1) and (2) are valid arguments, should we draw this kind of distinction?

One problem is that there are all sorts of words which appear to have properties of both the logical and non-logical vocabularies. Consider the example of 'inside' (cf. Sperber & Wilson 1995: 105).

- (3) a. x is inside y
 b. y is inside z

 (4) x is inside z

(3)–(4) is an example of a valid inference. The relation 'is inside' has certain formal properties: it is *transitive* (as just exemplified), *irreflexive* (nothing is inside itself), and *asymmetric* (if x is inside y , then it is never the case that y is also inside

² Fodor has recently suggested (personal communication) that no logical term can appear as a constituent of a closed, atomic sentence. This proposal is also somewhat unclear. Can quantifiers, for example, occur in closed atomic sentences? How about modals and tense? In the end it comes down to a stipulation of what counts as a logical term. I argue below that it is unlikely any such stipulation could be plausible.

x). These properties may be captured by a set of inference rules in the same way that the properties of a logical connective such as ‘or’ are captured. And yet it is patently not the case that ‘inside’ is a part of the logical vocabulary; at least part of its meaning derives from its other content (INSIDE is a full-fledged concept, with associated encyclopaedic information).

The point about the logical vocabulary is that its content *just is* its inferential role. But the content of ‘inside’ *cannot* just be its inferential role; if it were, then it would be synonymous with ‘below’, which has identical logical properties (like ‘inside’, ‘below’ is transitive, irreflexive, and asymmetric).³

There are even examples of words which have referential/descriptive content, but which also exhibit the properties of one or other of the standard logical constants. For example, adjectives such as ‘raw’ and ‘dark’ share the logical properties of the operator ‘not’, since ‘raw’ entails ‘NOT cooked’ and ‘dark’ entails ‘NOT light’. Connectives such as ‘since’ and ‘because’ share the logical properties of ‘and’, since ‘*x* since/because *y*’ entails ‘*x* AND *y*’. The words ‘sibling’ and ‘unless’ share the logical properties of ‘or’, since ‘sibling’ entails ‘brother OR sister’ and ‘*x* unless *y*’ entails ‘*x* OR *y*’. In none of these cases do these logical properties exhaust the meaning of the word, so it would not be possible to claim that these are logical words. Presumably, however, their logical properties must be accounted for in the same way as the corresponding logical constant—via inference rules. The possibility that some of the inferences entered into by non-logical words are content constitutive cannot, it seems, be ruled out.

A similar point can be made with regard to non-truth-conditional connectives such as ‘but’ and ‘although’. Are these logical or non-logical words? Most people agree that ‘but’ is logically equivalent to ‘and’, so that it would be governed by the same inference rules. Clearly ‘and’ and ‘but’ do not have the same meaning, however—in addition to its truth-conditional meaning, ‘but’ introduces some non-truth-conditional notion of ‘contrast’. Here is another case where the logical and non-logical vocabularies appear to be interdefined.

Even if it were possible to draw a neat distinction between logical and non-logical words, this might not be of much help. As Boghossian (1993: 354) points out, Quine’s arguments hold just as much of the logical vocabulary as the non-logical vocabulary. If we claim that the meaning of logical words is constituted by the inferences they enter into, we can still ask which inferences those are. How do we decide between an obviously true but non-defining sentence, and a sentence

³ Conversely, it cannot be the case that the logical properties of ‘inside’ are general to spatial prepositions, and hence do not need to be specified for each individual preposition, since not all spatial prepositions have the same logical properties (‘beside’, for example, is non-transitive, irreflexive, and symmetric). It would thus seem that information about their logical properties has to be stored with the individual concepts in these cases.

which implicitly defines the meaning of some logical term? Quine argues that we cannot make such a decision (see, for example, Quine 1935, 1954).⁴ So if (like Fodor) you are inclined to regard Quine's arguments against the analytic/synthetic distinction as ruling out content constitutive inferences, you should be wary of proposing an inferential account for the content of the logical constants.

We are in something of a dilemma at this point. On the one hand, we need to hang on to content constitutive inferences for some non-logical words (and certainly for the logical words). But on the other hand there appear to be good Quinean reasons for rejecting the notion of content constitutive inference. Before considering a possible way to resolve this dilemma, we need to look at the issues surrounding content constitutive inferences in more detail.

II

There is a venerable philosophical tradition which gives an account of a priori knowledge in terms of content constitutive inference (or implicit definition). According to this account, since certain inferences entered into by a concept are constitutive of its content, then these inferences must be valid (i.e. analytic).⁵ Moreover, since these inferences are analytic, we must be justified in asserting their validity, and this explains their apriority. Or alternatively: since certain statements (meaning postulates) define the meaning of certain of their constituent terms, then these postulates must be true, hence analytic. And since they are analytic, we must be justified in holding them true.⁶

Those who want to deny the implicit definition account of a priori knowledge therefore have two strategies available to them. They can either deny the relation between implicit definition (/content constitutive inference) and analyticity, or they can deny the analytic explanation of a priori knowledge.

Horwich (1992) adopts this latter strategy. He considers the possibility that the language faculty may contain certain meaning postulates (such as 'bachelor \rightarrow unmarried man' or ' x caused $y \rightarrow x$ preceded y ') which are "transmitted to that area of the brain in which beliefs are stored" (1992: 100). Horwich suggests that it is

⁴ Boghossian (personal communication) suggests $p \vee \neg p$ as a possible example (although he notes that since we lack a convincing theory that tells us how to pick out the defining sentences, any particular example will be controversial). Many people consider that $p \vee \neg p$ may well be false. Such examples underline the fact that it is difficult to say exactly which sentences contribute to implicit definition, even if their truth seems obvious.

⁵ Block (1993) terms this the 'Plausible-Sounding Principle' (and argues against it).

⁶ See also Horwich (1998: chapter 6) for discussion of these issues.

then a simple matter to characterise a notion of analyticity (what he terms ‘I-analyticity’) which is completely determinate in its application,⁷ and so immune from Quine’s arguments. By invoking a determinate notion of analyticity, Horwich is able to keep meaning postulates without having to associate the content of a concept with its entire inferential role. Content is associated with those inferences which are I-analytic, and this avoids the problems of holism which Fodor raised. The connection between I-analyticity and a priori knowledge is severed, because Horwich proposes that meaning postulates could arise or be revised in light of experience.

There are some problems, however, with Horwich’s account. First, the language faculty doesn’t seem to be the obvious place to locate these facts. It is certainly not a fact about the *English words* ‘cause’ and ‘effect’ that causes precede their effects. Rather, it is a nomological necessity, a fact about *causing* and *effecting*. The question is then *where* and *how* this fact is represented in the brain. It seems we must say that this is not a linguistic fact but a conceptual fact. Similarly, it is not a fact about the English word ‘kill’ (or about corresponding words in other languages) that killing someone entails that they die. It is a fact about *killing*, captured by properties of our concepts KILL and DIE.

In contrast with Horwich (1992), Boghossian (1993, 1994, 1997) adopts the first strategy, denying the relation between implicit definition and analyticity. He points out that for an inference to be content constitutive, it need not be analytic.⁸ An inference may be content constitutive for a subject provided the inference is *regarded as valid* by that subject. But regarding an inference as valid in no way requires that the inference actually *is* valid. *A fortiori*, it does not require that the inference is valid in virtue of the meanings of its constituent terms. We can imagine, for example, a pre-scientific community for whom the concept WHALE has its content constituted (in part) by the inference ‘WHALE \rightarrow FISH’,⁹ which members of this community take to be valid. This inference would be content constitutive for members of this community, even though the inference is not in fact valid (and so not analytic).

⁷ The characterisation given by Horwich is that “a sentence is analytic in a person’s I-language at time *t* if and only if it is a consequence of the principles of that person’s language faculty at time *t* that this sentence be taken as true regardless of evidence” (1992: 101).

⁸ Horwich (1998: 143) also discusses this point. In fact, Horwich (1998) seems to deny *both* the analytic explanation of a priori knowledge and the connection between content constitutive inference and analyticity.

⁹ Such inferences are more appropriately expressed in the form ‘*x* WHALE *y* \rightarrow *x* FISH OF A CERTAIN KIND *y*’. See Sperber & Wilson (1995: §2.4).

It is at least possible, then, to dissociate the notion ‘content constitutive inference’ from the notion ‘analytic inference’. But it is still necessary to say on what basis an inference is to be regarded as content constitutive. What is required is some means of cashing out ‘content constitutive inference’ which avoids Quinean problems.¹⁰ Here is a possibility.

III

Suppose that we follow Fodor (1975) in seeing content constitutive inferences as mentally-represented ‘meaning postulates’ (understood as inference rules rather than axioms). We will then have a uniform account of the logical properties of the so-called “logical” and “non-logical” words. The only difference between the two would be that the content of a logical concept such as AND is exhausted by its associated inference rules, whereas the inference rules associated with a non-logical concept such as DARK do not exhaust its content.¹¹ In both cases the inference rules constitute, at least in part, the content of the concept.

Adopting this approach leads us, like Boghossian, to deny the relation between content constitutive inference and analyticity. The reason is that content constitutive inferences are now to be conceived of psychologically—and while the majority of our psychologically represented inference rules are no doubt veridical, this is by no means necessary. What the approach we are proposing amounts to is this. We are specifying some property other than analyticity that can be used to determine *which* inferences are content constitutive, and thereby avoiding Quine’s objections. Let us call this property ‘psychosemantic analyticity’. Then an inference will be constitutive of the content of a concept just in case that inference is psychosemantically analytic—that is, just in case the inference is associated (as a meaning postulate) with the concept in question.

Whether or not an inference is psychosemantically analytic is an empirical question. In order to determine whether a given inference is psychosemantically analytic, hence content constitutive, we have to ask whether this inference is governed by an inference rule attached to a particular concept—that is, whether the inference is part of our mental logic. This can be discovered using standard

¹⁰ Boghossian (1994: 120) notes that Quine regarded obviousness as one possible basis in the case of logical principles (see Quine 1954). However, it seems that content constitutive inferences do not necessarily reduce to ‘obvious’ inferences (see also Horwich 1992, p. 104 on this point).

¹¹ Note that, *pace* Horwich, there is no suggestion that these meaning postulates inhabit the language faculty.

techniques of empirical investigation (inference rules will be fast, they will apply mandatorily, and so on).

There are a number of constraints on psychosemantic analyticity. If we possess a particular inference rule, then we must believe that the rule expresses a (nomological) necessity. This rules out the possibility that certain inferences (for example, ‘GERBIL \rightarrow PET’, which we would presumably not take to be nomologically necessary) are psychosemantically analytic. But clearly not all necessary truths are going to be represented as inference rules. We must surely represent most necessities in the form of propositional encyclopaedic knowledge, and since these are not inference rules, they will not be psychosemantically analytic. So how do we know whether a (nomologically) necessary truth is stored as an inference rule or as a proposition? Sperber & Wilson (1995: Chapter 2) suggest that only those inference rules which derive conclusions relevant to an organism (in the sense of improving its representation of the world) will be mentally represented. Using this criterion, Sperber & Wilson propose that the only inference rules attached to concepts are elimination rules—that is, there are no ‘introduction rules’ (e.g. \vee -introduction), since they will not derive conclusions of any relevance to the organism (an organism which believes that p will not improve its conceptual representation of the world by deriving from this $p \vee q$).

In summary, it is clear that the content of logical concepts is constituted by inference rules. It is also clear that no principled distinction can be made between logical concepts and the rest. It follows that the content of “non-logical” concepts may also be constituted, at least in part, by inference rules. We distinguish between those inferences that are content constitutive and those that are not on the basis of a property we are calling psychosemantic analyticity. This is a psychological property, and it is thus an empirical question which inferences have this property. A number of constraints on what could constitute a psychosemantic analyticity have been proposed.

References

- Block, N. (1993) Holism, Hyper-analyticity and hyper-compositionality. *Mind and Language*, 8: 1–26.
- Boghossian, P. A. (1993) Does an inferential role semantics rest upon a mistake? *Mind and Language*, 8: 27–40.
- Boghossian, P. A. (1994) Inferential role semantics and the analytic/synthetic distinction. *Philosophical Studies*, 73: 109–122.
- Boghossian, P. A. (1997) Analyticity. In B. Hale & C. Wright (eds) *A Companion to the Philosophy of Language*, 331–368. Oxford: Blackwell.
- Fodor, J. A. (1975) *The Language of Thought*. New York: Crowell.
- Fodor, J. A. (1987) *Psychosemantics*. Cambridge, Mass.: MIT Press.

- Fodor, J. A. (1994) *The Elm and the Expert: Mentalese and its Semantics*. Cambridge, Mass.: MIT Press.
- Fodor, J. A., M. F. Garrett, E. C. T. Walker & C. H. Parkes (1980) Against definitions. *Cognition*, 8: 263–367.
- Fodor, J. A. & E. Lepore (1991) Why meaning (probably) isn't conceptual role. *Mind and Language*, 6: 328–343.
- Horwich, P. (1992) Chomsky versus Quine on the analytic–synthetic distinction. *Proceedings of the Aristotelian Society*, 92: 95–108.
- Horwich, P. (1998) *Meaning*. Oxford: Oxford University Press.
- Quine, W. V. O. (1935/1966) Truth by convention. In Quine 1966: 77–106.
- Quine, W. V. O. (1953) Two dogmas of empiricism. In *From a Logical Point of View*, 20–46. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. O. (1954/1966) Carnap and logical truth. In Quine 1966: 107–132.
- Quine, W. V. O. (1966) *The Ways of Paradox*. Cambridge, Mass.: Harvard University Press.
- Sperber, D. & D. Wilson (1995) *Relevance*, 2nd edition. Oxford: Blackwell.