

*The Content of Logical Concepts**

RICHARD HORSEY

Abstract

Informational semanticists such as Jerry Fodor argue that the content of lexical concepts is constituted by nomological links between the concepts and the corresponding properties. While such an approach can account for the content of concepts such as CAT, BACHELOR, and maybe even DOORKNOB, it cannot account for the content of logical concepts: however much of an ontological liberal one may be, it is not plausible that the concept AND (say) gets its content from a nomological link with an abstract property of *conjunction*. The usual approach taken by informational semanticists has been to accept instead an inferential-role account of the logical terms. Recently, however, Fodor (2004a) has rejected this inferential-role view of the content of such terms, leaving open the question of how their content is constituted.

In this paper, I develop a psychologically-motivated account of the content of the logical connectives based on elimination rules. The analysis suggests that there is a degree of principled underspecification in the content of such connectives which corresponds closely to the range of interpretations found in natural language. It is shown that additional specification can be provided by pragmatic narrowing processes.

This account provides support to Sperber & Wilson's (1995) proposal that (deductive) mental inference makes use exclusively of elimination rules. The analysis also partly supports Fodor's rejection of the inferential-role account of logical content. At the same time, however, I argue that the elimination rules attached to logical concepts have to be seen as content constitutive, and therefore that Fodor is wrong to reject all content-constitutive inferences.

1 What is conceptual content?

It is a curious fact about the world that some things (sentences, thoughts, propositions) are *about* other things (tables, chickens, even other sentences/thoughts/propositions). Thus the sentence 'Boris the cat is black' is about Boris the cat; and the thought 'chickens are tasty' is about chickens; and the proposition $\lceil \text{John}'(x) \rightarrow \text{greedy}'(x) \rceil$ is about John. This is what is known as 'intentionality', and it wants explaining, since it is not immediately obvious *how* something can be *about* something else. It's not, for example, likely that intentionality is a fundamental property of things in the same way as mass or charge or spin.

* I would like to thank Deirdre Wilson for discussion of these issues and for her comments on an earlier draft.

There is clearly a link between intentionality and representation. The sentence ‘Boris the cat is black’ conveys a certain piece of information (represents a certain state of affairs), namely that Boris the cat is black. This same information could, it seems, also be conveyed in other ways. For example, a photograph or a drawing of Boris the cat could convey the same information (it too represents a certain state of affairs). There are, however, important differences between these two modes of representation. A fundamental difference is that while the photograph or drawing represents iconically (that is, it resembles what it represents),¹ the sentence (/thought/proposition) represents in an indirect or coded way (that is, by means of abstract meaning relations). This fundamental difference gives rise to a number of specific differences between these modes of representation.

One significant difference concerns the amount of information that is conveyed. A sentence, thought or proposition conveys highly specific information (for example, the information that the cat is black). A photograph, on the other hand, cannot convey one piece of information (the colour of the cat, say) without also simultaneously conveying an indefinite number of other pieces of information (the cat’s size, shape, orientation, and so on). Dretske (1981: 137) calls this a distinction between ‘digital’ and ‘analogue’ forms of representation:

...a signal (structure, event, state) carries the information that *s* is *F* in *digital* form if and only if the signal carries no additional information about *s*.... When a signal carries the information that *s* is *F* in analogue form, the signal always carries more specific, more determinate, information about *s* than that it is *F*.

This is important. Intentionality is the property that something has of being about something else. Sentences, thoughts and propositions have intentionality (so the thought that Boris is black is a thought *about* Boris). Photographs do not have intentionality in the same way (while a photograph of Boris might be about Boris, it

¹ Actually, things are a little more complex. Although photographs are iconic representations, in the sense that they resemble what they represent, they do not necessarily represent what they resemble. That is, their content is fixed not by resemblance but by their causal properties—a photograph of Boris is not a photograph that resembles Boris, however closely, but a photograph caused by Boris (so it needs to have been Boris that was sitting in front of the camera when the photograph was taken, and not, say, Boris’s identical twin). Compare this with drawings. Drawings are also iconic representations, but again their content is fixed not by resemblance but (postmodernist art critics notwithstanding) by reference to the intentions of the creator: a sketch intended to be of Boris is a representation of Boris, and not of Boris’s twin, *even if Boris’s twin posed for the artist* (to help get the shape of the face right, say). The sketch may not even resemble Boris (perhaps the artist is intentionally abstract or just not very skilled). The classic text on the different types of representation is Peirce (1931–1935), although the specific proposals that it makes have been rejected by many later theorists. See Goodman (1976).

might also be about many other things: cats in general, quadrupeds, fur, a table which happens to have a cat sitting on it, a room, and so on).

Another important difference concerns the possibility for conveying false information (misrepresentation). A photograph's meaning is tied to its information content, which it records through a deterministic process of photons striking light-sensitive chemicals, and so on. A photograph can only carry the information that p if the state of affairs p in fact pertains (/pertained). So a photograph of Boris the cat can only carry the information that Boris is black if in fact Boris actually is (/was) black. If the photograph is deliberately manipulated so that it depicts Boris as being white, it does not carry the information that Boris is white, since there is no such information for it to carry, and the photograph therefore does not mean that Boris is white. Thoughts and other intentional structures are not like this. The thought that Boris the cat is white can be false or true, depending on how the world is constructed, and the thought represents this state of affairs whether or not it happens to be true. So while representation and intentionality are related notions, they pull apart in important respects. (For detailed discussion of these issues, see Dretske 1981, 1986, Grice 1957, and Woodfield 1986.)

Let us now consider the more specific case of the propositional attitudes. As the name suggests, propositional attitudes are attitudes that we can take towards propositions. Examples of propositional attitudes are believing that a cup of coffee is on the desk before me, desiring to drink some coffee, intending to raise the cup to my lips, and so on. In each case, there is an attitude (belief, desire, intention) towards a proposition (expressed by a 'that...' or 'to...' clause). The proposition expresses what is known as the content of the attitude. A common way of talking about propositional attitudes, following Stephen Schiffer, is in terms of a series of 'boxes' corresponding to the various attitudes. On this analogy, the event of my believing that p comes about as a result of having in my belief box a representation that means p (the same representation appearing in my desire box would count as an event of my desiring that p , and so on *mutatis mutandis*). The analogy is rather direct: the content of the attitude corresponds to the contents of the box in question.²

I will assume that the representations expressing the contents of the attitudes are structured in broadly the same way as the natural-language clauses used to express them (this is the language of thought hypothesis). That is, mental representations of clauses are built up from lexical concepts, combined in accordance with the

² Note that propositions are abstract entities, so it's a representation (a token of a symbol that expresses the proposition) which appears in the attitude-box, not a proposition. See Fodor (1978). Attitude boxes are of course just a convenient metaphor. The idea behind this metaphor is that attitudes are typed with respect to their functional properties (for example, desires tend to cause action designed to bring about the conditions for their satisfaction). Importantly, this proposal is intended to be neutral on the question of how the *content* of the attitudes is to be determined.

syntactic principles of the language of thought (in the same way that natural-language clauses are built up from lexical items combined in accordance with the syntax of the natural language in question). It follows that the content of a clausal representation is built up from the contents of its constituent concepts, together with the syntax. (And in fact, nothing in what follows will depend on there being a language of thought; all we need is the much less contentious assumption that thought is compositional.) The content of propositional attitudes is therefore derivative; underived content is to be located at the level of lexical concepts.

All intentional states have the property of semantic normativity: they represent things as being a certain way, and as we have seen above, they do so independently of whether things actually are that way. Intentional states therefore have truth conditions (or, more generally, satisfaction conditions)—sets of conditions that determine whether a particular belief is true, or whether a particular desire is satisfied. Thus, for example, the belief ‘Boris the cat is black’ is true iff the world is arranged in such a way that the individual cat in question (Boris) has the particular property in question (blackness). Similarly, my desire to drink coffee is satisfied iff a certain state of affairs is brought about (*viz.*, the state of affairs expressed by the proposition in question: I drink some coffee).

Propositional attitudes have two important properties, then. They have propositional (conceptual) content, and they have satisfaction conditions. The question arises how these two properties are related. It is first worth noting that whatever propositional content is, it is something abstract. Propositional attitudes are instantiated because we are related in a certain way to a proposition, by having a mental representation expressing that proposition in our belief box, say. A token mental representation expresses an abstract proposition, so the content of that proposition is presumably itself something abstract. Next, notice that the satisfaction conditions of a propositional attitude reduce to the truth conditions of its propositional object: a belief is true if its propositional object is true, a desire is satisfied by making its propositional object true, similarly, an intention is realised when the truth of its propositional object is brought about. Truth conditions are also abstract. They are a characterization of how the world would have to be in order for the proposition to be true.

So, propositional attitudes have content, which is something abstract, and their propositional objects have truth conditions, which are also abstract. One obvious conclusion to draw from this is that propositional contents just are truth conditions. The content of a concept would then be seen as the contribution that concept makes to the truth conditions of propositions within which it occurs. And, indeed, this is a fairly widely held assumption. One possible objection to proposing that propositional contents and truth conditions are identical, however, is that this rules out the possibility that different propositional contents have the same truth conditions. If we wish to leave this possibility open, then, we would need to find

some abstract characterization of content that was more fine-grained than truth conditions—that is, some function from propositions to truth conditions. (For discussion of these points, see Stalnaker 1998.)

Fodor (1998) has a somewhat different solution. He holds that the contents of the propositional attitudes are truth conditions, but also accepts that two propositions can have the same content, and therefore the same truth conditions, and yet be *different* propositions. He does this by adapting the Fregean proposal that concepts are individuated by both reference and mode of presentation. According to Fodor's adaptation of this position, concepts are individuated not only by their contents, but also by the way in which this content is presented to thought, that is, its mode of presentation. For Fodor, modes of presentation are not Fregean senses, but rather language of thought expressions. And since, on present assumptions, propositions inherit their content from their constituent concepts (plus the syntax), this explains how different propositions can have the same content, and hence the same truth conditions.

The above gives some characterization of the abstract objects that might be propositional contents. A separate question is: in virtue of what does a particular representation come to have the content that it does? In other words, how is it that an abstract content becomes attached to a token mental representation? One answer that has been proposed is informational semantics. The difficulty, however, is that it is not obvious that informational semantics can work as an account of the content of the logico-syntactic apparatus. To see why, we need to first look at informational semantics in more detail.

2 Informational semantics

We are assuming that a token mental representation gets its content from its constituent concepts together with their mode of combination (the syntax). The problem of how a mental representation gets its content therefore reduces to the question of how primitive concepts do so.³

According to informational semantics (as set out in Dretske's seminal 1981 book *Knowledge and the Flow of Information*), primitive concepts get their content from their nomological relations to the entities that fall under them. Exactly how this

³ Dretske, though not Fodor, leaves open the possibility that some complex concepts get their content not from their constituents, but directly (that is, that the concept BLACK CAT, say, gets its content not from the contents of BLACK and CAT, but from its nomological links to black cats). It is not clear, though, that it is necessary—or wise—to do so. See Fodor (1990a: 58).

works is a somewhat complicated story, but in what follows I will try to summarize what are the essentials for our present purposes.⁴

As we noted above, there is clearly some sort of link between intentional content and representation. Part of the story as to how a mental state can be *about* black cats pretty clearly has to do with the fact that the mental state in question *represents* black cats. And representing black cats seems, intuitively at least, to have to do with *carrying information about* black cats. So it seems, at least *prima facie*, that the content of a mental state supervenes on the information that it carries. Since ‘information’ (at least in quantitative terms) is a well-defined mathematical concept, this looks to be a promising strategy for naturalizing intentional content—that is, for giving an account of intentional content in non-intentional terms. This is Dretske’s project.

There are two immediate problems that any attempt to reduce content to information must deal with, as we saw above. First, the content of a mental state is determinate, whereas the information carried is not. Thus, a mental state that carries the information *A* must also carry the information *A or B* and the information *A or B or C*, and so on; a mental state with the content *A*, however, does not also have the content *A or B*. Similarly, a mental state that carries the information that the temperature is 100 degrees also carries the information that the temperature is higher than 90 degrees, whereas a mental state with the content *the temperature is 100 degrees* does not thereby also have the content *the temperature is higher than 90 degrees*. Dretske’s solution to this problem is to introduce a distinction between digital and analogue forms of information carrying. Every signal carries information in both digital and analogue forms, but Dretske (1981: 137) stipulates that the most specific piece of information that a signal carries is the only piece it carries in digital form, with all other information being carried in analogue form. Dretske notes that sensory and cognitive processes can be distinguished on this basis: cognition consists in the extraction of certain pertinent information (in digital form) from the plethora of information presented by our senses (in analogue form). The content of a mental state, then, is related to the information carried by that state *in digital form*.

The second problem is how to account for false contents. Mental states have their contents independently of how the world is constructed, so the contents of mental states may be true or false. But it is not clear that informational semantics can accommodate this fact. If it is not the case that *s* is *F*, then there is by definition no information that *s* is *F*. So if the content of a mental state is related to the information carried by that mental state, it would not seem possible to account for false contents.

⁴ For a longer (and presumably more faithful) summary see Dretske (1983). A very clear and concise discussion of the issues is also to be found in Woodfield (1986).

A number of solutions to this problem have been proposed. Dretske's own solution (1981: Chapter 8) was to propose that there was a learning period during which mental states were trained to digitize a particular piece of the range of information presented in analogue form, perhaps through some kind of feedback mechanism. In this way, a particular mental state came to be associated with a particular piece of information, so that future mental state tokens of this type inherited this particular content, whether or not these subsequent mental tokens actually carried that information. This allows for the content of a token mental representation to be identified with the information carried (during the learning period) by the corresponding mental representation type, without it necessarily being the case that all tokens of that type (or indeed any, outside the learning period) carry that information. In this way, a token mental representation can misrepresent, and can therefore have a content that is false.

Fodor (1984) objects to Dretske's proposal, on various grounds. First, he points out that it is difficult to see how to draw a principled distinction between the learning period and the post-learning period. He also notes that Dretske's solution can only account for the misrepresentation of learned symbols, not innate symbols. Crucially, though, Fodor also points out that any situation that gives rise to a misrepresentation *B* after the learning period would also have caused a similar misrepresentation (that is, a false tokening of a mental representation 'A' that means *A*) if it had occurred during the learning period. How do we then declare that what was being learned was that the mental representation of type 'A' meant *A* rather than ($A \vee B$)? This is what Fodor calls the 'disjunction problem'. Notice that Dretske can't just do without the learning period and say that the content of a token mental representation is identified with the information typically carried (in digital form) by its corresponding representation type, because this leaves him immediately open to Fodor's disjunction problem again: if a representation of type 'A' typically carries the information that *A* and only occasionally the information that *B*, then this is better explained by that representation type meaning ($A \vee B$), which is the information that it carries even more typically than the information that *A*. In general, any apparent misrepresentation can always be subsumed by some sufficiently complex disjunction. So we are left without an account of how misrepresentation is possible.

One possible way around this problem would be to say that what a mental representation of type 'A' means is linked with the information it carries (in digital form) *in normal circumstances*. The difficulty then becomes to give a naturalistic account of what constitutes 'normal circumstances'. As Fodor (1984, 1987: Chapter 4, 1990a) points out, however, it's not at all clear that this is possible.⁵ The

⁵ One of the most thorough attempts to work out a theory of this kind was by Fodor himself (1984b/1990).

typical way of accounting for ‘normal circumstances’ is by appeal to teleology underwritten by Darwinian natural selection. Thus, the content of a mental representation of type ‘A’ can be identified with whatever it *normally* carries information about, which is a matter of what it was *selected* to carry information about. First, it’s not clear that natural selection always favours veridical representations—that is, that what’s selected for must ipso facto be true. Perhaps, argues Fodor, there are situations in which it is of greater survival value to an organism to represent something as false rather than veridically (think of repression of unbearable truths). Worse, it is often not clear exactly what information a mental representation type was selected to carry. Consider the famous case of the fly-eating frog, discussed in Fodor (1990a). Suppose (what is prima facie reasonable) that the frog tokens a mental representation of type ‘F’ in the presence of flies, and that this causes appropriate fly-catching behaviour. Suppose also that it’s not just flies, but moving black dots in general, that cause tokenings of ‘F’. Tokenings of ‘F’ caused by moving black dots other than flies are misrepresentations only on the assumption that in normal circumstances it’s flies that cause ‘F’-tokenings, and that this is so because natural selection favoured a mechanism that responded selectively to flies (so that ‘F’ has the content *fly*). But Fodor argues that there’s another equally valid way of telling this story: supposing that in the environment in which the mechanism evolved it happened that most moving black dots were in fact flies, we can say that natural selection favoured a mechanism that responded selectively to moving black dots. Thus, when the frog tokens ‘F’ in response to a non-fly this may result in a case of indigestion but it is not a case of misrepresentation (the content of ‘F’ in this case being *moving-black-dot*). How to choose between these two versions of the story? In terms of natural selection it doesn’t seem to matter. In fact, Fodor suggests that as far as natural selection is concerned, we could equally well say that the content of ‘F’ is *fly-or-inedible-moving-black-dot*. All that natural selection cares about is that in the frog’s environment (more accurately, in the environment in which the frog’s ‘F’-tokening mechanism evolved) what falls under FLY-OR-INEDIBLE-MOVING-BLACK-DOT tends to be flies. Fodor concludes from this that teleology can’t provide an answer to the disjunction problem.⁶

Having rejected both Dretske’s learning period solution and the teleology solution, Fodor then sets out his own solution to the disjunction problem (see 1990b). The basis of Fodor’s proposal is a fundamental asymmetry between false

⁶ There’s another problem for teleological theories that Fodor raises but which I won’t go into here. It’s that some tokens of ‘F’ that are not caused by F’s are nevertheless not errors (and therefore occur even in teleologically normal circumstances). For example, thinking about frogs may lead to thinking about flies (that is, to tokening FLY). This is not an erroneous tokening of FLY even though it’s been caused by something that isn’t a fly (viz., by a mental representation, in this case FROG). See Fodor (1990c).

tokens and true tokens. Suppose that cows cause COW-tokens, and so too (say) do some horses (distant ones in the evening light, perhaps). However, COW still means *cow* and not *cow-or-horse*. The reason for this, according to Fodor, is that those COW tokens that are caused by horses depend on the fact that there are COW tokens that are caused by cows, but not the other way round. As Fodor puts it, “noncow-caused COW tokens are *asymmetrically dependent upon* cow-caused COW tokens” (1990b: 91, original emphasis). Another way to think of this is that the nomological link between cows and COWs is in a certain sense more basic than the nomological link between distant horses in the evening light and COWs. The latter hijacks the mechanism that links cows with COWs (our cow-detector, as it might be), so if cows did not exist, the link between distant horses in the evening light and COWs would be severed, whereas if horses did not exist, cows would still cause COWs. This is the fundamental asymmetry that Fodor makes use of to try and solve the disjunction problem.

3 Problematic cases for informational semantics

A number of problems have been raised for Fodor’s asymmetric dependence theory as a solution to the disjunction problem. I will not discuss these here,⁷ but instead want to focus on a different kind of problem, which is that not all classes of concepts are amenable in any case to Fodor’s treatment.

As we have seen, according to Fodor’s informational semantics, concepts get their content from a nomological link with the property they express. This is fine for natural kind concepts like CAT and TREE, which plausibly express properties (*cathood* and *treehood*). And it may be that this treatment is extendable to nominal kind concepts such as BACHELOR which can be seen as expressing the property *bachelorhood*, and even to artefact concepts such as DOORKNOB which Fodor (1998) argues expresses the (mind-dependent) property *doorknobhood*. But it is not *prima facie* plausible that informational semantics could account for the content of proper names or the logico-syntactic apparatus, since CHOMSKY does not plausibly express the property of *chomskyness*, just as AND does not plausibly express the property of *andness*.

3.1 Proper names

First, consider proper name concepts. Unlike natural kind concepts, say, which express properties that any number of individuals can instantiate, proper name concepts are a species of individual concept: only one individual can ever fall under

⁷ See, however, Fodor (1990b), and the papers in Loewer & Rey (1991).

them.⁸ This is why it does not make intuitive sense to speak of a proper name concept (such as CHOMSKY) as expressing a property (*chomskyness*). There is no property, no hidden essence, possession of which will mean that an individual falls under the concept CHOMSKY, other than (trivially) the property of *being Chomsky*, a property that necessarily only Chomsky can instantiate. For familiar Kripkean reasons,⁹ even *the author of Syntactic Structures* won't do the trick, since it is impossible for Chomsky not to be Chomsky, while it is at least possible for Chomsky not to be the author of *Syntactic Structures*.

Fodor has at different times endorsed different accounts of the content of proper name concepts, usually without much discussion.¹⁰ In Fodor (1987: 84ff.), while noting that “The course of wisdom would be to reiterate the moral—viz., that names are a hard problem for everybody—and then to shut up and leave it alone”, he briefly sketches a variant of the description theory of names. According to this version, a concept such as CHOMSKY has as its content *the person named ‘Chomsky’*.¹¹ Importantly, this account is able to deal with Frege cases. For example, the concepts CICERO and TULLY are distinct, even though they both refer to the same individual, because they express different (linguistic) properties: the property of *being the person named ‘Cicero’* and the property of *being the person named ‘Tully’*. A well-known problem with this proposal, however, is that it implies “Cicero is named ‘Cicero’” is a necessary truth, when this is clearly not the case (he could have been called anything at all). This leads Fodor to reject this account as it stands.

Fodor's twist is to treat proper names like demonstratives, and claim that, for example, “Cicero was bald” says “he_{Cicero} was bald” and it presupposes that he is called ‘Cicero’. This allows “Cicero is named ‘Cicero’” to come out contingent, as it should be, since “he_{Cicero} is named ‘Cicero’” presupposes that he is called ‘Cicero’ only in this world, not in all possible worlds. It also allows us to maintain the intuition that, since *being Cicero* and *being Tully* are the same property, CICERO and TULLY have the same meaning. What makes them distinct concepts is that they differ in presupposition, and hence (at least on some accounts of presupposition) have different truth conditions.

⁸ It's possible, of course, for *no* individuals to fall under a proper name concept, as is the case with empty names (SANTA CLAUS, say).

⁹ See Kripke (1972/1980) and the extensive literature that this gave rise to. (For some recent attempts to defend descriptivism, see Stanley 1997, Sosa 2001, and Nelson 2002; Everett 2005 argues that these attempts are ultimately unsuccessful.)

¹⁰ Cain (2002: 116) suggests that Fodor endorses a causal-historical account of the content of proper names. While this is probably true of Fodor (1987) it is certainly not true of Fodor (1994), as we will see below.

¹¹ Such a theory was proposed by Kneale (1962), and has more recently been defended in Geurts (1997).

But note that this is not an informational semantic account, despite some suggestive language from Fodor (1987: 85, emphasis and bracketed sentence in original):

‘Cicero is Tully’ is informative because, although it doesn’t *say* that the guy who was called ‘Cicero’ was called ‘Tully’, it “carries the information” that he was. (For more on this notion of carrying information, see Dretske 1981 and Barwise & Perry 1981.)

What Fodor seems to be getting at is this: The sentence ‘Cicero is Tully’ expresses the proposition CICERO IS TULLY, which has the (referential) content *he_{Cicero} is he_{Tully}* (uninformative) but which presupposes *the person named ‘Cicero’ is the person named ‘Tully’* (informative). This is certainly not an informational semantic treatment. There is no nomological link or lawful correlation here between the concept CICERO and the information that Cicero is called ‘Cicero’ (as we have seen above, it can’t be the case that CICERO expresses the property of *being the person named ‘Cicero’*; if CICERO expresses any property, it’s the property of *being Cicero*). Rather, it seems that there could be a causal–historical link between the concept CICERO and the fact that Cicero is called ‘Cicero’. But to repeat, this is not a nomological link, and therefore this is not an informational semantic account in the sense proposed by Dretske and adopted by Fodor.¹² That it is a causal–historical account that Fodor has in mind is also made very clear by the fact that he explicitly likens his treatment of proper names to the treatment of demonstratives. And demonstratives in his view demand a causal–historical treatment if anything does: the only remotely plausible account of the content of the concept THAT BOOK (say) is whichever book *actually gave rise to that particular Mentalese token*.

In Fodor (1994: Appendix A) we get a rather different treatment of proper names, which Fodor explicitly contrasts with his earlier account outlined above,¹³ and which is basically an informational semantic account. Somewhere between 1987 and 1994 Fodor has realised that the above account doesn’t work: apart from anything else, as Fodor (1994:112f.) points out, it’s just not the case that there is anything specially metalinguistic about names. It’s true that “he is Cicero” invites the inference that he is called Cicero, but “that is a rose” can similarly invite the inference that that is called a rose, without anyone supposing that “rose” has metalinguistic properties.

¹² Fodor himself has been careful to stress that the causal–historical and the nomological should not be conflated, particularly in swamps. See (Fodor 1994: Appendix B).

¹³ See Fodor (1994: 111): “...I do want to stress the difference between this view and (what I’ll call) the Metalinguistic View, viz., that ‘Cicero’, but not ‘Tully’, means something like *is called ‘Cicero’*.”

According to Fodor's new account, it is plausible to assume that CICERO and TULLY carry the same information, since they express the property *being Cicero* and the property *being Tully*, which are plausibly the same property. Assuming an informational semantic account, CICERO and TULLY therefore have identical content. What, then, makes them different concepts? The normal way to proceed, as we have seen earlier, would be to propose that they have syntactically-distinct modes of presentation—that is, to say that one or other (or both) of them have complex modes of presentation.¹⁴ But this is implausible—why wouldn't CICERO and TULLY be syntactically primitive, just as the corresponding natural language words are? Fodor instead proposes that there must be some other (possibly neurological) difference between CICERO-tokens and TULLY-tokens which allows them to be type-distinct while having the same content and syntax. All that is left to explain is *how* the concepts CICERO and TULLY come to express the property of *being Cicero/being Tully*. That is, what mechanism is it that sustains this link? Here, Fodor (1994: 118f.) again adopts a causal–historical account, this time not as a metaphysical account of the content of proper names, but as an explanation for why the nomological link between the concept and the property holds. That is, the causal–historical properties of proper names are the mechanism that sustains the nomological link between proper name concepts and the corresponding properties. But, crucially, for Fodor it is the existence of such a link, not the details of the mechanism that sustains it, which underpins the metaphysics of proper name content.

3.2 Logical terms

The other class of concepts that presents difficulties for informational semantics is the logical terms (or, more broadly, the logico-syntactic apparatus). As was the case with proper names, it doesn't seem *prima facie* plausible that logical concepts get their content from a nomological link with the property they express. With proper names, this intuition stemmed from the fact that the entity they refer to seems not to be picked out by reference to any particular property that it possesses. In the case of logical concepts, it seems that they do not refer at all. It does not seem plausible, therefore, to adopt the usual informational semantic analysis and propose that AND (say) gets its content from a nomological link with a property of *andness* or *conjunction*.

¹⁴ Recall that modes of presentation are tokens of Mentalese. And tokens of Mentalese can be distinguished only by their syntax (in particular, they can't be distinguished phonologically or orthographically as natural language words can, because Mentalese has no phonology or orthography). Since primitive tokens of Mentalese are *ipso facto* syntactically identical, two concepts with the same content can be distinct only if one or both of them has a complex mode of presentation.

The usual approach taken by informational semanticists has been to cede the logical terms to inferential-role semantics. That is, it has generally been accepted that informational semantics is not the right approach for dealing with the logical vocabulary. The informational semantics dictum “meaning is reference” is fine for those concepts that refer, but for those concepts that do not refer, meaning must be constituted by something else, and inferential relations are an obvious candidate.

Recall that the metaphysical question, “what content does a particular concept have?” is answered by stating what contribution the concept makes to the truth conditions of propositions in which it occurs. In the case of logical terms, then, their content is the logical contribution they make to the propositions they occur in—that is, their logical properties. Take, for example, a subset of the logical terms, the logical connectives. The contribution that a logical connective makes to the truth conditions of a proposition in which it occurs is its particular (Boolean) function—that is, its truth table. This is to give a characterization of the abstract objects that might be the contents of the logical connectives. The question then is to give an account of how a token mental representation comes to have this abstract object as its content. This is the question which, in the case of logical terms, semanticists have traditionally answered with an inferential role account, typically in terms of implicit definitions or possession conditions which provide introduction and elimination rules for the concept.

For example, Peacocke (1992: 6) proposes that the possession conditions on the concept of conjunction can be identified with the transitions of the forms in (1), which a possessor of conjunction must find “primitively compelling”.¹⁵

- (1) a. $p, q / pCq$
 b. pCq / p
 c. pCq / q

These, of course, are the standard introduction and elimination rules for conjunction. As such, it is easy to show that together they uniquely specify the logical properties of conjunction (see below). Peacocke’s proposal is not this, which would be trivial, but rather that *grasping the rules in (1) in the right way just is to possess the concept of conjunction*. This is to claim that there is nothing more to having the concept of conjunction than finding its canonical introduction and elimination rules (primitively) compelling. In particular, there is no need to postulate mind–world links, as there is on an externalist account of concept possession. Recall that the informational semanticist views content as an abstract

¹⁵ For a thinker to find an inference primitively compelling, according to Peacocke, is for that thinker (i) to find the inference compelling, (ii) not to do so as a result of inferring it from something else, and (iii) not to necessarily take correctness of the inference as answerable to anything else.

object (a property), which individual token mental representations express in virtue of nomological links they have to this property. Contrast this with the view we are currently considering of the logical terms. On the inferential-role view, content is not an object at all—not even an abstract one—but rather a relational or dispositional state.¹⁶

However, Fodor (2004a, 2004b) now rejects this view of the content of logical terms, which he believes to be viciously circular. Instead he proposes that having the concept AND (say) doesn't depend on being disposed to accept the canonical AND-involving inferences (its introduction and elimination rules), but rather on being able to think conjunctive thoughts—that is, it depends on having a concept that means *and*. The details, however, are a little thin on the ground. In the following section, I develop an argument which demonstrates, in support of Fodor's position, that possessing the concept of a logical connective need not require accepting the full set of introduction and elimination rules for that concept. This casts doubt on the inferential-role view of concept possession according to which grasping the introduction and elimination rules is necessary and sufficient for possession of the concept. This still leaves open the crucial question of how a token mental representation comes to have the logical content that it does, to which Fodor did not give any detailed answer. In what follows, I will set out a psychologically plausible and independently motivated account which can provide an answer to this question (although not necessarily an account that a strict Fodorian—which Fodor sometimes is—would be completely comfortable with).

4 Logical connectives and their canonical inferences

Consider the truth-functional connective '∧'. This connective is governed by standard introduction and elimination rules as set out in (1). Is it necessary in order for a mind to have the concept AND that it grasps these three rules? According to inferential role semantics, the answer is 'yes'. On such an account, for a mind to have AND *just is* for that mind to grasp these three rules. I will argue, however, that the correct answer to this question is 'no'. It is possible for a mind that does not grasp all of these rules to nevertheless be able to think conjunctive thoughts and therefore to have the concept AND. In particular, it is possible for a mind which

¹⁶ Peacocke (2004: 98) dissociates himself from the view that contents are merely dispositional: "...it is very important to distinguish dispositionalism from relational individuation. Possession of a given concept is a relationally individuated state, in the sense that what makes it the state it is has to do with what judgements and transitions in thought a thinker is willing to make, when nothing interferes. But a relationally individuated state can be a categorical one, and not something merely counterfactual or dispositional." So be it.

does not grasp the introduction rule for conjunction to nevertheless have the concept AND. This is what I aim to demonstrate below.

First, consider why it is that grasp of the introduction and elimination rules in (1) is sufficient for grasp of AND. Given the propositions p , q , and $p * q$, grasp of the rules in (2) will be sufficient for grasping ‘*’ as AND.

- (2) a. $p, q / p * q$
- b. $p * q / p$
- c. $p * q / q$

Syntactically, ‘*’ operates as a connective. Grasping it as AND requires ruling out other inconsistent interpretations. In standard Boolean logic, there are 15 other possibilities which need to be ruled out, as shown in (3).

(3)

		And															
p	q	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
0	1	0	0	0	0	1	1	1	1	0	0	0	0	1	1	1	1
1	0	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	1
1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1

The three rules in (2) are sufficient for grasp of AND as they rule out the other 15 possibilities in (3). Since the content of a logical connective is just its truth table, for a logical concept to pick out a (unique) truth table is all that is required to constitute the content of that concept. The rules in (2) pick out the truth table for conjunction as follows. Rule a. states that if two propositions are true, connecting them with ‘*’ results in a proposition that is true. This rules out those functions that give a value of ‘0’ when both p and q have a value of ‘1’—that is, it rules out all the odd-numbered functions in (3). Rule b. states that if a proposition of the form ‘ $p * q$ ’ is true, then the constituent proposition p is true. This rules out those functions that give a value of ‘1’ when p has a value of ‘0’—that is, it rules out functions 5–16 in (3). Finally, rule c. states that if a proposition of the form ‘ $p * q$ ’ is true, then the constituent proposition q is true. This rules out those functions that give a value of ‘1’ when q has a value of ‘0’—that is, it rules out functions 3–4 and 7–16 in (3). The only function which has not been ruled out is function 2, which is the function for ‘and’. It follows that grasping the rules in (2) is sufficient for grasping ‘*’ as AND. No other Boolean connective is compatible with these three rules.¹⁷

¹⁷ I want to be clear on an important point. I am not proposing here a *psychological* procedure to check the consistency of candidate Boolean functions with the introduction and elimination rules attached to a concept (although it is plausible that our deductive device is supplemented by a procedure to monitor for contradictions—see Sperber & Wilson 1995: 102). This would be to

However, this in itself does not imply that grasping the three rules in (2) is *necessary* for grasping AND, even if there is no smaller set of rules which uniquely identifies ‘and’ among the possible Boolean connectives. To see why, suppose that there are other more general considerations which rule out some possible Boolean connectives. Suppose, for example, that our minds incorporate a general constraint which rules out truth-functional contradictions (that is, functions whose values are always false, regardless of the values of their inputs) as possible connectives.¹⁸ Such a general constraint, together with just the elimination rules in (2) (that is, just rules b. and c.), is sufficient to rule out the 15 other possibilities in (3). Notice that the only function eliminated by rule a. which is not also eliminated by rules b. or c. is function 1 (contradiction). Since this is ruled out by our hypothetical constraint, it follows that the elimination rules alone are sufficient to uniquely identify ‘*’ as AND. Thus, for any mind that rules out truth-functional contradictions on more general grounds, grasp of the elimination rules will be sufficient for grasp of AND.¹⁹

In a way, this supports Fodor’s (2004a, 2004b) position that having the concept AND (say) doesn’t depend on being disposed to accept its introduction and elimination rules, but rather on being able to think conjunctive thoughts (by having a concept that means *and*). I say ‘in a way’ because the considerations I have presented above show that accepting the introduction and elimination rules is not necessary for grasping the concept AND. This is consistent with Fodor’s position that having AND doesn’t depend on having these rules. But Fodor’s position is more general. It’s not that he thinks some of the canonical rules are unnecessary, it’s rather that he denies that *accepting inference rules is constitutive of concept possession*. We will return to this point below.

If our minds have a constraint such as the one I have proposed, then the standard inferential-role account of the content of AND is undermined. It is so far an open question, of course, whether our minds actually are like this, and whether the considerations set out above with respect to AND can be generalised to the other logical connectives. We also need to consider the question of whether a mind without AND-introduction could be effective at performing deductive inference. These are the questions I shall turn to in the following section.

propose that our minds represent the truth tables for all Boolean functions, and can select the appropriate truth table on the basis of the inference rules attached to a concept. I do not make this claim. Rather, I am claiming that the introduction and elimination rules associated with a logical concept are metaphysically sufficient to specify the content of that concept.

¹⁸ I will not motivate such an assumption here, because it is merely illustrative. I will demonstrate below, however, that a similar but more general assumption is well-motivated and psychologically real.

¹⁹ Whether it is also sufficient for deploying conjunction in mental inference is a separate issue that I will return to below.

5 Elimination rules and mental deduction

The reason that logicians have stated the meaning of ‘and’ in terms of the standard introduction and elimination rules is not just that these serve to fix the meaning of ‘and’ (that is, to uniquely specify the function it performs). After all, logicians can just specify the truth table to fix the meaning. The point about the introduction and elimination rules is that they are *rules*, and they are needed to support the process of deductive inference. The question then arises, if a mind dispensed with introduction rules, would it be able to perform deductive inferences?

It is not, in fact, implausible that our mental deduction device relies only on elimination rules, and no introduction rules. Sperber & Wilson (1995: Chapter 2) have argued convincingly that this is the case. Given a set of premises, there is an infinite set of conclusions that can be validly drawn using the standard introduction and elimination rules for the logical connectives. For example, from the assumption that p it is possible to derive an infinite number of conclusions of the form $p \vee q$ (for any q , regardless of truth value) using the rule of *or*-introduction. As Sperber & Wilson point out, this is unproblematic in an informal system of natural deduction, where it is left to the intelligent user of the system to determine which rules to apply at which point in a derivation. But in any characterization of our mental deduction systems, such profligacy is highly problematic. There are two specific problems which Sperber & Wilson raise. The first problem is that in any formal model of our deductive device in which they are incorporated, introduction rules will apply an infinite number of times to any set of assumptions, generating an infinite set of conclusions. The second problem is that the conclusions they derive are trivial in a certain (intuitive) sense: conjoining any arbitrary proposition to an assumption through (for example) *or*-introduction does not produce a conclusion that is useful to an organism in the sense of improving its representation of (that is, its understanding of) the world.

There are two reasons why it might be considered necessary to postulate both introduction and elimination rules. The first reason is that although introduction rules directly derive only trivial conclusions, they appear to be necessary inasmuch as these trivial conclusions are themselves needed as premises for the subsequent derivation of non-trivial conclusions. The second reason is that introduction rules appear to be required as they are constitutive of the content of logical concepts. We will consider each of these in turn.

First, consider the need for introduction rules in derivations. In order to derive the conclusions in (5) and (7) below from the premises in (4) and (6), a step of *and*-introduction or *or*-introduction seems unavoidable (examples taken from Sperber & Wilson 1995: 98).

- (4) a. If the trains are on strike and the car has broken down,
there is no way of getting to work.
b. The trains are on strike.
c. The car has broken down.
-
- (5) There is no way of getting to work
- (6) a. If the boiler needs repairing or the electricity has been
cut off, the house will be uninhabitable.
b. The boiler needs repairing.
-
- (7) The house will be uninhabitable.

The reason for this is that a standard logical derivation would proceed as in (8) and (9) respectively.

- (8) a. $(p \wedge q) \supset r$ [premise]
b. p [premise]
c. q [premise]
d. $(p \wedge q)$ [by \wedge -introduction from b. and c.]
 $\therefore r$ [by *modus ponens* from a. and d.]
- (9) a. $(p \vee q) \supset r$ [premise]
b. p [premise]
c. $(p \vee q)$ [by \vee -introduction from b.]
 $\therefore r$ [by *modus ponens* from a. and c.]

As can be seen, each of these derivations relies on the corresponding introduction rule. There is no reason to assume, however, that mental reasoning uses the standard rules of informal natural deduction. As Sperber & Wilson point out, in order to show that these introduction rules are required it would be necessary to show that the same derivations could not be carried out using alternative elimination rules, or that such rules were implausible on psychological grounds. But this is not the case. Sperber & Wilson show that the conjunctive and disjunctive versions of *modus ponens* in (10) and (11) obviate the need for any introduction rules. What is more, they argue on theoretical grounds that such rules are psychologically plausible, and cite experimental evidence (from Rips 1983) in favour of (11).

- (10) *Conjunctive modus ponens*
- a. *Input:* (if $(p$ and $q)$ then r)
 p
Output: (if q then r)
- b. *Input:* (if $(p$ and $q)$ then r)
 q
Output: (if p then r)
- (11) *Disjunctive modus ponens*
- a. *Input:* (if $(p$ or $q)$ then r)
 p
Output: r
- b. *Input:* (if $(p$ or $q)$ then r)
 q
Output: r

I now turn to the second reason why it may be considered necessary to postulate introduction rules in mental logic, the fact that they have generally been taken to be constitutive of the content of the logical connectives. We have already seen above that the standard introduction and elimination rules governing ‘and’ are sufficient to uniquely specify the appropriate Boolean function. The same is true of ‘or’ and ‘if...then’, as I will now briefly show.

The canonical rules governing the use of ‘or’ are given in (12).

- (12) a. $p / p * q$
 b. $q / p * q$
 c. $p * q, \neg p / q$
 d. $p * q, \neg q / p$

Together, these rules uniquely specify the corresponding Boolean function (function 8 in (13)). Rule a. eliminates functions with a value of ‘0’ when p has the value ‘1’—that is, functions 1–3, 5–7, 9–11 and 13–15. Rule b. eliminates functions with a value of ‘0’ when q has the value ‘1’—that is, functions 1–5, 7, 9–13, 15. Rules c. and d. both eliminate functions with a value of ‘1’ when p has a value of ‘0’ and q has a value of ‘0’—that is, functions 9–16.²⁰ The only function which has not been eliminated is function 8, the correct result.

²⁰ The fact that both of these rules eliminate the same functions means that one can be considered redundant in constituting the content of OR. We will see below, however, that once introduction rules are eliminated, both elimination rules are necessary to fix the content.

(13)

		Or								If...then							
<i>p</i>	<i>q</i>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
0	1	0	0	0	0	1	1	1	1	0	0	0	0	1	1	1	1
1	0	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	1
1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1

As regards ‘if...then’, consider the introduction and elimination rules in (14) below.

- (14) a. $\neg p / p * q$
 b. $q / p * q$
 c. $p * q, p / q$
 d. $p * q, \neg q / \neg p$

The introduction rules I have given in a. and b. are not the standard ones, although they do suffice to uniquely specify the appropriate function and logicians have claimed that “it is reasonable to require of any theory of the conditional that it explain why they hold”.²¹ The standard introduction rule for ‘if...then’ is the conditional proof, which can be expressed in the following form: “if q can be derived from the assumption that p , by some sequence of rules of the deductive device, then it may be inferred that $p \supset q$, whether or not it is in fact the case that p ”. This rule cannot be stated in terms of the truth values of p , q and $p * q$, and so cannot be used as a basis for eliminating candidate Boolean functions. In fact, however, the choice of introduction rules will not be important in what follows, since I propose that such rules can be eliminated in any case.

The rules in (14) specify the function corresponding to ‘if...then’ (function 14 in (13)) as follows. Rule a. eliminates functions with a value of ‘0’ when p has the value ‘0’—that is, functions 1–12. Rule b. eliminates functions with a value of ‘0’ when q has the value ‘1’—that is, functions 1–5, 7, 9–13 and 15. Rules c. and d. eliminate functions with a value of ‘1’ when p has the value ‘1’ and q has the value ‘0’—that is, functions 3, 4, 7, 8, 11, 12, 15, 16.²² The only function which has not been eliminated is function 14, the correct result.

One additional consideration for OR and IF...THEN, which did not arise for AND, is that the introduction and elimination rules attached to these concepts involve negation, and therefore presuppose the availability of NOT. It is not clear what to do about NOT. Giving a suitable implicit definition or possession conditions for NOT is

²¹ See Martin (1987: 14). In fact, these two rules, although valid in propositional logic, have been regarded by many as counterintuitive as applied to the natural language expression ‘if...then’.

²² Again, once introduction rules are eliminated we will see that both rule c. and rule d. are necessary for fixing the content of IF...THEN. (Cf. footnote 20.)

a difficult problem which has not been adequately solved (see Peacocke 2004). The standard rules of double negation introduction and elimination ($\lceil p / \neg\neg p \rceil$ and $\lceil \neg\neg p / p \rceil$), while they can be used to derive potentially useful rules such as modus tollens with a negated consequent ($\lceil p \supset \neg q, q / \neg p \rceil$), cannot be employed to introduce or eliminate single instances of negation, nor to specify the truth table for negation. In what follows I will assume that NOT is antecedently available to our minds—a reasonable, if unexplained, assumption.²³

We have seen briefly how the introduction and elimination rules for the logical connectives ‘and’, ‘or’ and ‘if...then’ can fix the content of the corresponding concepts. We are now in a position to return to the question of whether it is possible for a mind that does not grasp the full set of canonical inference rules for a given logical connective to nevertheless possess the corresponding concept. We have already seen that given a certain general assumption about the mind, it is possible to grasp the concept AND by grasping only its elimination rules. I will argue in what follows that given a more general assumption about the mind (which I will show to be motivated), this argument can be extended to all of the standard logical connectives (conjunction, disjunction and implication).

The assumption is that the meaning postulates (that is, inference rules)²⁴ attached to a concept constrain the interpretation of that concept in two distinct ways, as set out in (15).

(15) a. **Input consistency constraint**

For a concept C with attached meaning postulates MP_1, MP_2, \dots, MP_n , an interpretation of C that makes any of MP_1, MP_2, \dots, MP_n vacuous, in the sense that the input conditions for that postulate are contradictory, is ruled out as a possible interpretation of C .

b. **Postulate validity constraint**

For a concept C with attached meaning postulates MP_1, MP_2, \dots, MP_n , an interpretation of C that makes any of MP_1, MP_2, \dots, MP_n invalid is ruled out as a possible interpretation of C .

These constraints, although they may have psychological justification (see below), are taken to be metaphysical rather than psychological in nature. That is, it’s not

²³ Since ‘ \vee ’ and ‘ \rightarrow ’ can be defined in terms of ‘ \wedge ’ and ‘ \neg ’, it follows that a mind which grasps the latter has all the resources necessary for grasp of the former. It does not, however, follow that such a mind has the concepts OR and IF...THEN. Possession of a logical concept is not a question not of whether a mind has the necessary resources *in principle*, it is a question of whether a mind possesses a concept with the appropriate meaning (that is, in the case of logical concepts, whether it has a concept with the appropriate logical properties).

²⁴ I use the term “meaning postulate”, in the sense of Fodor (1975) and J. D. Fodor (1977), to refer to mental rules of inference.

that the mind *uses* the constraints to pick out the relevant truth table, and then reasons on the basis of this truth table. Rather, the mind reasons according to the meaning postulates attached to a concept, which has the (metaphysical) effect of ruling out some truth tables as interpretations of this concept.

A different way of stating the input consistency constraint is that any truth-table which makes the input conditions for any of the inference rules attached to a concept inconsistent is ruled out as a possible truth table for that concept. The justification for the input consistency constraint is this. A meaning postulate whose input conditions were contradictory would have no utility, as there would be no logically possible circumstances under which it would apply. Not only would such a meaning postulate be useless, but its origin would also be obscure. A postulate could conceivably either be innate or acquired. But it is difficult to see how such a meaning postulate could be innate, given that it does not - nor could it ever have - served any purpose. It is similarly difficult to imagine how such a postulate could come to be acquired.

The postulate validity constraint is just a restatement of the status of meaning postulates: that is, any interpretation which would make the meaning postulates attached to a concept invalid, in the sense that the output could be false when the inputs were all true, is ruled out. As such, this constraint requires no further justification.

Given these two constraints, it can be shown that the contents of the concepts of conjunction, disjunction and implication are constituted by their elimination rules, without the need for introduction rules. A summary of the following results is provided in (20) below.

Consider first the concept of conjunction, which has (at a minimum) the meaning postulates in (16).

(16) a. *Input:* $p * q$
 Output: p

 b. *Input:* $p * q$
 Output: q

By the input consistency constraint, any truth table which makes the input conditions for any of the meaning postulates contradictory is ruled out. In this case, since each meaning postulate has only one input condition, any truth table on which that input is self-contradictory will be ruled out. This eliminates truth table 1.

By the postulate validity constraint, any truth table which makes one or more of the meaning postulates invalid is ruled out. This means that any truth table on which the input of one of the meaning postulates is true and the output false is ruled out. As we have already seen, this eliminates truth tables 3–16. The only remaining truth table is 2, the truth table for AND.

Next consider the concept of disjunction, which has minimally the meaning postulates in (17).

(17) a. *Input:* $p * q$
 $\neg p$
Output: q

b. *Input:* $p * q$
 $\neg q$
Output: p

The input consistency constraint rules out any truth table on which the two inputs of either meaning postulate are contradictory. This rules out truth tables 1–6. The postulate validity constraint rules out any truth table on which either of the meaning postulates is invalid, eliminating truth tables 9–16. This leaves two truth tables remaining (7 and 8), those for OR (inclusive ‘or’) and XOR (exclusive ‘or’), an interesting result.

A grasp of the introduction and elimination rules for ‘or’, as given in (12), is sufficient for grasp of OR, but excludes XOR, as we saw earlier (see the table in (13)). That is, the full set of canonical inference rules uniquely specifies inclusive ‘or’. A grasp of the elimination rules only, together with the general principles I have proposed, can eliminate all non-disjunctive possibilities, but leaves open both OR and XOR interpretation. Let us call this underspecified concept ‘DISJUNCTION’. This result is interesting, because it has often been claimed that English ‘or’ can have both inclusive and exclusive interpretations.²⁵ On the current proposal, both of these interpretations are left open.

This suggests one of two possibilities. The first is that DISJUNCTION could be genuinely ambiguous at the semantic level between inclusive and exclusive readings. There is no reason to believe that this is the case, however. On the assumption that English ‘or’ expresses the concept DISJUNCTION, the ‘ambiguity hypothesis’ would lead us to expect that sentences with ‘or’ would be regarded by speakers as ambiguous in the same way that sentences with polysemous lexical items are, which is not the case. The second possibility is that DISJUNCTION expresses a general, underspecified,²⁶ meaning, which can be further specified through pragmatic processes such as concept narrowing. This is the much more

²⁵ There has been fairly extensive discussion of this question in the literature. No clear consensus has been reached, and the three obvious possibilities have all been argued for: that English ‘or’ is always inclusive (Pelletier 1977, Lepore 2000), that it is (virtually) always exclusive (Lakoff 1971), and that ‘or’ has two possible meanings, the one to be adopted being determined by pragmatic factors (Hurford 1974). See Evans et al. (1993: Chapter 5) for a review.

²⁶ Note that to say that the concept is *underspecified* is not to say that it is *indeterminate*. This distinction is discussed in more detail in relation to implication (see footnote 34 below).

likely possibility. For example, we could suppose that a pragmatic principle assigns the more general meaning—that is, the one giving rise to fewer entailments, in this case OR—in situations such as this (cf. Chierchia & McConnell-Ginet 2000: 78–79) and that the narrower meaning—in this case, XOR—could arise via conversational implicature (specifically, a scalar implicature). See Grice (1967), Sperber & Wilson (1995: Postface), Carston (1998), and Noveck (2004).

It could also be the case that additional meaning postulates are the mechanism through which concept narrowing is achieved. For example, additional elimination rules such as those expressed by the meaning postulates $\lceil p * q, p / \neg q \rceil$ and $\lceil p * q, q / \neg p \rceil$ would act in this way, as they serve to rule out the inclusive reading and therefore specify the (narrower) concept XOR. This could be a way to account for languages such as Latin, which had separate lexical items ‘vel’ and ‘aut’, according to some accounts corresponding to inclusive and exclusive disjunction, respectively.²⁷ Such an account could also be adopted if English ‘or’ can be demonstrated to have only an exclusive interpretation.²⁸ Notice, however, that no additional set of meaning postulates based on elimination rules can perform the opposite function—that of eliminating the narrower meaning (XOR) in favour of the more general meaning (OR). It follows that if meaning postulates are restricted to elimination rules, as we have been supposing, the concept OR cannot be uniquely specified (other than by pragmatic means).²⁹

Finally, consider the concept of implication, which has at least the meaning postulates in (18).³⁰

²⁷ This analysis, although the received wisdom in most introductory logic texts, is controversial. See Jennings (1994: 239–251).

²⁸ Cf. footnote 25 above.

²⁹ The following is an informal proof of this. Suppose that some set of meaning postulates based on elimination rules could rule out XOR. Now, on the approach we are taking there are two possible ways to rule out a truth table: the input consistency constraint, and the postulate validity constraint. Consider each in turn. First, it can be shown that the input consistency constraint cannot rule out XOR. The reason is that there is no condition where XOR is false when any one of $p, q, \neg p, \neg q$ are true, and therefore no possibility for a postulate with inconsistent inputs. Next, it can be shown that the postulate validity constraint cannot rule out truth tables satisfying the condition $\langle |p| = 1, |q| = 1, |p * q| = 0 \rangle$ (the only condition on which OR and XOR differ), because this constraint works by ruling out truth tables that make the output of the postulate false when its inputs are all true. But, on the assumption there are no introduction rules, the string $\lceil p * q \rceil$ must be an input to any postulate, implying that this constraint only looks at conditions where $|p * q| = 1$, which it is not in the case we are considering. By analogous reasoning, it is not possible (other than by pragmatic means) to uniquely specify the concepts CONDITIONAL, REVERSE-CONDITIONAL or NAND via elimination rules.

³⁰ The psychological evidence suggests, in fact, that while the first of these meaning postulates, corresponding to modus ponens, is directly represented and highly accessible, the second meaning postulate, corresponding to modus tollens, is not directly represented, or at least not highly accessible (see Evans et al. 1993: Chapter 2). The effects of modus tollens, and hence the

(18) a. *Input:* $p * q$

p

Output: q

b. *Input:* $p * q$

$\neg q$

Output: $\neg p$

The input consistency constraint for these two meaning postulates rules out truth tables 1, 2, 5, 6, 9 and 13. The postulate validity constraint rules out truth tables 9–16. As with disjunction, this leaves two truth tables remaining—this time, 10 and 14, the truth tables for the biconditional and the conditional.

This is also an interesting result. The difference between the conditional and the biconditional is that the former, but not the latter, is true when the antecedent is false and the consequent is true. This is precisely the condition that people have most intuitive difficulty with.³¹ Indeed, natural language utterances containing ‘if...then’ can often have both conditional and biconditional interpretations.³² Consider for example the utterance in (19).³³

metaphysical constraints it places on possible truth-functional interpretations, can be obtained from modus ponens together with a form of reductio ad absurdum: given the premises of modus tollens ($p * q, \neg q$) assume that p , from which by modus ponens we can conclude that q , contradicting one of the initial premises and falsifying our assumption that p (see Evans et al. 1993: 14–15 and Braine & O’Brien 1991 for discussion and detailed proposals; for psychological evidence in support of the proposal that in evaluating conditionals people create an imaginary world that includes the assumption that p , or at least focus on the possibility that the antecedent is true, see Hadjichristidis et al. 2001, Over & Evans 2003 and Evans et al. 2003, as well as Sperber et al. 1995). Alternatively, it cannot of course be ruled out that the only meaning postulate in this case is modus ponens, and that there are no other metaphysical constraints in play. This would give rise to an underspecified concept allowing AND, Q-IDENTITY, CONDITIONAL and BICONDITIONAL interpretations. Perhaps Q-IDENTITY can be generally excluded on pragmatic grounds: why use a connective when the antecedent is always irrelevant to the truth value, rather than just stating q ? The remaining three possible interpretations could be further specified pragmatically (by narrowing, assuming a general principle assigning the most general interpretation). Interestingly, this might explain the characteristic pattern of errors associated with conditional reasoning, including those of children, who seemingly go through three developmental stages: they initially interpret conditional statements as existential conjunctions, later as biconditionals, and finally as true conditionals (see Evans et al. 2003, Barrouillet & Lecas 1998, and Barrouillet et al. 2000). Sperber et al. (1995) show how pragmatic considerations can give rise to conjunctive interpretations for conditional statements.

³¹ Although there are several possible reasons for this, as the natural language uses of ‘if...then’ are not limited to expressing truth-functional relations.

³² The situation is similar to that of disjunction (see footnote 25 above), and the corresponding literature on implication is vast. Some consider that English ‘if...then’ expresses the conditional

(19) If the train is on time then I'll be at your office at four o'clock

This utterance can be interpreted on either a conditional or a biconditional sense of 'if...then'. It could be interpreted as stating merely a *sufficient* condition (the train being on time) for arriving at four o'clock, in which case 'if...then' is to be given a conditional interpretation (even if the train is late, I could rush to the office and still be there at four o'clock). Alternatively, the utterance could be interpreted as stating a *necessary* condition for arriving at four o'clock (I'll be rushing already, so if the train is late I certainly won't be at the office at four o'clock).

As with disjunction, then, the present analysis captures the range of observed truth-functional interpretations. Again, there are two possibilities: semantic ambiguity or underspecification. For the same reasons as previously discussed, the latter seems preferable.³⁴ I'll refer to the underspecified concept as IMPLICATION. Additional meaning postulates or pragmatic concept narrowing could provide further specification of IMPLICATION. Pragmatic narrowing would work analogously to the case of DISJUNCTION—again, we have a more general meaning (CONDITIONAL) and a narrower meaning (BICONDITIONAL), and the latter could arise as an interpretation of the former via scalar implicature (see above). Alternatively, additional elimination rules such as those expressed by the meaning postulates $\lceil p * q, q / p \rceil$ or $\lceil p * q, \neg p / \neg q \rceil$ would serve to rule out the material implication reading and therefore specify the (narrower) concept BICONDITIONAL.

(material implication) and never the biconditional (material equivalence). Others allow that it may also (and perhaps always in young children) express the biconditional, which can explain why subjects often draw the 'fallacious' conditional inferences—denial of the antecedent and affirmation of the consequent. Many, however, consider that it has a significant non-truth-conditional component. See Evans et al. (1993: Chapter 2) for a review.

³³ Cf. Partee et al. (1993: 102–104).

³⁴ To say that the concept is *underspecified* is not to say that some of its truth conditions are *indeterminate*. To see this, compare a classical two-valued logic of the kind assumed here with a three-valued logic allowing the values true, false and indeterminate. In the system proposed here, all cells in a truth table must, when filled at all, be filled either with a '0' or a '1'. However, some truth tables are underspecified in the sense that the metaphysical constraints imposed by the meaning postulates attached to a concept may not be sufficient to determine the values of these cells, which may be done by pragmatic concept-narrowing processes. By contrast, it has been suggested by Wason (1966), Johnson-Laird & Tagart (1969), Evans & Over (2004) and others that certain truth tables might be defective, in the sense that the values of certain cells could be indeterminate (particularly the cells corresponding to a false antecedent in the truth table for the conditional) because a conditional statement is irrelevant in such circumstances (see Evans et al. 1993, Chapter 2). This is to propose a determinate three-valued logic, since there is no lack of specification here—the indeterminate cells are not awaiting a value, they already have one (viz., 'indeterminate', in the sense of irrelevant). This point is important, because the 'defective truth table' account has certain highly implausible implications (see Johnson-Laird 2005).

As with DISJUNCTION, no set of meaning postulates based on elimination rules could rule out the BICONDITIONAL meaning.³⁵

All the above results for conjunction, disjunction and implication are summarized in the table in (20) below.³⁶

(20)

		Contradiction	And	p and not q	p-identity	q and not p	q-identity	Xor	Or	Nor	Biconditional	Not-q	Reverse-conditional	Not-p	Conditional	Nand	Tautology
p	q	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
0	1	0	0	0	0	1	1	1	1	0	0	0	0	1	1	1	1
1	0	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	1
1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1

Conjunction	ICC	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	PVC	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Disjunction	ICC	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>							
	PVC	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>					
Implication	ICC	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	PVC	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

It has been argued by Sperber & Wilson (1995) that our mental deductive device plausibly operates with elimination rules but no introduction rules. I have argued above, on independent grounds, that given a reasonable assumption about the mind the content of the logical connectives can be determined (with a certain degree of principled underspecification) purely on the basis of the proposed meaning postulates (which correspond to the canonical elimination rules for these connectives). This undermines inferential-role accounts of logical content, according to which grasping a logical concept *just is* to grasp both its introduction and elimination rules. In doing so, the present account also provides support for Fodor’s (2004a, 2004b) claim that possessing a logical term is a matter not of being disposed to accept the term’s canonical inferences, but of possessing a concept with the right meaning. His claim, then, is more general, as he does not believe that *any* meaning postulates are constitutive of logical content.

³⁵ Cf. footnote 29 above.

³⁶ In this table, ‘ICC’ stands for the input consistency constraint, and ‘PVC’ stands for the postulate validity constraint.

In the following section I address the question of whether meaning postulates are to be seen as content constitutive.

6 Are meaning postulates content constitutive?

What can we now say about whether meaning postulates are content constitutive? First, note that the discussion above is at least suggestive of meaning postulates being constitutive of the content of logical terms. After all, we have seen how meaning postulates impose metaphysical constraints on the content of the concept to which they are attached, which would tend to indicate that meaning postulates are content constitutive. In this section, we will consider whether such a claim can be supported.

In deciding on the constitutivity of meaning postulates, two tests are important:

1. Can we have cases where we possess a logical concept *without* having the associated meaning postulates? and
2. Can we have cases where we possess a logical concept with the associated meaning postulates, but where these meaning postulates *fail* to make a contribution to content?

I will discuss each of these tests in turn.

6.1 Content without meaning postulates?

The first test is important because if it can be shown that the content of a logical concept may be constituted in the absence of any prescribed set of meaning postulates, this could suggest that meaning postulates are not content constitutive. After all, if an inference rule is not necessary for the possession of a concept, it is difficult to see how that inference rule could be constitutive of the content of that concept.

Inferential-role accounts of content, for example, deny that we can possess a logical concept without possessing/grasping the canonical inference rules (that is, meaning postulates) for that concept. According to inferential-role accounts, the identity conditions for a concept are the same as the possession conditions for that concept. It follows on such an account that if a given inference rule is not among the possession conditions for a concept, then that inference rule cannot be content constitutive. Or, to put things the other way round, on an inferential-role account any meaning postulate which is content constitutive is also required for possession of the concept.

In fact, however, it is clear from the earlier discussion that content can be constituted without the need for any particular set of meaning postulates. To see this, first note that there is no a priori reason why we need meaning postulates at all in order to fix the content of a logical connective (although we may of course need them for other reasons, such as for mental deduction). There is nothing incoherent about the idea of a mind which represents logical properties in some other way—say, by directly representing the relevant truth table.

Beyond this, we have seen that different combinations of meaning postulates and general principles can serve to pick out the same truth table. For example, a concept ‘*’ will be the concept AND if it has attached meaning postulates corresponding to the canonical introduction and elimination rules for conjunction, given in (2) above. It is clear, however, that these three meaning postulates are not *required* in order to fix the content of AND. We have already seen above that, given a certain plausible assumption about the mind, we can dispense with the meaning postulate corresponding to ‘and’-introduction. Even restricting ourselves to elimination rules, there can be different sets of meaning postulates that fix the content of a particular connective.

Given that there is no particular meaning postulate or set of meaning postulates that is required in order to fix the content of a connective, must we conclude that meaning postulates are not content constitutive? Perhaps not. The other possibility is to allow that, *contra* inferential-role theorists, something can be constitutive of the content of a concept without being *necessarily* present (that is, without being a possession condition). I will argue in favour of such a proposal, and against the inferential-role account.

The position taken by inferential-role theorists is influenced by the fact that in cases such as conjunction there do not seem to be any alternatives: if not the canonical introduction and elimination rules, then what other rules could determine the content of AND? If another non-equivalent set of rules could be found which could also determine the content of AND, then the identification of content-constitutive inferences with possession conditions would be undermined. So, to allow that a meaning postulate can be content constitutive without necessarily being present is just to allow that there can be more than one way to constitute the content of a concept.

In fact, of course, the previous section demonstrated just that: different combinations of meaning postulates and general principles can serve to fix the content of a logical connective. In which case, there is no justification for identifying content-constitutive inferences with possession conditions. All that we can say of meaning postulates is this: *if* a particular meaning postulate is present, then it constitutes (partly or wholly) the content of the concept. But this is not to say that such a meaning postulate *must* be present.

There is a potential objection to this line of reasoning, which is that although meaning postulates may play a critical role in fixing content, this does not mean that they have to *constitute* content. After all, it is perhaps not *prima facie* implausible to see meaning postulates as just another kind of sustaining mechanism. Compare: a mechanism to recognize tigers may play a critical role in fixing TIGER-content, but this doesn't mean that such a mechanism is constitutive of TIGER-content. On such a view, one could regard the abstract logical properties of the concept (that is, its truth table) as giving the content, and the meaning postulates as merely providing a means of what Fodor would call "semantic access" to the content. It's *that* the right logical properties are picked out, not *how* they're picked out, that's important. In the same way, a tiger-detector provides a sustaining mechanism linking TIGER with its content, the abstract property of *tigerhood*. It's *that* the concept picks out the right property, not *how* it does so, that's important to informational semantics.

On reflection, however, it's not clear that this will do. According to an informational semantic account, the concept TIGER (and the concept BACHELOR, and even the concept DOORKNOB) means what it does because of a nomological link between the concept and the corresponding property. As such, it's the existence of the link, not the mechanisms by which the link is sustained, that determine the content. But the situation is different for the logical terms. It would be distinctly odd to claim that it is in virtue of some nomological link between AND and the property *conjunction* that AND means what it does.³⁷ It is not plausible to see AND and other logical concepts as *referring* to some abstract logical property. Rather, AND means what it does because it possesses the logical properties of conjunction—a token of AND is an *instance* of conjunction.³⁸ Compare this with TIGER, which refers to tigers, but which doesn't have the properties of tigers (unlike tigers, TIGERs have no stripes)—a TIGER-token is certainly not an *instance* of tigerhood. The right question to be asking about AND, therefore, is not how it gets linked to the property it expresses, but what gives it the properties that it possesses. Sustaining mechanisms can't provide an answer to this question, but meaning postulates just might.

I have argued that there is no privileged set of content-constitutive meaning postulates, but that this does not mean that meaning postulates are not content constitutive. Rather, meaning postulates are content constitutive where present, but no particular meaning postulate is required to be present. Should you find this argument unconvincing, however, the second test proves to be conclusive.

³⁷ Cf. Prinz & Clark (2004) and Fodor's (2004b) reply.

³⁸ Assuming the usual caveats concerning cases where 'AND' is mentioned rather than used.

6.2 Meaning postulates without content?

The second test is in fact the crucial one for determining whether meaning postulates are content constitutive. If meaning postulates can be present without being (partly or wholly) constitutive of content, then meaning postulates clearly cannot be content constitutive.

Suppose for the sake of argument that meaning postulates could be present without being constitutive of content. This would be to claim that although the mind might make use of meaning postulates for reasoning, they do not fix the content of the concepts to which they are attached. But it is difficult to see how this could be the case. Suppose that a mind has a concept ‘*’ with the meaning postulates $\lceil p * q, \neg p / q \rceil$ and $\lceil p * q, \neg q / p \rceil$ attached. These meaning postulates serve to pick out truth tables 7 and 8 in (20), given the constraints set out in (15). Suppose now that the mind acquires a new set of inference rules, $\lceil p * q, p / \neg q \rceil$ and $\lceil p * q, q / \neg p \rceil$, which it attaches to ‘*’ as meaning postulates. The crucial point to note is that the effect of attaching these meaning postulates is to *thereby* change (/narrow) the content of ‘*’ from DISJUNCTION to XOR. For now the set of meaning postulates attached to ‘*’ picks out a different set of truth tables—that is, just truth table 7. And it just cannot be the case that the content of a connective is different from the truth table with which it is associated.

This same example, however, also raises a potential difficulty. For notice that in order to narrow the content of ‘*’ from DISJUNCTION to XOR we have proposed the addition of two further meaning postulates, $\lceil p * q, p / \neg q \rceil$ and $\lceil p * q, q / \neg p \rceil$. While it is reasonable to assume the addition of both of these postulates, in order that our mental deduction can be systematic, either one of these postulates alone is sufficient to narrow the content of ‘*’ to XOR. Since one of these postulates will be unnecessary for specifying the content, does this mean that only one is content constitutive? Which one?

In fact, this seems to be the wrong way of looking at things. Whether a meaning postulate is content constitutive or not is a question of whether it constrains the (semantic) interpretation of the concept or not. It is irrelevant whether the constraints it imposes are spurious—what determines whether a postulate is content constitutive is whether it has an *input* into content, not whether that input has any *effect* on the content in a particular case. I would therefore argue that the right thing to say in a situation such as the one we have been considering is that both meaning postulates are content constitutive, and that their contribution to content (that is, the constraints they impose) happens to be identical. If either postulate was removed, the constraints imposed by the other would be substantive.

We can see from the above considerations that, when present, meaning postulates cannot fail to be constitutive of content. This I take to demonstrate conclusively that the meaning postulates attached to a concept are content constitutive.

References

- Barrouillet, P., N. Grosset & J.-F. Lecas (2000) Conditional reasoning by mental models: chronometric and developmental evidence. *Cognition* 75: 237–266.
- Barrouillet, P. & J.-F. Lecas (1998) How can mental models theory account for content effects in conditional reasoning? A developmental perspective. *Cognition* 67: 209–253.
- Barwise, J. & J. Perry (1981) *Situations and Attitudes*. Cambridge, Mass.: MIT Press.
- Braine, M. D. S. & D. P. O'Brien (1991) A theory of *if*: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review* 98: 182–203.
- Burge, T. (1979) Individualism and the mental. In P. A. French, T. E. Uehling & H. K. Wettstein (eds) *Studies in Metaphysics*. Midwest Studies in Philosophy, vol. IV, 73–121. Minneapolis: University of Minnesota Press.
- Cain, M. J. (2002) *Fodor: Language, Mind and Philosophy*. Cambridge, UK: Polity.
- Carston, R. (1998) Informativeness, relevance and scalar implicature. In R. Carston & S. Uchida (eds) *Relevance Theory: Applications and Implications*, 179–236. Amsterdam: John Benjamins.
- Chierchia, G. & S. McConnell-Ginet (2000) *Meaning and Grammar*, 2nd edition. Cambridge, Mass.: MIT Press.
- Dretske, F. (1981) *Knowledge and the Flow of Information*. Cambridge, Mass.: MIT Press.
- Dretske, F. (1983) Précis of *Knowledge and the Flow of Information*. *Behavioral and Brain Sciences*, 6: 55–90. (Includes commentaries and author's response)
- Dretske, F. (1986) Misrepresentation. In R. Bogdan (ed.) *Belief*, 17–36. Oxford: Oxford University Press.
- Evans, J. St. B. T., S. J. Handley & D. E. Over (2003) Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29: 321–335.
- Evans, J. St. B. T., S. E. Newstead & R. M. J. Byrne (1993) *Human Reasoning: The Psychology of Deduction*. Hove, UK: Lawrence Erlbaum.
- Evans, J. St. B. T. & D. E. Over (2004) *If*. Oxford: Oxford University Press.
- Everett, A. (2005) Recent defences of descriptivism. *Mind & Language* 20: 103–139.
- Fodor, J. A. (1975) *The Language of Thought*. New York: Crowell.
- Fodor, J. A. (1978) Propositional attitudes. In *Representations*, 177–203. Cambridge, Mass.: MIT Press.
- Fodor, J. A. (1980) Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences* 3: 63–109.
- Fodor, J. A. (1984a) Semantics, Wisconsin style. *Synthese* 59: 231–250.
- Fodor, J. A. (1984b/1990) Psychosemantics or: Where do truth conditions come from? Ms., MIT. (Published in W. G. Lycan (ed.) *Mind and Cognition*, 312–337. Oxford: Blackwell, 1990.)
- Fodor, J. A. (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, Mass.: MIT Press.
- Fodor, J. A. (1990a) A theory of content, I: The problem. In *A Theory of Content and Other Essays*, 51–87.
- Fodor, J. A. (1990b) A theory of content, II: The theory. In *A Theory of Content and Other Essays*, 89–136.
- Fodor, J. A. (1990c) Information and representation. In P. Hanson (ed.) *Information, Language, and Cognition*, 175–190. Vancouver: University of British Columbia Press.
- Fodor, J. A. (1991) A modal argument for narrow content. *Journal of Philosophy* 88: 5–26.
- Fodor, J. A. (1994) *The Elm and the Expert: Mentalese and its Semantics*. Cambridge, Mass.: MIT Press.

- Fodor, J. A. (1998) *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J. A. (2004a) Having concepts: a brief refutation of the twentieth century. *Mind and Language* 19: 29–47.
- Fodor, J. A. (2004b) Reply to commentators. *Mind and Language* 19: 99–112.
- Fodor, J. D. (1977) *Semantics: Theories of Meaning in Generative Grammar*. New York: Crowell.
- Geurts, B. (1997) Good news about the description theory of names. *Journal of Semantics* 14: 319–348.
- Goodman, N. (1976) *Languages of Art*, 2nd edition. Indianapolis: Hackett.
- Grice, H. P. (1957) Meaning. *Philosophical Review* 66: 377–88.
- Grice, H. P. (1967) *Logic and Conversation*. William James Lectures, reprinted in Grice 1989: 1–143.
- Grice, H. P. (1989) *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Hadjichristidis, C., R. J. Stevenson, D. E. Over, S. A. Sloman, J. St. B. T. Evans & A. Feeney (2001) On the evaluation of ‘if p then q’ conditionals. In J. D. Moore & K. Stenning (eds) *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 409–414. London: Lawrence Erlbaum.
- Hurford, J. R. (1974) Exclusive or inclusive disjunction. *Foundations of Language* 11: 409–411.
- Jennings, R. E. (1994) *The Genealogy of Disjunction*. Oxford: Oxford University Press.
- Johnson-Laird, P. N. (2005) If bears eat in the woods...? Review of Evans & Over 2004. *Trends in Cognitive Sciences* 9: 43–44.
- Johnson-Laird, P. N. & J. Tagart (1969) How implication is understood. *American Journal of Psychology* 82: 367–373.
- Kneale, W. (1962) Modality de dicto and de re. In E. Nagel, P. Suppes & A. Tarski (eds) *Logic, Methodology and Philosophy of Science. Proceedings of the 1960 International Congress*, 622–633. Stanford: Stanford University Press.
- Kripke, S. (1972/1980) *Naming and necessity*. Cambridge, Mass.: Harvard University Press.
- Lakoff, R. (1971) If’s, and’s and but’s about conjunction. In C. J. Fillmore & D. T. Langendoen (eds) *Studies in Linguistic Semantics*, 114–149. New York: Holt, Rinehart & Winston.
- Lepore, E. (2000) *Meaning and Argument*. Oxford: Blackwell.
- Loewer, B. & G. Rey (eds) (1991) *Meaning in Mind: Fodor and His Critics*. Oxford: Blackwell.
- Martin, J. N. (1987) *Elements of Formal Semantics*. London: Academic Press.
- Nelson, M. (2002) Descriptivism defended. *Noûs* 36: 408–436.
- Noveck, I. A. (2004) Pragmatic inferences linked to logical terms. In I. A. Noveck & D. Sperber (eds) *Experimental Pragmatics*, 301–321. Basingstoke: Palgrave Macmillan.
- Over, D. E. & J. St. B. T. Evans (2003) The probability of conditionals: the psychological evidence. *Mind & Language* 18: 340–358.
- Partee, B. H., A. ter Meulen & R. E. Wall (1993) *Mathematical Methods in Linguistics*. Dordrecht: Kluwer.
- Peacocke, C. (1992) *A Study of Concepts*. Cambridge, Mass.: MIT Press.
- Peacocke, C. (2004) Interrelations: concepts, knowledge, reference and structure. *Mind and Language* 19: 85–98.
- Peirce, C. S. (1931–1935) *Collected Papers*, vol. II. Cambridge, Mass.: Harvard University Press.
- Pelletier, F. J. (1977) Or. *Theoretical Linguistics* 4: 61–74.
- Prinz, J. & A. Clark (2004) Putting concepts to work: Some thoughts for the twenty-first century. *Mind and Language* 19: 57–69.

- Putnam, H. (1975) The meaning of “meaning”. In K. Gunderson (ed.) *Minnesota Studies in the Philosophy of Science, Volume VII: Language, Mind, and Knowledge*, 131–193. Minneapolis: University of Minnesota Press.
- Putnam, H. (1983) ‘Two dogmas’ revisited. In *Realism and Reason (Philosophical Papers, vol. 3)*, 87–97. Cambridge: Cambridge University Press.
- Rips, L. (1983) Cognitive processes in propositional reasoning. *Psychological Review* 90: 38–71.
- Sosa, D. (2001) Rigidity in the scope of Russell’s theory. *Noûs* 35: 1–38.
- Sperber, D., F. Cara & V. Girotto (1995) Relevance theory explains the selection task. *Cognition* 52: 3–22.
- Sperber, D. & D. Wilson (1995) *Relevance*, 2nd edition. Oxford: Blackwell.
- Stalnaker, R. (1998) What might nonconceptual content be? *Philosophical Issues* 9: 339–352.
- Stanley, J. (1997) Names and rigid designation. In B. Hale & C. Wright (eds) *A Companion to the Philosophy of Language*, 555–585. Oxford: Blackwell.
- Wason, P. C. (1966) Reasoning. In B. M. Foss (ed.) *New Horizons in Psychology*, vol. 1, 135–151. Harmondsworth: Penguin.
- Woodfield, A. (1986) Two categories of content. *Mind and Language* 1: 319–354.