

INTENTION, MEANING AND TRUTH-CONDITIONS

(Received 17 April, 1978)

In an interesting paper called 'The meaning-nominalist strategy', Jonathan Bennett¹ attempts to show how the conventional meaning of an utterance-type can be explained by appeal to what a speaker means by an utterance. He calls his strategy a nominalist strategy because he thinks that certain ideas of Grice's² can be used in turn to explain what a speaker means by an utterance — what I shall call speaker meaning — without unexplained appeals to conventions or rules or anything similar. In a sense, then, Bennett wants to turn Grice upside down, for he wants to explain non-natural meaning by appeal to a kind of natural meaning.

Here is the plot. First we explain along Grician lines what I shall call pre-conventional meaning: cases in which an utterer *U* communicates successfully with an audience *A* in a way which doesn't depend on the existence of any rules or conventions. David Lewis' account of conventions³ is then imported to show how pre-conventional cases could lead to the establishment of a convention between *U* and *A* with the result that *U*'s utterance-type comes to have a conventional meaning. I shall begin by briefly expounding this plot, and then I shall re-write it in the light of a certain very fundamental criticism which I shall make.

I

A. The pre-conventional case

Bennett takes from Grice the following conditional.

- (GC) If *U* utters *e* intending thereby to get *A* to believe that *P*, and relies for the achievement of this upon the Grician Mechanism (GM), then *U* means by *e* that *P*.

Here is what we shall understand by the Grician Mechanism:

(GM) *A* recognizes *U*'s intention to get *A* to believe that *P*, and is led by that recognition – through trust in *U* – to believe that *P*.

This is a somewhat simplified version of Grice's most recent account, but we needn't worry too much about sophisticated qualifications here, for we want only a sufficient condition, and our main application will be to a very primitive case of communication. Bennett claims that this condition could be satisfied by *pre-linguistic U* and *A*, i.e., *U* and *A* who have no convention-governed means of communication at their disposal at all, leave alone a shared means of convention-governed communication.⁴ Bennett imagines the following sort of case.

Suppose we observe the following: Utterer performs a crude enactment of a man being hit on the head by a heavy object, and Audience sees this and steps out from under the palm tree just before a coconut falls on the place where he was standing. (Bennett, p. 147).

What Bennett needs to show is that in a case like this we could reasonably credit *A* with the belief that *U* intends to get *A* to believe that he risks being hit on the head by a coconut. He tries to make this plausible by breaking the matter up into easy steps, as follows: *U* intended by his performance

- (a) to achieve something;
- (b) to change his environment;
- (c) to change Audience;
- (d) to change Audience either visually or in some way arising from the visual change;
- (e) to get Audience to believe something;
- (f) to get Audience to believe something about someone's being hit on the head;
- (g) to get Audience to believe that he risks being hit on the head by a coconut.

(Bennett, p. 147)

What matters, as Bennett says, is just that *A* believes *U* intended (g). The point is not that *A* reasons through these steps. The point is rather to see how behavioral data might license attribution of the crucial belief to *A*.

Bennett says,

...it isn't hard to see what might justify the shift from (b) to (c): Audience can see that Utterer can see that his performance won't get trees felled or fires lit or fish caught, and that in general it hasn't a hope of affecting – in a manner answering to Utterer's interests – anything except Audience himself. Similar considerations underlie (a), and

each move from (a) through (d). Nor is there any special problem about the move from (d) to (e), except for those who object *in principle* to non-linguistic beliefs about intentions to produce beliefs. The shift from (e) to (f), obviously, depends on the fact that Utterer has acted in a way which could serve as a natural reminder or pointer (for Audience) to the idea of someone's being hit on the head. The final move from (f) to (g) rests on Audience's knowing that the only belief about someone's being hot on the head which Utterer could want Audience to acquire at that time would be the belief that Audience himself was in danger of being hit on the head. In the circumstances there just isn't anything else about head-hitting that Utterer could think it worthwhile to try to get Audience to believe.
(pp. 147–8)

There are many problems with this. As Bennett notes, there are those who would object in principle to 'pre-linguistic' beliefs about intentions to produce beliefs. I will return to this briefly in Section III. Also, there is a general problem about attributing mental dispositions on behavioral grounds which derives from the fact that radically different mental dispositions can have the same behavioral manifestations. For instance, someone who prefers vanilla to chocolate but prefers to disguise his preferences might well exhibit the same choice behavior as someone who simply prefers chocolate. Bennett avoids this line of objection by introducing a methodological principle: other things being equal, the less sophisticated attribution is always to be preferred to the more sophisticated attribution. This plausibly allows him to restrict his attention to showing that the data could be such as to allow no simpler attribution than the one satisfying (GC) (see p. 148). In any case, I propose to swallow this whole for the time being, for I have a different ax to grind.

B. *Conventions*

The second stage in Bennett's strategy is to import David Lewis' account of conventions to show how a convention could emerge between *U* and *A* concerning *U*'s utterance. Here is a simple illustration of convention à la Lewis. Suppose you and I are talking on the telephone, and the connection is broken. Assuming we each want to re-establish the connection, what should we do? If we both hang up, and one dials while the other waits, we will succeed. But if we both dial, or both wait, we will fail. Perhaps we shall fail, or perhaps we shall succeed by blind luck. Suppose the latter: perhaps I am distracted by a student, and you dial. Now in future if the same thing happens, I may reason as follows: last time, I waited and you dialed. Further, this is the only information bearing on the problem which we share. Hence,

doing as I did before — waiting — gives the best available chance of success. It may not be very good, but it is better than acting at random. (If the previous case occurred just five minutes ago, it is a pretty good bet.) You reason analogously, and dial. Result: we succeed. After several such successes it will become virtually certain that I will wait while you dial. We have a convention governing this situation. In general, when a group achieves coordination in a certain situation by acting in a certain way, and they act that way because each actor knows, and knows that the others know,⁵ that that is how coordination has been achieved in the past, then the group has a convention governing that situation.

As it stands, this analysis applies to a case involving coordination of action, whereas our target involves a kind of coordination between *U*'s action and *A*'s beliefs. But the account is easily generalized, for adopting a belief is, in the relevant sense, something *A* does. It is something he does in the relevant sense because *A* can have as his reason for adopting the belief that *U* intends him to believe that *P* in uttering *e* the fact that he knows, and knows that *U* knows, that in the past *U*'s intention in uttering *e* has been to get *A* to believe that *P*. If *A* is then led, through trust in *U*, to believe that *P*, we have a case which comes to satisfy the Gricean Conditional because *U*'s utterance of *e* is governed by a convention existing between *U* and *A*. This yields the following account of conventional meaning.

- (CM) Utterance-type *E* conventionally means that *P* when uttered by *U* to an audience *A* if (a) in the past *U* has uttered tokens of *E* to *A* only when he meant that *P*, and (b) this fact is mutually known to *U* and *A*, and (c) because of this mutual knowledge it continues to be the case that when *U* utters tokens of *E* he means and is understood to mean that *P*.⁶

We can now put the pre-conventional case and the conventional case together in an obvious way. Suppose *U* intends to get *A* to believe that a coconut is about to fall on *A*, and *U* goes through a certain performance which results in *A* recognizing *U*'s intention and, via trust in *U*, adopting the belief that he is about to be hit on the head by a coconut. Here we have a pre-conventional case in which communication occurs only because conditions are especially favorable and because *U*'s performance has a certain natural suggestiveness. Next time, however, the mechanism of convention will set in. As repetitions occur, special conditions will no longer be neces-

sary, and *U*'s performance can be streamlined, by a process akin to stimulus substitution, to the point where it need have no natural suggestiveness at all beyond the fact that *A* and *U* perceive it as 'of the same type' as its ancestors. The account thus allows for the obvious but theoretically crucial fact that any sign may, as far as its physical characteristics go, have any meaning whatever.

This concludes my brief sketch of Bennett's meaning-nominalist strategy. But before turning to criticism and reconstruction I want to take note of an ingenious argument which Bennett uses to block the possibility of natural language counter-examples to (GC). A counter-example to (GC) would involve *U* uttering an expression having *Q* as conventional meaning while satisfying the Grician condition for meaning *P*. Bennett upstages such cases as follows. First, he borrows Donnellan's account of why one cannot mean anything one likes by what one says. The crucial point is that to intend to achieve *T* by doing *x* one must believe that doing *x* has an appreciable chance of achieving *T*. Hence, for *U* to utter *e* intending thereby to get *A* to believe that *P*, *U* must think there is an appreciable chance that his utterance of *e* will lead *A* to believe that *P*. Hence, *U* cannot *mean P* in uttering *e* unless he thinks there is an appreciable chance his utterance of *e* will lead *A* to believe that *P*. Free choice of meaning requires free choice of intention, but free choice of intention requires free choice of belief, which we do not have. This point allows Bennett to give an account of why *U* is forced to mean by uttering *e* what *e* conventionally means (except in extraordinary circumstances). For if *e* has a conventional meaning, then there is mutual knowledge that speakers have meant that *P* by uttering *e* in the past, and this provides *A* with a powerful reason for thinking that in uttering *e* now, *U* intends *A* to believe that *P*. *U* knows this, hence cannot think there is a substantial chance of getting *A* to believe that *A* by uttering *e*, except in extraordinary circumstances. Hence, *U* cannot mean that *Q* by his utterance of *e* except in extraordinary circumstances. (Of course, special conventions may take precedence over more general conventions, as when husband and wife have a tacit convention whereby 'What time is it getting to be?' said at a party means, 'Let's go home'. But this simply shows that the conventions which are characteristic of English are not the only conventions governing English utterances.) If this is right, then it is no good looking to convention-governed utterances for counter-examples to (GC), and that is probably enough to guarantee that no knock-down counter-example will be forth-coming.

II

Can Bennett's meaning-nominalist strategy yield a theory of meaning? By a theory of meaning I mean a theory which (at least) assigns a meaning to each sentence⁷ of a language. Now it is obvious that there are infinitely many sentences in a natural language like English which have distinct meanings. Indeed, it is obvious that there are infinitely many *non-compound* sentences having distinct meanings. From this it follows that a theory which takes sentence-meaning as primitive hasn't a chance of being a theory of meaning in my sense for a natural language. Evidently, Bennett's strategy must fail to yield a theory in my sense, for it begins and ends with sentence-meaning. His strategy must treat the meaning-convention for each non-compound sentence as primitive,⁸ and it is therefore doomed from the start, for there cannot be infinitely many primitive meaning conventions.⁹ If we are to avoid this problem, we must find a way of introducing meaning-conventions for units of less than sentential size, for it is at least possible that the meaning of every sentence is determined by meaning-conventions governing a finite number of sub-sentential units. This is the possibility I wish to pursue in the remainder of this paper.

In his book *Speech Acts*, John Searle presents an account of predication which suggests that the primary intention of Utterer in predicating \emptyset is to raise the question of the truth or falsity of \emptyset of some identified thing.¹⁰ I propose to adapt this suggestion to present purposes by supposing that U 's primary intention is to get A to consider of some identified thing whether \emptyset is true or false of it. This gives us an intention to produce an effect in A which has some chance of forming the core of a Gricean account of what an utterer means by predicating \emptyset of some thing. Here is a First Pass at formulating a Gricean Conditional for these less-than-sentential meanings.

- (FP) If U identifies x for A and utters \emptyset intending thereby to get A to consider (of x) whether \emptyset is true or false of x , relying on the Gricean Mechanism, then U predicates \emptyset of x and means by \emptyset ...

Our immediate problem is: How are we to specify what U means by \emptyset ? If U 's intention in uttering \emptyset is to get A to consider whether x is red, then it seems that what U means by uttering \emptyset (on that occasion) is: *red*.¹¹ So the way to complete (FP) is by specifying in determinate from what U intends to get A to consider. The whole point of considering predication as opposed to

statement making or the like is that there is some hope that we might be able to do this for each predicate, there being, we hope, only finitely many. More precisely, we hope that specifying the meaning of each of a finite number of predicates will suffice as a base from which we can build up the meanings of any given sentence. Here, then, is a Second Pass.

- (SP) If U identifies x for A and utters \emptyset intending thereby to get A to consider (of x) whether x is red, relying on the Grician Mechanism, then, on that occasion, U predicates \emptyset of x and means 'red' by \emptyset .

We are to imagine that there is a statement like this for each of a finite number of predicates. Notice that when we specify what U intends in determinate form we eliminate the semantic predicate 'true of' from our formulations.

(SP) contains the phrase ' U identifies x for A '. This phrase must be explained if clauses like (SP) are to form the basis of a theory of meaning which is nominalistic in Bennett's sense. Since (SP) states only a sufficient condition, we need only concern ourselves with formulating a sufficient condition for identification.

- (ID) If U intends to get A to believe that U intends to predicate something of x , and if A recognizes this intention (i.e., comes to believe of x that U intends to predicate something of x), then U identifies x for A .¹²

The idea underlying this formulation is familiar: the point of identifying something for an audience is to get the audience to know what one intends to talk about.

Unfortunately, this formulation makes use of the expression 'predicate something of x ' and this is part of what (SP) is supposed to explain. It appears that an expansion of (SP) based on (ID) will be circular. Appearances are misleading, however. (SP) requires that U intend (i) to get A to consider whether x is red, and, further, that U rely upon GM to achieve this end. Now for U to rely on GM to achieve this end is for U to intend to get A to recognize intention (i), i.e., for U to intend (ii) to get A to believe that U intends to get A to consider whether x is red. Intention (ii) is just the primary intention required of U by (ID). So the only thing the contemplated expansion would add to the condition expressed by the antecedent of (SP) is that A actually recognizes the intention (i), i.e., that A actually comes to believe

that U intends to get A to consider whether x is red. Here, then, is what we have after performing the expansion and eliminating redundancies.

- (TP) If U utters \emptyset intending thereby to get A to consider (of x) whether x is red, relying for this on GM, and if A recognizes via (GM) U 's intension to get A to consider (of x) whether x is red, then U identifies x for A , predicates \emptyset of x , and means 'red' by \emptyset .¹³

If (TP) is to play a role analogous to the role of (GC) in Bennett's meaning-nominalist strategy, then it must be possible for U to have the required intentions, and for A to recognize them, in a pre-conventional context. Just how plausible this can be made depends a good deal on how we understand ' U intends to get A to consider whether Mars is red' and the like. We have a certain amount of freedom here, for expressions of the form 'consider whether x is \emptyset ' are evidently being used here in a technical sense: it is being used to express whatever it is that A *does* (or occurs in A) the doing of which (the occurrence of which) constitutes his *understanding* a predication, assuming that there is such a thing. Attempting to spell this out would be a project unto itself. I propose to simply set this whole matter to one side. Perhaps when we know in outline what sort of theory might emerge from our strategy, we will know better whether the fine details of its articulation and defense are worth further attention. The next move is to use (TP) and (ID) to give an account of what it is for a predicate or singular term to have a conventional meaning. This we do by assuming that for an expression to have m as its conventional meaning is for there to be a convention whereby one means m by uttering that expression, and then borrowing the Lewis-Bennett account of what it is for there to be such a convention.

- (ST) There is a convention whereby N refers to x in U 's language if (a) in the past U has uttered N only when he intended to identify x , and (b) this fact is mutually known by U and his hearers, and (c), because of this mutual knowledge it continues to happen that when U utters N he identifies x .
- (P) There is a convention whereby \emptyset means 'red' in U 's language if (a) in the past U has uttered \emptyset only when he meant 'red', and (b) this fact is mutually known to U and his hearers, and (c) because of this mutual knowledge it continues to happen that when U utters \emptyset he means (and is understood to mean) 'red'.

Actually, once the relevant conventions get going, it seems possible to tone down the meaning of 'means' in (P) in accordance with Searle's observation that *U*'s primary intention is typically just to get *A* to 'understand'. The idea would be to allow that *U*'s primary intention is simply that *A* recognize that the utterance is subject to a certain convention. There are tricky details here, however, and given the weak nature of the intention required by (TP) in comparison to the original Grice-Bennett account, the objection may not be worth meeting.

So far, we have only enough machinery to provide meanings for utterances having the first order from $\emptyset a$ where the meaning-conventions for each schematized component are primitive. What we need is a way of cooking up the meaning(s) of complex expressions given meaning-conventions for their constituents. This is accomplished in two phases. First we assert a relationship between conventional meaning and satisfaction conditions, and second we import Tarski's technique for combining satisfaction conditions.

- (S) 'The *i*th member of the sequence *s* is red' gives the satisfaction condition for a token consisting of \emptyset applied to¹⁴ the *i*th variable JUST IN CASE the (or a) conventional meaning of that token of \emptyset is 'red'.¹⁵

(S) allows us to go back and forth between satisfaction conditions and conventional meaning. If we start with cases for which a convention exists *via* (P), we get satisfaction conditions for the primitive predicates (by moving from the meaning part to the satisfaction part). We can then use Tarski's method to get satisfaction conditions for any first order combination of the primitive predicates. Using (S) again (moving from the satisfaction part to the meaning part), we wind up with conventional meanings) though not with meaning-conventions – for complex predicates.

In the case of closed sentences, our meanings will be, as Davidson claimed,¹⁶ truth conditions.¹⁷ This account thus has the consequence that to know a language is (in part, anyway) to know the sort of information codified in a theory of truth, and this would motivate the Davidsonian constraint on logical form – assign no logical form that requires an infinity of terms to be treated as semantically primitive – a constraint I criticized as unmotivated in "Truth and logical form."¹⁸ There I distinguished between (i) logical form: a canonical description that specifies whatever properties are relevant to the purposes of logic and truth-theory, and (ii) 'u-form': a can-

noncial description that specifies whatever properties are relevant to an account of how an expression is learned and understood. My claim was that there was no particular reason to suppose that logical form and *u*-form must (at some level of abstraction) coincide. The present approach would provide such a reason. What (TP) does is build a logical form into the specification of the propositional attitudes required by a Grician analysis. The point of this is to allow for conventional meanings attaching to sub-sentential expressions. But if there is to be a convention between *U* and *A* governing the use of \emptyset , then both *U* and *A* must recognize \emptyset as a constituent of the utterances in which it is instanced. For a convention between *U* and *A* governing the use of \emptyset requires that *U* and *A* reason about the use of \emptyset .

It might, of course, prove possible to work out an account like mine which takes something other than the terms of truth-theory as the significant constituents. I can't imagine how to do this because I can't see how to fit propositional attitudes to anything but sentences or sentential functions. But perhaps I have a weak imagination. If such an account were worked out, the combinatorial machinery of truth-theory would have to be replaced in (S) by something else — perhaps that of transformational grammar. From the point of view of an account like *that*, the relevance of truth-theory to learning and understanding a language would be coincidental at best.

This concludes my sketch of a meaning-nominalist strategy which deals in sub-sentential meanings. The project has been to re-write Bennett's recipe so that the basic ingredients are pre-conventional Grician meanings of a sort which might have some hope of forming a finite basis for a theory which is capable of assigning a meaning to any sentence of the language. The rewritten recipe is fundamentally unclear at many points: I have been high-handed with objections, and there are doubtless many important points which simply escaped notice. Still, there are a number of interesting things we can say about the result.

Notice first that the commitment to an infinity of conventions has been avoided, an essential point if conventions are anything like what Lewis says they are. (S) gives us indefinitely many meanings but only as many conventions as there are primitive terms.

Notice second that this account entails a version of the doctrine that predication and identification cannot be independently understood. It allows for predication without identification, but not in the pre-conventional case, and it is that case which explains the others.

Notice third that, given Bennett's point concerning the constraint the conventional meaning of an expression places on what one can mean by it, the account sketched here entails that one can only mean something simple when flouting all existing conventions, for the account provides no room for complex pre-conventional meanings.

Notice finally that (S) appears to provide an explanation of satisfaction of just the sort I called for in 'The philosophical problem of truth-of'.¹⁹ There I suggested that the or a central problem posed by Tarskian semantics is that we don't know what sort of facts the facts codified in satisfaction theory are. Satisfaction theory à la Tarski tells us what satisfies what, but leaves us without a clue as to how the fact expressed in a statement of satisfaction conditions should or could be explained. Now that (S) tells us is that a statement of satisfaction conditions for a simple predicate expresses the fact that a certain meaning convention exists, and *that* fact we explain in turn by appeal to Lewis' account of conventions and a meaning-nominalist construction modeled after Bennett.

III

This whole approach is fundamentally wrong-headed if the attribution of propositional attitudes only makes sense in connection with beings that speak a sufficiently rich language. If²⁰ we cannot explain '*A* believes that *P*' without appeal to facts about *A*'s language or linguistic dispositions, then formulations like (GC) are implicitly circular. Bennett spends a good deal of space in *Linguistic Behavior* (chapters 2–4) trying to show how non-linguistic behavior could reasonably lead us to a teleological account attributing beliefs and intentions. I think Bennett is on the right track: *A* believes that *P* iff a certain sort of functional analysis of *A*'s capacities is true.²¹ The only question is whether the right sort of account could be true of a being that knows no language. If genuinely non-linguistic behavior can really justify such an analysis, then the matter is settled. Such a thing is hard to prove, but a reasonable story should convince reasonable opposition.

The trouble is that some of the opposition is not likely to be reasonable in this case because the independence of propositional attitudes from language that is in question seems incompatible with a genuine truth embedded in the opposition position. This truth, as Fodor has seen²², is that attribution of propositional attitudes to *A* does presuppose that *A* 'has a

language' in a sense: the sort of analysis of *A*'s capacities that must hold if *A* is to have propositional attitudes is an analysis that attributes to *A* a structure capable of assuming states interpretable as representations of sentences in a fairly rich language, i.e., a structure the states of which are tokens of the sentences in Fodor's 'language of thought'. So this much of the objection is right: where there is no principled way of mapping mental states onto sentences of a non-trivial language, there are no propositional attitudes. What the objector does not see, and what Bennett and Fodor do see, is that a structure admitting such a mapping is required to explain a mass of capacities unquestionably had by beings that do not speak a language, the most notable among these being the capacity to learn a language. Mainly responsible for obscuring this point is the idea that there can be no 'language of thought' because there are no little men in our heads to speak it. Again, this is quite right in a way: the 'language of thought' is not literally a language, as Fodor is careful to point out (pp. 78–9). But this is quite beside the point. The argument concerns the kind and richness of internal states our account of non-linguistic capacities will force us to attribute to people. Everyone seems more or less agreed that the kind and richness in question is that exhibited by a natural language. But to suppose that the states of an organism are mappable in a principled way onto a natural language does not involve the assumption that people have the cognitive capacities they do because a little man or a part of the brain uses/understands the internal states in question as expressions in a language.²³ In short, talk of a 'language of thought' is (here, at least) merely a way of specifying indirectly the kind and degree of complexity an adequate functional analysis of non-linguistic human capacities will exhibit. Unquestionably, an analysis of such complexity would be difficult (though perhaps not impossible) to justify if our subjects could not speak to us. But even if linguistic evidence were the only evidence, it would take a very crude verificationism to allow us to conclude that our analysis must after all 'really' be an analysis of linguistic capacities.

IV

Grician accounts explain '*U* means *m* by *e*' by appeal to propositional attitudes. Such accounts are at bottom Lockean accounts of communica-

tion.²⁴ *U* communicates with *A* when *U* produces the appropriate mental state in *A*. The linguistic utterance is the (or a) means by which this is accomplished. Since *A* has understood *U* (*U* has communicated with *A*) when *A* knows what *U* meant by his utterance, it is inevitable from this point of view to think of a theory of meaning as a theory that (i) says what each expression characteristically communicates, i.e., distributes meanings over expressions by specifying what state *A* must get into to count as having understood *U*'s utterance, and (ii) explains what it is for expressions to have the meanings they have. This is evidently what an account of the kind suggested above is supposed to do. But this sort of account leaves several traditional questions about meaning untouched, e.g., what is it to understand (mean) *red* as opposed to *green*?

It seems to me that there are several usefully distinguished components to a philosophical theory of meaning.

1. Recursive semantics: a theory that specifies how the semantic properties of complex utterances depend on the semantic properties of their parts.

2. Pragmatics: a theory that explains how contextual factors co-operate with the semantic properties of an expression to yield the semantic properties of an actual utterance.

3. An analysis of meaning: a theory that explains what it is for an utterance to have the semantic properties assigned to it by 1 and 2.

4. An analysis of *understanding*: a theory that explains what it is for an audience to understand an utterance or an expression.

None of these things taken by itself is likely to be very satisfying. For instance, the type-3 account developed in this paper helps itself to the likes of *considering whether x is red* without explaining what this is or how it differs from *considering whether x is green*.²⁵ A likely complaint is that we have at best mapped the utterances of public language onto the language of thought, thus repeating the familiar mistake of substituting an obscure problem about the 'semantics' of the 'language of thought' for the clearer problem about the semantics of the public language. But I am inclined to agree with Locke and Fodor in thinking that this is not a mistake but an inevitable step toward solving the problem about the public language. There is an obvious sense in which the words ought to drop out of an account of the difference between the meaning of 'Hesperus' and the meaning of 'Phosphorus'.²⁶ It is no more likely that a theory about this will turn out to be about the words than that

a theory about the difference between greek and italian pizza will turn out to be about Greeks and Italians. No doubt we will have to observe the words (Greeks and Italians) to find out what we want to know, but this needn't lead us to think we are theorizing about the words (Greeks and Italians).

Similarly, Tarskian semantics (type-1) fails to satisfy because it tells us at most what satisfies what without telling us what satisfaction is. And speech act theories such as Searle's seem hollow because they help themselves to 'propositional meanings', i.e., to the semantic properties of the expressions of the language.

The cure for this sort of dissatisfaction lies in seeing how the four components described might mesh. A great virtue of the old conceptualist approach is that it allows us to separate questions about the language from questions about communication and from questions about understanding. The price we pay is that we are tempted to criticize answers to the separate questions as if they were intended to do the whole job. Some people do hammer with a wrench, but this is no fault in the wrench. Another price we pay is that we are forced to allow mental states generally, and propositional attitudes in particular, to have a life of their own independent of the public language and its use. But if Fodor is right, and I think he is, we are stuck with this anyway.

University of Wisconsin-Milwaukee

NOTES

¹ Jonathan Bennett, 'The meaning nominalist strategy', *Foundations of Language* 10 (1973), pp. 141–168 (hereafter, simply 'Bennett' followed by a page number). Also: *Linguistic Behavior* (Cambridge University Press, Cambridge, 1976).

² H. P. Grice, 'Meaning', *Philosophical Review* 66 (1957), 377–388; 'Utterer's meaning and intentions', *Philosophical Review* 78 (1969), pp. 147–177.

³ David Lewis, *Convention* (Cambridge University Press, Cambridge, 1969).

⁴ In its application to the pre-conventional case, we needn't put too much emphasis on the fact that (GC) purports to state a sufficient condition for *meaning*. So far as I can see, nothing important would be lost if, in considering the pre-conventional case, we substituted the following for (GC): If *U* utters *e* intending thereby to get *A* to believe that *P*, relying for the achievement of this on GM, and if *U* achieves this via GM, then *U* communicates the belief that *P* to *A*.

⁵ See Bennett's treatment of 'mutual knowledge': (Bennett, p. 150) for a precise formulation of this epistemic condition.

⁶ This is a slightly re-formulated version of something Bennett claims to borrow from Schiffer (see Bennett, p. 152, and Stephen Schiffer: *Meaning* (Oxford University Press, Oxford, 1972).

⁷ Throughout this discussion, 'sentence' is to be understood to mean 'sentence type'.

⁸ The meaning convention for an expression *e* is primitive just in case the conventions determining the meanings of utterances not containing *e* do not suffice to determine the meanings of utterances containing *e*.

⁹ The strategy *could* yield a theory of meaningfulness, and it is interesting that Grice began in a way which made this the primary target.

In Linguistic Behavior (Chapters 8–9), Bennett does discuss at some length how we as theorists could come to know what *U*'s terms mean on the basis of knowledge of what *U*'s sentences mean. This is no doubt possible. But we could not find out what *U*'s terms mean unless *U*'s terms had meanings. As far as I can tell, Bennett's discussion of how we might discover what *U*'s terms mean does not provide any hint as to how we should understand '*U* means such-and-such by \emptyset '. We therefore require an analogue of (GC) for terms.

¹⁰ John Searle, *Speech Acts* (Cambridge University Press, Cambridge, 1969), pp. 121–127.

¹¹ Of course, *U* might mean more than this, for *U* might also intend by the same utterance to consider whether *x* is scarlet: you say, 'He didn't wear a red shirt'. I reply, 'His shirt was scarlet!' Both intentions seem to underlie the reply. This is not a problem, though, for nothing like (FP) could endorse a move from what *U* *does* mean to a conclusion about what *U* *doesn't* mean.

¹² The Grician Mechanism is conspicuously absent from (ID). This is because, by my lights, identification is an act of communication which is essentially perlocutionary in nature: identification needn't involve *meaning* anything, though, of course, it often does. The point needn't be insisted on, however. Introducing (GM) into (ID) would complicate matters a good deal but would not, so far as I can tell, effect subsequent developments in any essential way.

¹³ The schematic '*x*' here yields, case by case, to a metalinguistic specification of what *U* identifies. So far as I can tell, there would actually be no *harm* in treating '*x*' as a genuine variable to be bound by a universal quantifier. To emphasize the transparent reading, we could reformulate using the constructions 'intends of *x*...*x*...' and 'considers of *x*...*x*...'.

¹⁴ 'applied to': this is to be understood as a syntactic predicate linking a predicate and its argument(s), e.g., linking the predicate and the singular terms or variables that fill its places.

¹⁵ The parallel formulations for closed atomic sentences is straightforward; e.g., (*s*) (*s* satisfies a token consisting of \emptyset applied to *N* iff Mars is red) JUST IN CASE the (or a) conventional meaning of *N* is 'Mars' and the (or a) conventional meaning of \emptyset is 'red'.

Why not write (S) thus: *s* satisfies \emptyset applied to the *i*th variable iff *s_i* is red JUST IN CASE the conventional meaning of \emptyset is 'red'? Because we could move from

s satisfies 'cordate' applied to the first variable iff *s₁* is renate

which is true, to

the conventional meaning of 'cordate' is 'renate',

which is false. (I owe this point to Philip Hugly.)

¹⁶ Donald Davidson, 'Truth and meaning', *Synthese* 17 (1967), pp. 304–323.

¹⁷ So far, we have been introducing, not meanings, but the meanings of the propositions expressed in illocutions. To get meaning proper, we must add illocutionary forces. Since Searle's account presupposes propositional meanings and builds illocutionary acts on to these, advocates of the account above can do as well by illocutions as Searle. The sympathetic reader will want to hold that there are Lewis-type conventions for taking an utterance having a specified propositional meaning as having one or another illocutionary force, depending on the satisfaction of Searlian conditions.

On the face of it, the approach just described will apply only to artificial languages, for (S) requires a specification of the logical form (a Tarskian structural description, in fact) of *e* to work on. To be of any help in connection with a natural language, therefore, this approach awaits progress toward completion of a Davidson-style truth-characterization for that language. (S) would actually appeal to such a theory at two points: (1) to determine the logical form to be associated with particular utterances of a given expression (given speaker, time and whatever other contextual factors are involved), and (2) to provide the satisfaction clauses for the terms discriminated by the associated logical form.

¹⁸ The Journal of Philosophical Logic 4 (1975), pp. 29–44.

¹⁹ The Canadian Journal of Philosophy 5 (1975), pp. 103–122.

²⁰ As Davidson alleges. See his 'Thought and talk', in Samuel Guttenplan (ed.), *Mind and Language* (Clarendon Press, Oxford, 1975), pp. 7–23.

²¹ For my account of functional analysis, see 'Functional analysis', *The Journal of Philosophy* 72 (1975), pp. 741–65.

²² J. A. Fodor, *The Language of Thought* (Crowell, New York, 1975), especially Chapter 2.

²³ For more on this sort of point, see my 'Programs in the explanation of behavior', *Philosophy of Science* 44 (1977), pp. 269–87.

²⁴ As Fodor sees: having described a Lockean (translation) account of communication, Fodor notes the connection with Grice on pp. 103–4, footnote 3.

²⁵ A theory like Field's ('Logic, meaning and conceptual role', *The Journal of Philosophy* 74 (1977), pp. 370–409) can help here by explaining what facts about *A* are constitutive of meaning/understanding *green* rather than *red*. Field suggests that we specify the meaning of *t* by specifying reference and conceptual role, where the latter is given by the conditional subjective probabilities of the sentences (in an extended language) containing *t*.

²⁶ See Field's account of this difference, *op. cit.*, pp. 390–1.