

AS



Functional Role and Truth Conditions

Author(s): Ned Block and John Campell

Source: *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 61 (1987), pp. 157-181+183

Published by: Blackwell Publishing on behalf of The Aristotelian Society

Stable URL: <http://www.jstor.org/stable/4106836>

Accessed: 07/02/2009 18:07

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



*The Aristotelian Society and Blackwell Publishing are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society, Supplementary Volumes*.*

<http://www.jstor.org>

FUNCTIONAL ROLE AND TRUTH CONDITIONS

Ned Block and John Campell

I—Ned Block

I

This paper is about truth-conditional content from the point of view of a use theory of meaning. I shall be taking up the question of how a use theory of meaning should deal with cases where uses of words might be said to be the same, but truth conditions differ—specifically Putnam’s ‘twin earth’ cases.¹ The background issue is whether there is a legitimate notion of meaning or content that is purely ‘inside the head’—not dependent on the external context of things referred to or on the language community. This paper defends the legitimacy of this ‘narrow content’ conception in the following limited way: I argue that use theories of meaning must deal with truth conditions in a way that leaves room for narrow content.

By ‘use theory of meaning’, I mean the type of theory that is known in philosophy as functional role semantics or conceptual role semantics, and in psychology and artificial intelligence as procedural semantics. The functional role approach has much to be said for it, and also much to be said against it. I will mention just a few of the things to be said for it (partly to give a better picture of what the theory says), and none of the things to be said against it.²

In a broad sense of ‘use’, functional role semantic theories agree that meaning is use. One approach takes use to include public use, especially in communication, whereas another excludes public use, restricting itself to private use, especially in thought. (Here ‘function’ is more appropriate than ‘use’.) The

¹ Hilary Putnam, “The Meaning of ‘Meaning’”, in *Language, Mind and Knowledge*, Keith Gunderson (ed.). University of Minnesota Press: Minneapolis, 1975. The twin earth case will be briefly described later.

² Some of the things to be said against it can be found in Jerry Fodor’s *Psychosemantics*, Bradford Books, MIT Press: Cambridge, 1987, Ch 3, and in Fodor’s ‘Banish disContent’, in *Language, Mind and Logic*, Jeremy Butterfield, ed., Oxford University Press: Oxford, 1986.

former tack leads to functional role semantics as a theory of truth conditional meaning and content, whereas the latter leads to a theory of 'narrow' meaning and content ('inside the head'). Offhand, it would seem that only the latter tack could provide a place for narrow meaning or content. I will argue that—contrary to initial appearances—it makes little difference whether the functional role semanticist focuses on public or private use: the two projects result in equivalent theories. Further, the theories resulting from the two projects are equivalent to a 'two factor' theory, a theory that has one component for functional role inside the head, and another component relating words to the world. Since such a theory has a component for functional role inside the head, there is a natural place for narrow meaning and content within it.

Here is my plan: I will try to explain and motivate functional role semantics by arguing that it can explain five relatively uncontroversial facts. Then, after a discussion of private versus public use, I will consider how a functional role theory can capture truth conditions, arguing that the ways of so doing appropriate to public use theories and private use theories converge. But first some preliminary points.

The functional role approach is very vague, with many possible roads of development. With the exception of the public/private difference (and also differences in degree of commitment to reductionism), I shall not be concerned with differences among the different approaches, exploring what the most plausible versions have in common, rather than any specific version. I should admit at the outset that since the theory is so vague, it does not deserve any honorific associations that may come with the word 'theory', as opposed to 'approach' or 'direction'.

Functional role semantics is not really a competitor of the various approaches to formal semantics, or of Katzian semantics. These approaches are concerned with explaining particular phenomena of particular languages, e.g., with why 'The temperature is rising' and 'The temperature is 90°' does not entail '90° is rising'. Functional role semantics, by contrast, aims to say what meaning is, and to cast light on issues that depend on what meaning is. For example, reductionist versions of functional role semantics attempt to solve the analogue of the mind-body

problem for meaning, attempting to exhibit the place of meaning in the physical world. One way of seeing the contrast is to note that many of the approaches mentioned have attracted linguists who can subscribe to and work within those paradigms. Functional role semantics (and its genuine competitors—see footnote 14), by contrast, does not provide a framework for work in linguistics.

The contrast is clouded by the fact that the formal semantics approaches tend to have answers to the question of what meaning is: for example, that the meaning of a sentence is a function from possible worlds to truth values. *This view is no competitor of the functional role approach, however, if it is in virtue of the sentence's functional role that it determines the function from worlds to truth values.* In general, when a formal semantics theory entails that the meaning of a sentence is X, I say that it is in virtue of the sentence's functional role that it means X. Some formal semanticists advocate a theory of meaning that does genuinely compete with functional role semantics, but such theories are not part and parcel of the formal semantics approach. For example, Barwise and Perry³ put forward a version of 'indicator semantics' that has also been advocated, independently of any formal semantics, by Dretske, Stampe, and Fodor.⁴

Semantics is usually understood as concerned with language, and if with thought, only derivatively. We talk more naturally of the meaning and reference of words than of concepts, and of the truth value of sentences than of states of mind.⁵

However, assuming, as I shall, that one thinks in a representational system, a theory of the meaning of elements of the representational system in which one thinks can do double

³ *Situations and Attitudes*, MIT Press: Cambridge, 1983.

⁴ See Fred Dretske, *Knowledge and the Flow of Information*, MIT Press: Cambridge, 1981; Dennis Stampe, 'Towards a Causal Theory of Linguistic Representation', in *Midwest Studies in Philosophy*, vol. 2: 42-63, 1977; Jerry Fodor, *Psychosemantics*, *op. cit.*

⁵ On the other hand it seems just as natural to speak of a thought being true, contradictory, or entailed by another as of a sentence. I believe these latter intuitions stem from the fact that there is a process/product or rather state/object ambiguity in words like 'thought' and 'belief', namely they can refer to states of minds on the one hand or the intentional objects of these states of mind on the other. In the intentional object sense, thoughts can be said to have truth value, etc.; but it is the state sense that raises the objections—at first blush, to speak of a state of mind as being true is what philosophers some years ago would have branded a category mistake (though, on the other hand, it is

duty as a theory of the intentional content of thought. (In talking of a representational system in which one thinks, I do not mean to assume anything like sentences in the brain, though I will sometimes speak this way. For my purposes, the representational system can consist in structured states.) Functional role semantics in the version that I favour—a version in which functional role is thought of as inside the head—identifies the meaning of a sentence with its role in thinking, problem solving, deliberating and the like—and, in general, in mediating between sensory inputs and behavioural outputs. (I'll get to the possibility of expanding functional role to include role in communication later.) To know the meaning of 'The cat is on the mat,' is for this expression to function appropriately in one's thinking, problem solving, etc.

Let us say that the inferential role of a sentence is a matter of what sentences can be inferred from it, and what it can be inferred from. Functional role includes inferential role plus more (e.g., causal relations to inputs and outputs), since the role of language in reasoning, problem solving, deliberating, etc., includes more than inference. Just what more I won't try to say. Corresponding to the functional role of a sentence we can speak of the functional role of the thought-state or event that consists in the thinking of that sentence. E.g., the inferential role of a state of thinking is a matter of what thought states would be (should be? ideally would be?) produced by it by a process of inference, and what thought states it would be produced by via inferential processes.⁶ I will travel back and forth between thought and language, between content and meaning, depending on the issue at hand.

I said that functional role includes inferential role plus more. Actually, breaking up functional role in this way involves a theoretical decision that a reductionist version of functional role

natural to speak of states of mind as about something). It is commonplace these days for philosophers to talk as if semantic properties apply equally directly to language and mind. I think it is important to remember that semantic properties apply in the first instance to language, and only on the basis of a controversial theory to mind.

⁶ Though I rely here on the empirical assumption that we do think in language, there is a case to be made that functional role semantics could make do with purely counterfactual roles—roles language would have if we did think in it.

semantics will reject. The reductionist version of functional role semantics requires a specification of functional role in non-intentional terms, so explicit talk of inference in any sense that involves propositional attitude talk is illegitimate. Note incidentally, that not even a functional role semanticist who allows intentional terms can afford to cash inference in semantic terms, e.g., in terms of the truth of one statement making the truth of another more probable, for the use of semantic terms in a semantic theory would threaten to make the theory uninteresting for purposes of explaining semantic notions. (Actually, if some semantic properties can be independently explained, they can be used to explain others; since I don't see how this could be done in a functional role context, I shall ignore this possibility.) If one wants to avoid talk of inference altogether, one can try to cash functional role in causal terms, that is in terms of what representations cause and are caused by what other representations, and also causal links to inputs and outputs. This would be the version of functional role semantics closest to traditional functionalism as a theory of mind. I will talk in terms of inferential role, but I don't want to preclude the purely causal option. Indeed, I've adopted the 'functional role' terminology (instead of the more common 'conceptual role') to emphasize the connection with functionalism as a theory of mind. I want to keep the causal option just mentioned alive, and emphasize the fact that the functional role of an internal representation includes not only its relations to other items of language, but also to sensory inputs and behavioural outputs.

Let me be a bit more specific about the issue of reductionism. One way a theory can fail to be reductionist is to explicitly use semantic notions. For example, suppose one offered the following use theory of the semantics of 'snow', namely that it is used to refer to snow. This would be an uninteresting theory, and one way to avoid it is to prohibit semantic notions in the specification of uses.

A second way a theory can fail to be reductionist is to make essential use of intentional content. Suppose, for example, that we give the meaning of 'Snow is white', in part, by saying that it can be used to infer that snow is not green. Since the intentional content (that snow is not green) is as much in need of explication as the original meaning, this approach does not get us very far.

Of course, one could try to eliminate all the contents/meanings at once via a Ramsey sentence, or some equivalent device, but this would involve a large and controversial further step. This step would lead to what I am calling a reductionist theory, one that makes no use of semantic or intentional notions. Perhaps an analogy will help. A functionalist in philosophy of mind can define 'disappointment' as the feeling one gets when one realizes one will not achieve what one wants. This might be useful for some purposes, whether or not the intentional notions in the definition are themselves functionally definable. A reductionist functionalist, however, will want to define all intentional states in terms of the abstract structure constituted by the relations that the states all have to one another and to inputs and outputs.⁷

The reductionist version of functional role semantics is far more ambitious and interesting than non-reductionist versions, but it is also much less likely to succeed.

II

Five Facts.

One motivation for functional role semantics lies in a cluster of facts about concepts and words that can be explained by the theory. I will mention five of these, in part to motivate the theory, and in part to give a bit more of an idea of what it is supposed to be. I hope that what the approach of functional role semantics comes to will become clearer as I say how it explains the facts that follow.

First, grasping some concepts involves grasping others. Could someone have the concept of Newtonian force without Newtonian mass? Or buying without selling? The fact that grasping some concepts requires grasping others is a fact in the realm of the mind; corresponding to it in the realm of language, we have the fact that knowing the meanings of some words requires knowing other meanings.

Secondly, having a thought involves the capability of

⁷ See p. 177 of my 'What is Functionalism?', in *Readings in Philosophy of Psychology*, Vol I, Harvard University Press; Cambridge, 1980: 171-184; and p. 263 of Sydney Shoemaker's 'Some Varieties of Functionalism', *Philosophical Topics*, 12, I, 1981: 83-118.

grasping a range of related possibilities (and understanding the meaning of a sentence requires a corresponding capability). A person would not have the thought that Reagan is in Washington if he could not grasp such ideas as that Reagan is elsewhere, that he will be in Washington later, and that someone else is in Washington. More generally, we have Evans' principle that if one has the ability to grasp that a thing has *F* he must have the ability to grasp that it has *G* (for every property *G* of which he has a conception), and that other things (of which he has a conception) have *F*.⁸ Moreover, it may be that in some cases, having one thought involves *actually* having another (not just the capability of having another). The most plausible cases are systems of thought such as scientific theories. This point underlies the observation that one comes to understand a system of linked statements not piecemeal, but together. As Wittgenstein said, 'When we first begin to believe anything, what we believe is not a single proposition, it is a whole system of propositions. (Light dawns gradually over the whole.)'⁹

Thirdly, as Kuhn has observed, modern scientists cannot understand the phlogiston theory without learning old concepts (elements of an old language).¹⁰ The reason for this is that the key terms of the phlogiston theory (e.g., 'phlogiston') cannot be simply translated into the language of modern science.

Fourthly, in translating a person's words, we cannot attribute irrationality without limit. As has often been noted, to the extent that the beliefs we attribute are incoherent, our translation hypotheses lose plausibility.

Finally, grasping a thought involves grasping its structure (grasping the meaning of a sentence involves grasping its structure)—the component ideas, and how the thought puts them together.

What follows is a sketch of the functional role account of each of the five facts. These accounts are of course controversial, and I don't pretend for a moment that these brief sketches settle the hash of competing accounts. The point of the sketches is to indicate that the functional role semantics approach has some

⁸ *The Varieties of Reference*, Oxford University Press: Oxford, 1982, p. 104.

⁹ *On Certainty*, Harper and Row: New York, 1972, p. 21e (# 141)

¹⁰ 'Commensurability, Comparability, Communicability', in *PSA 1982* Vol. 2, P. Asquith and T. Nickles (eds.), Philosophy of Science Association: Ann Arbor, 1983.

claim to explanatory power, and in so doing to give a better idea of what the approach comes to. I will start with the second commonplace mentioned above (because its explanation is more straightforward), and then go to the first.

(2) *Thoughts*

According to functional role semantics, for me to have a thought is for something (a representation) to have a certain functional role in me, that is, for it to be inferable in certain contexts from certain other representations with their own functional roles, and for other representations with their own functional roles to be inferable from it. But to be able to have these other representations with their own functional roles is to be able to have other thoughts corresponding to these other representations. Thus having one thought involves being able to have others.

(1) *Concepts*

The explanation of (1) is along the same basic lines as (2): To have concept X is to have a representation with a certain functional role; the functional role of this representation requires having other representations with their own functional roles; having one of these other representations constitutes having concept Y; so having concept X requires having concept Y. However, there is more room for differences between different versions of functional role semantics in explaining (1) than (2), for it is more obvious what the inferential roles of whole sentences or whole thoughts might be than what inferential roles of words or concepts might be. One position is that words only have functional roles to the extent that they contribute to the functional roles of sentences.¹¹

(3) *Incommensurability*

Functional role semantics give an account of why *phlogiston*

¹¹ Given this view, one might hold that what it is for a representation to have the functional role appropriate for expressing the concept of buying involves the presence of other conceptual representations that could appear in sentential representations inferable in some specifically simple way from sentential representations containing the representation of buying. These latter conceptual representations include one for selling, so grasping buying requires grasping selling.

can only be acquired by learning elements of a new way of talking, and not by translation. One way of putting Kuhn's point is that the functional roles of *phlogiston*, *principle*, etc., are very different from the functional roles of any concepts or complexes of concepts of modern physics, and hence there is no substitute for acquiring the 18th Century functional roles if we want to grasp the 18th Century ideas. Kuhn's claim is often rejected on the ground that we can define 'phlogiston' in our language via something like a Ramsey sentence, e.g., phlogiston can be identified with the referent of a description along the lines of 'The something that is a constituent whatnot of combustible materials and metals, that is given off in burning, and when a metal is converted to a calx, that renders air unbreatheable when saturated with it, that charcoal is especially rich in, . . .' ('Something' replaces 'phlogiston' in the original theory, and 'whatnot' replaces 'principle'.) But the functional role approach allows us to see that to incorporate this sort of Ramsey sentence into our language is just to incorporate the 18th Century functional roles—the Ramsey sentence is simply one common way of representing them.

(4) *Translation*

The fact to be explained is that in translation, we cannot rationally attribute irrationality without limit. The functional role semantics explanation is based on the idea that when we translate, we try to 'match' the functional roles of words in the alien's language with ours, and that accepting a poor match involves attributing bizarre thought processes to the alien. Suppose we have a translation hypothesis that leads to a terrible match of roles, one that, were it accepted, would require ascription of gross irrationality to the alien. There are two possibilities: another translation hypothesis will lead to a better match, or none will. In the former case, another translation hypothesis will lead to a better match, so we should certainly not accept the one we have. In the latter case, there is no minimally acceptable translation hypothesis, so we should accept none of them, concluding that the alien's conceptual system is not intelligible within

our own; there is no way of understanding the alien without acquiring his language together with the functional roles of his terms. In either case, it is irrational to accept the translation hypothesis that requires attribution of gross irrationality.

Note that the functional role semantics account does not support versions of the principle of charity advocated by Davidson (truth maximization) and Dennett (rationality maximization).¹²

(5) *Structure*

That *a* is *F* entails that something is *F*, that *a* has some property, and that something has some property. That a bachelor has *F*, entails that an unmarried person has *F*. Each concept in a structure and the structure itself are exploited by entailment relations, and these entailment relations lay bare the structure of the thought. Hence, one cannot appreciate a thought's inferential role without grasping its component concepts and how they fit together.¹³ Note that I am asserting no explanatory priority of functional role over structure. (I think neither has explanatory priority.) My point, rather, is that the functional nature of meaning explains why the state of grasping meaning involves grasping structure.

Even if functional role semantics can explain the five items just mentioned, that would not show that the doctrine is right, of course. Inference to the best explanation requires a comparison with other theories. But the facts just mentioned do shift the burden to the competitors.¹⁴ If the competitors cannot explain

¹²Note also that translation requires preservation of reference and truth value (indeed, actual practice often treats these matters as criterial). Hence, any notion of functional role that leaves out reference and truth value will inevitably leave out much of what is relevant to translation. I will be proposing a 'two factor' view, in which one factor does leave out reference and truth value, while the other factor includes reference and truth value; the two factors must be used together to account for translation.

¹³This needs to be qualified, since extraction of some 'component concepts' requires conceptual creativity. See p. 167, John Campbell, 'Conceptual Structure', in C. Travis (ed.), *Meaning and Interpretation*, Blackwells: Oxford, 1986, 159–174.

¹⁴What are the competitors? One competing family of theories sees meaning, not in terms of use, but as arising from other features of communication. See Campbell, *op. cit.*, for a critique of the communication-theoretic approach as a theory of the relation

these facts, functional role semantics has an advantage over them.

My aim so far has been to provide some reason for taking functional role semantics seriously. I will now explore how functional role semantics can be a theory of truth-conditional content.

III

Functional role and Truth-Conditional Content

Suppose the sentence 'Aluminum conducts electricity' has different truth conditions as used on earth and twin earth. (I assume it's the same sentence type uttered in both places.) On earth, 'aluminum' is used to refer to the element aluminum. On twin earth, what they call 'aluminum' is not an element at all, but rather an alloy composed of three different elements, none of which is known on earth. (Let's suppose this is the only relevant difference between the two language communities.) So the truth condition of 'Aluminum conducts electricity' depends on whether it is uttered by a speaker of the earthian dialect or the twin-earthian dialect. The earthian utterance is true iff aluminum conducts, the twin-earthian utterance is true iff the alloy conducts. (This claim is open to question, of course; I won't question it here.)

If a functional role semantics theory is to be a theory of truth conditional content, it must somehow reflect this difference between earthian and twin-earthian truth conditions. My

between grasp of conceptual structure and grasp of truth conditions. Another family of theories that—in some versions—is a competitor ('Gricean semantics') reduces meaning to mental content. Another ('indicator semantics'), sees meaning as co-variance: the paradigm of meaning is the relation between the number of rings on the tree stump and the age of the tree. Another (causal theory of denotation, plus, perhaps, an account of truth in terms of denotation of the sort advocated by Field) sees meaning as derived from the causal chains connecting words to the world. See pp. 652–661, 665 of my 'Advertisement for a Semantics for Psychology', in P. A. French, T. E. Uehling, and H. K. Wettstein, *Midwest Studies in Philosophy* Vol. X, University of Minnesota Press: Minneapolis, 1986: 615–678, for a discussion of Gricean and indicator semantics, and Michael Devitt, *Designation*, Columbia University Press: New York, 1981, Ch 3, 5, for a discussion of the causal theory. The last three doctrines can be held in versions that are compatible with or even equivalent to the 'two factor' functional role approach that I advocate.

strategy will be to start with the explanation of such *differences* in truth conditions rather than the explanation of truth conditions themselves. I will explore a number of options for developing functional role semantics so as to provide truth conditions for the sentences in the earthian dialect that are different from the truth conditions in the twin-earthian dialect.

To begin with, let us consider the possibility of doing this by putting the burden on the thoughts functionally related to 'aluminum'-thoughts. That is, if we have a thought, *t*, whose truth-conditional content is to be captured by functional relations (including inferential relations) to thoughts *a*, *b*, and *c*, perhaps this can be done by exploiting the truth-conditional content of *a*, *b*, and *c*. Of course, this strategy would involve a non-reductionist construal of functional role, one that explains the truth-conditional content of some items in terms of others. This project is certainly less interesting than one using only non-intentional resources, but let us persevere to see what can be achieved.

The picture behind this way of proceeding is the one behind a common way of looking at dictionary definitions. Suppose we don't know the referent of 'Mars'. Consulting the dictionary, we are told that Mars is the planet between Earth and Jupiter. We can then learn the referent of 'Mars', provided we know the referent of 'Earth' and 'Jupiter'. More generally, we use the dictionary to break out of the circle of words by knowing the semantic values of many of the words in the definitions already.

We infer (non-demonstratively, of course) from 'This is light, grey, does not tarnish, and conducts electricity moderately well, though less well than copper' to 'This is aluminum', and from 'This is aluminum' to 'This is metal'. Whatever inferential role is exactly, some such inferences will no doubt be involved in the inferential role (and therefore, the functional role) of 'aluminum'. But are there any *differences* between earthian and twin-earthian inferences in virtue of which the referent of 'aluminum' in earthian English differs from the referent of this word in twin-earthian English?¹⁵ Of course not, because the words 'light',

¹⁵ Recall that we are considering earth and twin-earth at a stage of inquiry at which no one has any knowledge that distinguishes aluminum and twin-aluminum. This could be the case because, though aluminum (and its twin) has been named, no one has examined

'grey', 'tarnishes', 'conducts', 'copper', and 'metal' have the *same* extension in the earthian language as in the twin earthian language.

If this is not obvious, note that the extension of a term includes all the things in the whole universe that it applies to, whether or not the speaker knows about them. Thus the extension of 'grey', for example, as used on both earth and twin earth is exactly the same, namely all the grey things in the universe, and this extension includes both aluminum and twin aluminum. The word 'aluminum' as used in the two places has different extensions because the extension is fixed by something that differs in the two language communities. Not so for 'grey', 'tarnishes', 'conducts', 'metal', etc. So we cannot distinguish the truth conditions of 'aluminum' utterances by appealing to a difference in the truth conditions of the non-'aluminum' utterances.

But what if 'aluminium' has a synonym in each of the two language communities? E.g., suppose both communities have a convention that you get a synonym if you spell a word backwards. Reply: Then the reasoning just given will apply to the pair, 'aluminum', 'munimula'. But what if there are *lots* of terms, all semantically related to 'aluminum', some of which can be used to specify the truth conditions of others? Only in this circumstance, can there be any semblance of success. The special circumstances of the success show its limitations.

Much more important, however is the manifest limitation of the non-reductionist approach itself. The non-reductionist approach explains truth-conditional properties well enough for some purposes, e.g., explaining the five facts mentioned earlier. But it is of no help at all in explaining the difference between the referential and truth-conditional properties of words on earth and twin earth. There is no point in explaining the difference in reference between 'aluminum'_{earth} and 'aluminum'_{twin-earth} in terms of the difference between 'munimula'_{earth} and 'munimula'_{twin-earth}, since the latter difference is every bit as much in need of explanation as the former one.

it very carefully. All the twin earth case requires is that people on earth and twin earth are in the *same* epistemic state with regard to *different* substances, not that the actual and counterfactual substances are *really* observationally indistinguishable.

Let's try another tack: give up on a pure functional role account, and adopt a so-called two factor theory.¹⁶ One factor, the functional role inside the head factor, is aimed at explaining facts such as the five mentioned at the beginning of the paper, and at capturing the aspect of meaning common to the sentences 'Aluminum conducts electricity' in the two dialects, viz., the controversial narrow meaning. One rationale for this idea is that someone who holds true the aforementioned sentence and who also holds true 'Bus tokens are made of aluminum', will be inclined to stick a bus token in the fuse box in sufficiently desperate circumstances, independently of whether he lives on earth or twin earth. The common behaviour of sticking the bus token in the fuse box is arguably to be explained in terms of a common aspect of content (or meaning of the sentence that is thought) in the same way on earth and twin earth, and the functional role factor is supposed to capture this aspect of content and meaning. Of course, this factor is aimed at capturing the aspect of meaning in common to the two dialects, not the aspect that differs (i.e., not reference or truth-conditions).

The second factor is aimed at capturing the difference between the two dialects, and must therefore go beyond internal functional role. For example, if a causal theory of reference seems promising, one might want to approach the difference between the dialects as a matter of one causal chain leading from aluminum to our word 'aluminum', and another causal chain leading from twin aluminum to their word 'aluminum'.

How are the two factors related? I claim that the functional role factor determines the nature of the referential-truth-conditional factor (without determining reference or truth conditions themselves). See 'Advertisement for a Semantics for Psychology', *op. cit.*, p. 643–644.¹⁷

I believe that opposition to the two factor account stems from

¹⁶ Hartry Field, 'Logic, Meaning and Conceptual Role', *Journal of Philosophy* 74, 1977: 379–409; Colin McGinn, 'The Structure of Content', in *Thought and Object*, Andrew Woodfield (ed.), Oxford University Press: Oxford, 1982; Brian Loar, 'Conceptual Role and Truth Conditions', *Notre Dame Journal of Formal Logic* 23, 1982: 272–283.

¹⁷ Actually, the best versions of a causal theory of reference include internal functional role; see Devitt, *Designation, op. cit.* On this account, the internal factor is simply part of the referential factor, and thus, trivially, the referential factor would also determine the functional role factor.

the fact that it seems to many to involve narrow content. But accepting the two factor account does not literally require accepting narrow content. Indeed, a version of the two factor account is compatible with a 'no content' view.¹⁸ Further, even if the two factor theorist accepts content, he is not compelled to accept narrow content; a two factor theorist can regard functional role inside the head as specifying a component of or determinant of content that it is important to distinguish from the referential component, but does not itself determine a special kind of content. But suppose the distinction between narrow content and narrow determinant of content is a distinction without a difference. What's wrong with narrow content? Many philosophers hold that for a representation to express a thought is for it to represent the world. In order for a representation (that expresses a thought) to represent the world, it must represent it as being one way or another, so it must impose some sort of conditions—truth conditions—that must obtain if the representation is to represent correctly. Hence, without truth conditions, there can be no genuine representation (at least representation of thoughts). Hence, a conception of meaning that attempts to abstract from the world—meaning inside the head—is nonsense. The idea of narrow content is also reviled as a revisionist notion that has been given no clear sense,¹⁹ and, perhaps incompatibly, as overly holistic.²⁰ Another line of argument against narrow content associates it with a picture of mind that holds that in perception, what is before the mind is not public objects, but rather internal representations of them. Thus, narrow content is alleged to come between us and the objects of our thoughts, leading the narrow content theorist to scepticism.²¹

I don't agree with these criticisms, but I will not argue against them or try to otherwise justify narrow meaning or content.

¹⁸ Stephen Stich, *The Case Against Belief*, MIT Press: Cambridge, 1983.

¹⁹ See the papers by Burge referred to in footnote 30.

²⁰ See, for example, Hilary Putnam's 'Computational Psychology and Interpretation Theory', in Putnam's *Realism and Reason*, Cambridge University Press: Cambridge, 1983.

²¹ See Gareth Evans, *The Varieties of Reference*, *op. cit.*; and John McDowell, 'De Re Senses', *Synthese* 1986. Narrow content is defended against this and other arguments by Evans and McDowell in a paper by Gabriel Segal, 'Conceptions of Content: Inner or Outer,' that is in his MIT Ph.D. thesis, 1987.

Rather, my strategy is to sneak up on the issue by taking a look at the kind of functional role semantics an anti-narrow content view leads to, arguing that it leads to a conception of functional role that is equivalent to the two factor theory, a conception of functional role that is made for narrow content (though it does not literally require narrow content).

Use theories of meaning bifurcate into, on the one hand, those that favour public use, most notably in communicating, but also in requesting, demanding, etc., and, on the other hand, those that restrict themselves to private use, use in thinking, deliberating, problem solving, and the like. The well-known use theorists—a list of which would contain Wittgenstein and Sellars, and perhaps the pragmatists—have mainly been concerned with public use. For example, Wittgenstein's dictum that in matters of meaning one should look to the use certainly refers to public use. I don't know of any use (or functional role) theorists who have clearly advocated private use (or function).²²

If you like use theories but hate narrow content, no doubt you will want to investigate a theory of truth conditions that appeals to public use. I will argue that some appeals to public use will not help in solving the problem at hand; another sort of appeal to public use will help, but will result in a theory equivalent to a two-factor account.

Before I do this, let me just briefly mention the metaphysical picture that I think motivates the predilection for private use theories. There is a fundamental difference between two types of representations. Representations of the sort that traffic between people, e.g., those on this page, must be heard or read to be understood. Mental representations, by contrast, do not require being seen or heard to be understood. Representations on the page can be understood only in virtue of their production of or transformation into mental representations, which do not, themselves, require such transformation. Searle describes mental

²² Except my 'Advertisement . . .', *op. cit.* A focus on use in thinking and the like was favoured by Harman in 'Language, Thought and Communication' in *Language, Mind and Knowledge*, Keith Gunderson, ed., University of Minnesota Press: Minneapolis, 1975. Harman also quotes Chomsky quoting Humboldt, favouring this idea. However, Harman's publications on functional role semantics, e.g., 'Conceptual Role Semantics', *Notre Dame Journal of Formal Logic*, 23: 242–256, 1982, have adopted a conception of use that includes public use.

representations as having *intrinsic* meaning, whereas representations on the page have only *observer-relative* meaning, the latter being inherited from the former. Thus a theory of meaning should apply in the first instance to intrinsic meaning, and observer-relative meaning can be derived via its relation to intrinsic meaning.²³

What does the intrinsic/observer relative distinction come to? The difference between representations in the head and representations on the page is that the representations on the page don't *do* anything; whereas the representations in the head, by contrast, are, in Haugeland's phrase, *semantically active*. Not just any kind of activity is semantic activity. To use Haugeland's example, if you write words on the back of ants, you get a frenzy of linguistic movement, but that's not semantic activity. You have semantic activity if, for example, sentences of the form '*p*' and '*p*→*q*' tend to produce '*q*'.²⁴

Of course, public words aren't *inert*. As Wittgenstein said, 'Every sign *by itself* seems dead. *What gives it life?*—In use it is alive'.²⁵ The point is that public tokens have a life only to the extent that we deliberately use them, whereas tokens in the head seem to have a life of their own. (That's why 'use' is more apt for public representations, 'function' for private ones.) Perhaps the best contrast for purposes of illustrating this picture is that between non-linguistic mental representations such as experiences and images on the one hand, and words on the page on the other. In sum, according to this picture, a use theory of meaning will apply in the first instance to use in the head, and only derivatively to use in communication.

Notice that the picture just described does not commit me to any sort of developmental priority to inner language, or to the idea that outer language is not used in thought. I can accept the picture just sketched, and also hold that we think in English, and

²³ See Searle's reply to critics of his 'Minds, Brains and Programs', *Behavioural and Brain Sciences* 3: 450-457, and John Haugeland, 'Programs, Causal Powers, and Intentionality', in the same issue, pp. 432-433.

²⁴ See Haugeland, *Artificial Intelligence: The Very Idea*, MIT Press: Cambridge, 1985, Ch 3.

²⁵ *Philosophical Investigations*, Macmillan: New York, 1953, section 432. The passage continues 'Is life breathed into it there?—or is the *use* its life?' I guess Wittgenstein favoured the latter option. I favour the former—as well as the latter (unless the two are understood so as to be incompatible.)

that in learning English, we learn new tools for thinking.²⁶

Though I believe there are good grounds for pursuing private use functional role semantics, I shall examine the public use option. Actually, what I will examine is not a 'pure' public use theory, a theory restricted only to public use. Rather, I shall examine a public plus private use theory, a theory that includes both. My reason for giving pure public use theory short shrift is that I cannot see how a public use theory could get off the ground without including use in thinking out loud. But once use in thinking out loud is included, I can see no non-behaviourist rationale for excluding thinking quietly. I will tend to use 'public use' to mean public plus private use.

Turning now to an examination of public plus private use, note that one cannot appeal to public use in the sense in which it could be recorded by a movie camera to solve the present problem, since the public uses on earth and twin earth are, on this way of individuating use, exactly the same. If we take movies of goings on involving the use of 'aluminum' on earth, we cannot distinguish them from movies of goings on involving the use of 'aluminum' on twin earth.

Of course, one could appeal to the fact that on earth they use 'aluminum' to refer to aluminum, whereas on twin earth they use 'aluminum' to refer to twin aluminum. Further, on earth, 'aluminum' is used to communicate information about aluminum, whereas on twin earth, 'aluminum' is used to communicate information about twin aluminum. But appealing to semantic differences such as differences in reference and what communication of information is about to distinguish functional roles is question-begging as mentioned earlier. One might as well regard '*Snow is white*' is used to mean that *snow is white* as a use theory of meaning of '*Snow is white*'.

The only way I can see to take advantage of the difference in public use within the confines of a functional role theory is to include relations to aluminium and twin aluminum as part of the inputs and outputs of the functional roles. Functional roles are often taken to stop at the skin (in sense and effector organs),

²⁶ A sketch of how this picture can be adapted to cope with the possibility that internal representations are entirely distinct from outer language can be found on p. 663-634 of my 'Advertisement for a Semantics for Psychology', *op. cit.*

but there is no principled reason to construe conceptual role in this way. Instead, functional roles can be understood as extending into the world of the things we sense and affect. Why not long-arm roles instead of short-arm roles? Aluminum is what people who think aluminum thoughts on earth see, hear, touch, and toss in the trash, whereas twin aluminum is what people who think twin-aluminum thoughts on twin-earth see, hear, touch, and toss in the trash. This fact could hardly be irrelevant to the fact that one group is thinking aluminum thoughts while the other is thinking twin-aluminum thoughts. If these types of effects of aluminum on the sense organs, and effects of the body on aluminum, are counted as part of the functional role of the earthian word 'aluminum', then, it would seem, we would have the sought-after difference between the truth conditions of 'aluminum'-thoughts on earth and twin-earth. Indeed, this is Harman's version of a functional role semantics aimed at truth-conditional content.

So far, so good. I shall now argue that a theory along these lines is equivalent to a 'two factor' theory, in which one factor is purely internal functional role, and the other is reference.

The key is to consider matters that have traditionally fallen in the theory of reference. Consider a 'visiting' twin-earth case. An earthian space-man lands on twin earth and radios home 'Good place to get aluminum'. Though the pieces of metal that he has just seen, clanked together, and held in his hands are not aluminum, one wants to take him as having said something false of aluminum, in his message home, not as having said something true of twin-aluminum. Twin earth is not in fact a good place to get aluminum. If his functional role of 'aluminum' at t is a matter of what he interacts with under the heading 'aluminum' at t , long-arm functional role semantics gets the wrong result.

Causal theories of reference have handled this sort of fact by talking about context of language acquisition, or, on more sophisticated versions, saying something to the effect that the dominant causal source (the ground of most of the causal chains) of his word 'aluminum' has been aluminum, and though he is now seeing and touching something else in connection with using the word 'aluminum', the dominant causal source remains

aluminum, so that's what he's referring to.²⁷ If he comes to live on twin earth for many years running a twin-aluminum cookware business, the dominant causal source of his word will shift to twin-aluminum, and so the reference of his word 'aluminum' will shift to twin-aluminum.

If a Harman-like functional role semantics theory of truth-conditional content is to be maintained, we have to understand input, not simply in terms of what is now causing patterns of energy to strike the body, but in some way that will have the effect gained by the 'dominant causal source' manoeuvre in the causal theory of reference. In fact, just such an idea is suggested by Harman for this kind of case. He uses a notion of 'normal context' in which the relevant inputs and outputs are not the ones that our explorer experiences just before radioing home, but rather the ones that obtained back home in his usual environment. If he stays on twin earth, the normal context of his 'aluminum' utterances will presumably shift *precisely* when the dominant causal source shifts. The 'normal context' manoeuvre and the 'dominant causal source' manoeuvre are analogous responses to the same problem within different theoretical contexts. Thus the theory of the input and output ends of the 'normal context' long-arm functional roles contains a twist familiar from the study of reference.

Further, other such manoeuvres will have to be made in long-armed conceptual role terms to handle other facts falling under the theory of reference. 'Moses was wise' has different truth conditions on earth and twin earth because the word 'Moses' refers to different prophets on the two planets. Any theory of reference will have to handle chains of borrowed reference to things past. Presumably, this will be done by pointing to the distinct man named 'Moses' in the two places, the reference of the word 'Moses' or some ancestor of it, in the mouths of those who saw these men, the borrowed references in the words of those who transmitted the words of those who saw them, etc. The truth-conditional functional role semanticist would have to use

²⁷ See Gareth Evans, 'The Causal Theory of Names', *Proceedings of the Aristotelian Society*, Supp. vol. 47: 187-208, 1973; Michael Devitt, 'The Semantics of Proper Names: A Causal Theory,' Ph.D. Thesis, Harvard University, 1972; 'Singular Terms', *Journal of Philosophy* 71: 183-205; *Designation*, Columbia University Press: New York, 1981, Chapters 2, 5, 7.

some manoeuvre to this effect as part of the theory's specification of the inputs and outputs of the twins who say 'Moses was wise'. Thus the functional role of my word 'Moses' will have to extend through my sensory inputs back thousands of years into the past. (The point also holds for kind terms, e.g., 'manna' rather than 'Moses'.)

(Notice that long-arm functional roles are beginning to look less deserving of being described as functional roles at all. I do not wish to make much of this, however, since whether long-arm roles ought really to count as functional roles is an issue about the proper application of what is, after all, just a technical term.)

Likewise for referential facts involving the division of linguistic labour, as in Putnam's elm/beech case. Presumably, any reasonable theory of reference would say that the references of my word 'elm' and 'beech' are fixed and maintained in virtue of something to do with the activities of people who really know about the relevant trees.²⁸ A functional role semantics theory of truth-conditional content will require some such device as part of what defines the functional roles, and in fact Harman regards roles in the minds of experts as part of the functional role of my word 'beech'.

What is the upshot? On the 'public use' version of functional role semantics that we are considering, as on the 'private use' two factor version, the difference in the truth conditions of 'aluminum conducts' between earth and twin earth has nothing to do with any difference in the *internal* part of the functional roles of the word 'aluminum' in the two communities. Rather, the difference is entirely a matter of factors outside the sense organs and motor effectors, factors involving relations between the metals in question and the bodies of the speakers of the two languages. Since the internal part of the functional role of 'aluminum' is the same for earthians and twin-earthians, *it can be identified with the narrow meaning* of 'aluminum'—if narrow meaning exists. The theory of this internal part would then be a theory of narrow meaning and content. The internal part is at least partly inferential, but not so with the external part. There

²⁸ More specifically, the references can be fixed originally by people who don't know the difference (e.g., a forester who sees elms but not beeches, and another forester who sees beeches but not elms). Experts are required to avoid the kind of widespread error that would eventuate in reference change. See Michael Devitt, *Designation*, *op. cit.*

need be no inferences of mine involved in the causal chains connecting aluminum with the energy impinging on my sensory surfaces.

So if we pursue a public use functional role theory of truth-conditions in the way I have sketched, what we find ourselves with is a theory that can be seen as consisting of two distinct (though not independent) sub-theories: one of these is at least partly inferential, and is not about the relations between words and the world; the other is not inferential,²⁹ and is about the relations between words and the world. The latter is about matters outside the head, and indeed looks very like a traditional theory of reference; the former mentions facts about what goes on inside the head, and looks like a functional role theory of narrow meaning. The whole thing is equivalent to the two-factor theory we arrived at starting with private use.

Whether the resulting theory should be regarded as a one-factor or as a two-factor theory depends on the explanatory advantages of one stance over the other. Is 'theory of reference' really a distinct subject matter? Or is referring merely one of a number of interlocking linguistic activities that must be understood together, or not at all? Is narrow content an explanatory notion, either in scientific psychology or everyday life?³⁰ I haven't tried to answer these questions.³¹ Rather, I have tried to show that within the general functional role approach,

²⁹This is over-simple. A better formulation is given in the last sentence of the last paragraph. There can be inferences of mine involved, there are almost certainly inferences of other people involved, and there can be 'public' inferences (e.g., you give the premise, I give the conclusion) involved. On my view, public inferences are just the public manifestation of private inferences.

³⁰Jerry Fodor, *Psychosemantics*, MIT: Cambridge, 1986, argues for the former, for narrow content as an explanatory idea in science, and Brian Loar, 'Social Content and Psychological Content', forthcoming in the Proceedings of the 1985 Oberlin Conference, for the latter. Tyler Burge argues against the former view in 'Individualism and Psychology', *Philosophical Review* XLV, 1, January, 1986: 3-45; and against the latter in 'Individualism and the Mental', *Midwest Studies* IV, P. French, T. Uehling, and H. Wettstein (eds.), University of Minnesota Press: Minneapolis, 1979: 73-121; 'Two Thought Experiments Reviewed', *Notre Dame Journal of Formal Logic*, XXIII, 3 (July, 1982): 284-293; and 'Intellectual Norms and Foundations of Mind', *Journal of Philosophy* LXXXIII, 12, December, 1986: 697-720.

³¹I'll tell you what I think, though: one scheme is better for some purposes, the other for other purposes. The short-arm roles are better for psychological explanation, long-arm roles are better for thinking about translation and communication. See Loar, *op. cit.* for related views.

the two sides must deal in equivalent ways with the task of accommodating truth conditions, and the way they have of accommodating truth conditions leaves room for narrow content.

APPENDIX

The reader will notice that I have approached the issue of how to connect truth conditions to what is in the head via thinking about the Putnam twin earth argument, ignoring the rather different type of consideration introduced in an important series of papers by Tyler Burge. (See the references in footnote 30.) The considerations raised by Burge's arguments are complex, and cannot be discussed in any detail here; but I do feel I should say something about the matter. A brief reminder of a Burgeian example: Fred has arthritis, uses the word 'arthritis' freely and normally, with one exception: one day he feels a pain in his thigh of the same sort he gets in his joints, and he says 'Good Lord, I have arthritis in the thigh'. In fact, what he says is impossible—the term applies only to rheumatoid disease of the joints. Imagine Fred, unchanged in all relevant non-intentional respects, in a counterfactual situations—let's speak of Fred_c for 'Counterfactual Fred'—in which the word 'arthritis' is used by his linguistic community so as to apply to arthritis in our sense of the word plus ailments of the thigh of the sort Fred has. Burge argues that the *de dicto* ascription ('believes that he had arthritis in the thigh') that applies to Fred does not apply to Fred_c. What Fred_c says is true; what Fred says is false. Since Fred and Fred_c are in relevant respects non-intentional duplicates, we have another argument that meaning isn't in the head. The relevant feature of this case is not the causal relations to natural kinds in the world (as in the aluminum case discussed earlier), but the determination of meaning by the linguistic community.

Can the facts of this case be explained within the theory of reference? Perhaps it will be said that either one can trace the difference between Fred and Fred_c back to different causal sources, or else one should simply hold that they deserve the same *de dicto* ascriptions. But this won't do. Suppose, for example, that the causal chain from the coining of the term 'arthritis' to Fred's and Fred_c's uses is exactly the same; the difference between the actual and counterfactual situation is a matter of language

change in the counterfactual language community—‘arthritis’ comes to be used more broadly than at first, applying to rheumatoid inflammations of bones as well as joints. But suppose this language change never happens to causally impinge on Fred_c or on any of the links in the causal chain leading to the original source of the word. It seems obvious that Burge’s arguments would apply in this case, unchanged.

The Burge cases can only be handled by a view that explicitly takes into account facts about the language community. Both the long-arm and short-arm functional roles that I have discussed are keyed to individual persons. The obvious suggestion is to introduce the quite different notion of functional role in a language. Functional role in a language would not be any kind of ‘average’ of functional roles of individuals in the language community. It is possible for almost everyone in a language community to misuse a word. What is crucial to a use’s being a misuse is the structure of deference in the language community. If almost everyone in the language community is strongly inclined to accept corrections in such a way that allows a few persons to determine what everyone would accept as correct, then, in that community, massive misuse of a word is perfectly possible. (This may actually be the case for ‘brisket’, another one of Burge’s examples.) So the functional role in a language must be cashed in a way that takes the structure of deference into account; perhaps this will be best done in terms of an idealized individual functional role.

Once we have a notion of functional role in a language, we have a number of options. One option (my construal of a suggestion by Tony Anderson) is that we focus on the fact that Fred_c’s role for ‘arthritis’ matches that of his language, whereas Fred’s role does not match that of his (i.e., our) language. In virtue of these facts, Fred_c’s sentence is true, whereas Fred’s sentence is false. But this can’t be right, because Fred’s sentence would be false even if his functional role matched that of the language perfectly. (You are wrong if you say you have arthritis in your thigh even if you know perfectly well what arthritis means.)

The main issue, it seems to me, is whether only the functional roles of words in the language are relevant to truth conditions, or whether the roles in individuals matter too. To explain the

difference between Fred and Fred_c, we will have to appeal to Fred_c's language community's roles—they, after all, make the difference between the actual and counterfactual cases. But there is reason to think the individual's functional roles are also relevant. Suppose A says 'B will arrive in precisely a fortnight', but A thinks a fortnight is ten days. In some contexts, it would be reasonable to report A as believing that B will arrive in ten days.³² The individual's functional role will loom large if, for example, we are interested in explaining why A goes to the train station in ten days, not two weeks. The relevance of individual functional role would explain another fact: Suppose Fred uses 'accident' the way the rest of us use 'mistake'. If he says 'I shot the hunter, because I thought he was a deer; it was an accident', we have some inclination to report his views using 'mistake'. Similarly, if we had a word 'tharthritis' (applying to arthritis and also rheumatoid inflammation of the thigh) in English, we would have some inclination to report Fred's 'Perhaps I have arthritis in my thigh' as: Fred thinks he may have tharthritis in his thigh. (Note, however, that the similarity in orthographic form itself makes for a certain similarity in functional role. For example a person who says 'I have arthritis' will react differently to 'Do you have arthritis?' than a person who says 'I have tharthritis.') It looks as if the close match between the functional role of 'arthritis' in Fred's mouth and 'tharthritis' in the language inclines us to report the former with the latter.³³

³² I am indebted to Sarah Patterson for pointing out to me that Burge makes this point himself. See p. 132 of Burge's 'Belief and Synonymy', *Journal of Philosophy* LXXV, 78, 1981: 119-138.

³³ I am very grateful to Sylvain Bromberger, Michael Devitt, Paul Horwich, Georges Rey, and Gabriel Segal for their helpful comments on earlier drafts, and to audiences at the University of Toronto and the University of Minnesota for stimulating discussions of the paper.

FUNCTIONAL ROLE AND TRUTH CONDITIONS

Ned Block and John Campbell

II—John Campbell

Due to circumstances beyond his control, John Campbell's contribution to the symposium was not available at the time of going to press. It will be published in the *Proceedings of the Aristotelian Society* 1987/8.