

Do: x is a necessary thing=DF There is an attribute that is such that (1) everything is necessarily such that there is something having that attribute, (2) x is necessarily such that it has that attribute, and (3) that attribute is not necessarily had by everything.

It seems to me conceivable that there is an attribute of being a K such that: (a) it is necessary that there are K s (so, everything necessarily has the attribute of there being something that is K in accordance with (1)), (b) any given K contingently exists (for example, on certain views of time it is necessary that there are things in time, but for anything in time it is not necessary that it exist), (c) any K is necessarily a K , and (d) there are non- K s.

On p. 52 Chisholm defines an ordered pair in terms of an attribute that orders a pair of things as follows:

Do: R is an attribute that *orders* something x to something y =DF R is an attribute that has among its instances (1) an attribute that has x as its only instance and (2) has x and y as its only instances.

It would seem to follow from this definition that if there is some attribute F uniquely possessed by a , and some attribute G that only a and b have, then the attribute of being an attribute orders a and b ; a curious result.

On p. 58 Chisholm contends that x is tenselessly F is equivalent to x was, is or will be F on the grounds that:

(U) There tenselessly are dinosaurs,

is equivalent to:

(T) There were dinosaurs or there are dinosaurs or there will be dinosaurs.

An advocate of the tenseless reading of 'are' would presumably reply that (T) does imply (U), but that is because it concerns dinosaurs, things necessarily in time, rather than things that are, arguably, timeless such as numbers.

In summary, Chisholm's book gives a valuable, if demanding, insight into his current thinking on a wide range of metaphysical issues.

THE UNIVERSITY OF KEELE

ANDRE GALLOIS

PHILOSOPHY OF LANGUAGE

Representations, Targets and Attitudes

By ROBERT CUMMINS

MIT, 1996. x + 154 pp. £25.00

Robert Cummins has produced an important book on representational content. The theory he defends is not new; it is a picture theory. But his way

of defending the theory is, and in the course of defending it he provides insights into the failings of alternative approaches, in particular: conceptual role semantics and covariance theories (the subject matters of Chs. 4 and 5 respectively). That said, he has produced just the bones of a theoretical approach to representational content—something I suspect he would be happy to acknowledge. In what follows, I shall try to give a feel for the approach and indicate potential problems it may face.

One thing to mention quickly is his adherence to a particular way in which to conduct the debate. This may be kindly described as *austerely naturalistic*. He is unwilling to take seriously the idea that there is no need to provide a definition of representational content in non-semantic and indeed non-cognitive terms so long as one can show how our scheme of attributing content may be subject to a non-semantic, non-cognitive characterisation of the conditions for its application (pp. 3–4). For Cummins, either one provides such a definition or no intellectual progress can be made. Those unsympathetic with this way of looking at the matter will still find the book of interest. They can find out whether an approach of the kind Cummins prefers is likely to be successful.

Cummins identifies four plausible desiderata that any naturalistic account of representational content should meet (pp. 85–86):

- (1) The content of a representation must be distinct from its target.
- (2) The content of a representation is distinct from the contents of the attitudes in which it figures.
- (3) The content of a representation must be independent of its use or functional role.
- (4) Representational content should be neither holistic nor atomistic.

Quite rightly, Cummins takes the major problem facing naturalistic accounts of representational content to be that of accommodating error, a problem which has plagued all previous accounts in one form or another. The constraints (i) to (iii) stem largely from an attempt to resolve the matter. The difficulty is this. Suppose that there is a representation in my head that we would find it natural to interpret as *there is a mouse*. If one adopts conceptual role semantics (for instance), the reason why the representation should have this interpretation is that it plays the right role in my reasoning, judgements based on perception and behaviour. But everybody concedes that I can sometimes make mistakes, for instance judging a rat to be a mouse. This suggests that the representation plays a slightly different role from the one envisaged, namely that which would go with the representation being interpreted as *there is a mouse or a rat*. Apparent error gives rise to grounds for reinterpretation which in turn banishes the presence of error. The problem of error for naturalistic accounts is that they can't explain how there is any.

Cummins considers various ways in which the proponents of previous theories have tried to deal with this and finds them wanting (pp. 41–47; 59–62). His solution is to distinguish between the target and the content of representation. The *target* of a representation is the thing which it is the

representation's function to represent. Representations obtain functions by being produced by mechanisms (dubbed 'intenders') which have that function (pp. 8–10). The content of an attitude like belief is determined by the target and the representational content applied to the target. For instance, suppose I have a mouse detector in my head. It might just produce the representation to be interpreted as *present* whenever there is a mouse. It is legitimate to ascribe to me a belief *that there is a mouse* because the *mouse* detector has indicated *present*. This would be so even if the mouse detector misfired in the presence of a rat. So the possibility of error is preserved (pp. 6–7).

But what makes something a mouse detector? Cummins suggests that the target of an intender is "what Σ [a cognitive system] expects to find when r [a token representation] is accessed" (p. 18). These expectations are design assumptions of the system. The thought is developed in Cummins' discussion of Target Fixation in Ch. 8. He writes,

"(TF) t is the target of an application of r just in case it is the function of that tokening of r to represent t " (p. 118).

He suggests that the function of the intender is fixed either by appeal to natural selection or design. Design theories "define the function of a mechanism or process in terms of its functional role" (p. 116). It is reasonably clear how the latter type of theory fits with Cummins' earlier remarks concerning target fixation. What I am not clear about is why he allows that selective functions may fix targets too. A design theory focuses on how the system actually works but the appeal to selective function concerns how the system should work. These may come apart. So it seems that Cummins is really committed to the design approach.

This is not just an idle question of clarification since those who press the point that naturalistic accounts have a problem with error will naturally look at this issue in order to see whether the difficulty once more arises. And here I think Cummins faces something of a dilemma. I am not clear how, if a system is likely to make mistakes, a capacity to detect mice (say) is not, counterintuitively, rather better described as a capacity to detect *mice or rats in certain circumstances*. So I think Cummins needs a notion of function which does not concern the actual role played by an intender. This is where an appeal to the selective approach might be thought to help. But such an appeal, although it might deal with this difficulty, appears ruled out by Cummins' endorsement of Fodor's objection to adaptational theories of content (p. 46). If being F and being G are correlated in an organism's environment (being a black speck and being an insect, for example), then, Cummins argues, it would be equally adaptive to represent either. So there would be no way of assigning one content rather than another. But if this is so, it is not clear to me how there would be a way of distinguishing between an intender's being a black speck detector and its being an insect detector either. So there are different possible patterns of error assignments with no way of deciding between them.

One might think that Cummins' discussion of valuable inaccuracy might help us resolve the issue (pp. 116–118). He suggests that the key to understanding why an intender could have the function of detecting Fs rather than Gs—even when its representations are more accurate as representations of Gs—is that it is the degree of fit between the representations and Fs which underlies the intender's contribution to the system. But this is no help. The issue is not how something could have the function of detecting Fs when it produces inaccurate representations of them. The issue is how something could have the function of detecting Fs rather than Gs when the detection of either is equally adaptive in the circumstances. So (contrary to advertisement) it does not seem that the 'target' element gives us a satisfactory treatment of error.

I have not yet discussed the main component of Cummins' theory, the picture element—the subject matter of Ch. 7. He adopts this as a way of avoiding a straightforward functional account of representative content along the lines of Dretske and Millikan, which faces the problem concerning error already discussed (pp. 55–57). Suppose one took the content of a representation to be that which, when it was true, enabled the intender to successfully perform its function. This would ignore the fact that successful performance of a function *need not* involve *accurate* representation. If enough little ambient black particles in a trout's environment are insects, then an insect detector which registers these as insects to the trout would be functioning successfully even if there is the odd case of misrepresentation. It would be more costly in processing and response time to require greater discrimination (p. 45).

Cummins' idea is that,

“R represents C if and only if R and C are isomorphic” (p. 90).

Only structures represent. A consequence of this is that sentences of natural languages do not represent. They only have their semantic properties conventionally. They have *meaning for* a language user but they do not have *meaning* in Cummins' sense (pp. 86–87; Ch. 10). He thinks that the pictorial scheme has a number of merits:

- (1) It breaks the connection between the use of a representation and its content (satisfying desideratum (iii)). A representation may represent something inaccurately by lacking some of the structural properties the thing to be represented has, yet perform its proper function because it is accurate enough for the purposes to hand (p. 93).
- (2) Structural properties may be explanatory of cognitive processing (satisfying what Cummins calls the explanatory constraint, p. 2, pp. 94–96).
- (3) It is neutral between holism and atomism. Graphs are pictorial schemes which are holistic. Cartoon faces are pictorial schemes which are atomistic. So clearly a pictorial scheme can be either (satisfying desideratum (iv), pp. 75–76; p. 97). This is in contrast with conceptual role semantics or causal accounts.

A natural objection to make to Cummins' approach is that a representation will be isomorphic to many things. So we need an account of which isomorphism is relevant to the content of the representation (p. 97). Cummins considers this objection and his reply is, basically, to deny that this is necessary. He invites us to consider as an example a card with a slit cut in it that will, as it is 'read', direct a car through a maze. He suggests that the card represents all the things which are isomorphic with the slit pattern (including the correct path) (p. 99; pp. 128–129). He does not think that this is a problem but rather indicates how a representation may play different explanatory roles depending on how it is exploited and the target that it was produced to service. However, he does concede that the content of attitudes is just a function of representational content plus target. If this is so, then we would have nothing like the intuitive assignment of belief contents. The target will not resolve the fact that the representation has multiple interpretations corresponding to its multiple isomorphisms. So while Cummins might have produced an account of representational content that can make sense of error, it is very much a revisionary account of the assignments of content that we would wish to make in spite of the claims that Cummins makes for his overall picture (p. 14, for example).

A second worry concerns whether Cummins can really hold that isomorphism is a case of non-conventional 'intrinsic' representation (p. 93, for example). If this is so, doesn't it follow that any object represents any other object with which it is structurally similar? The obvious reply is to suggest that this is not the case because only certain objects are used as representations by being produced by intenders. But if use comes in here to determine when isomorphism is relevant, it is tempting to think that one should appeal to use to settle which of the many isomorphisms are relevant to the interpretation of a representation so resolving the previous difficulty I identified. One could do this while still leaving room for error.

As I said, I think Cummins only offers the bare bones of an approach to these issues. However, these bones are well worth picking over.

THE UNIVERSITY OF NOTTINGHAM

PAUL NOORDHOF

Coming to Our Senses: A Naturalistic Program for Semantic Localism

By MICHAEL DEVITT

Cambridge University Press, 1996. x + 338 pp. £35.00 cloth, £13.95 paper

All too often, fundamental questions go unasked. Answers to such questions as 'What exactly do we need a theory of X for?' are either taken for granted or else completely ignored. In his *Coming to Our Senses*, however, Michael Devitt asks just such fundamental questions about semantics: What are the semantic tasks? Why are they worthwhile? And how should we accomplish them? By asking such fundamental questions, and by providing a naturalistic method for answering them, *Coming to Our Senses* makes a substantial contribution to both semantics and the philosophy of language. The book is well organised and carefully argued. It is reader friendly to philosophers working

outside of semantics. Given the fundamental nature of the issues addressed, it is broad enough to function well as a core text for a graduate seminar.

One of the main contributions of the book is the methodology offered. Systematic methodologies for semantics are virtually non-existent and so, by providing one, Devitt both draws attention to a deficit while addressing it. The proposed methodology affords several additional advantages. It is naturalistic. It is geared to avoid ad hoc moves and it allows a place for intuitions without giving them free rein. As a result, semantics relies less on armchair intuitions and is thus more like other empirical sciences.

Using this methodology and presupposing a realist metaphysics, Devitt argues for a representationalist (i.e. truth referential) semantics coupled with semantic localism (i.e. *only some* of the inferential properties of a word constitute its meaning). He argues against Cartesianism (i.e. linguistic competence requires privileged knowledge of meanings). Along the way he argues against two-factor and verificationist theories of reference as well as against direct reference (i.e. the meaning of a term is merely the property of referring to its bearer). He also argues against revisionism and eliminativism. He questions the explanatory usefulness of ascribing narrow meanings and offers compelling grounds for questioning the dogma that a token has but one meaning.

Along the way, Devitt bolsters his position by arguing that his opponents have offered insufficient reasons for abandoning or revising his position. As a defensive strategy, this is perfectly adequate. It is unclear, however, whether Devitt uses such a strategy offensively. Suppose that a certain position is assumed to be the default position and it is argued that no one is justified in holding any other position since there are insufficient reasons for abandoning the default position. This form of argument is unfair. By assuming that a position is the default position important questions are probably begged. It also tacitly raises the requirements since it no longer suffices for an opponent of the default position to establish the viability of her position. Instead, she must show that her position is clearly better *and* that it is clearly better on grounds that proponents of the default position would accept.

Devitt treats realism as a default position. First and to his credit, Devitt makes his commitment to realism explicit at the very start (p. 2). He is a "common sense" realist which involves a plausible combination of metaphysical and epistemological realism theses. If Devitt thinks that this particular combination is the only tenable one (as it sometimes seems he does) then I disagree but cannot argue the point here.

Second, Devitt claims that his realist presuppositions are unproblematic. Presumably, this is because "The chances that discoveries about meanings will cast doubt on this realism are, in [his] view, just about nil" (p. 2). If Devitt thinks that his realism plays no role in what follows, I disagree. His realist presuppositions certainly play a role when he argues against verificationist theories of meaning by taking an alleged commitment to non-realism as a *reductio* of such theories (p. 194). In arguing for a principled basis for his localism, Devitt's realism comes into play. "Putative meanings, like planets

and everything else, are the way they are and not some other way. That is the way the world is and nothing more needs to be said" (p. 117).

Third and finally, it is sometimes difficult to determine just which realism theses Devitt is taking for granted. Consider, for example, metaphysical realism. Metaphysical realism maintains that there are objective facts about the individuation of things and collections of things. Since an objective causal structure requires this sort of realism, Devitt's discussion of a causal theory of reference appears committed to it. Yet, Devitt makes claims directly at odds with metaphysical realism. He claims that "we are free to choose which properties to name" (p. 102). Metaphysical realism must maintain that some, indeed most, properties are utterly ineligible to be referred to by us. It is the non-realist who contends that we have more referential freedom. Presumably, Devitt is a metaphysical realist who is not quantifying over *all* properties when he says that we are free to name any. Rather, he is best understood as quantifying over just those properties that are objectively eligible (according to metaphysical realism) to be named by us. Once again, his realist presuppositions play a significant role.

Everyone has to start somewhere and Devitt plays fair to the extent that he makes his presuppositions explicit. He also chooses a very plausible set of starting points. It is just that I think he is wrong to think that his are the *only* plausible starting points. He is wrong to be so dismissive of various non-realisms.

THE UNIVERSITY OF MARYLAND, BALTIMORE COUNTY MARY KATE MCGOWAN

PHILOSOPHY OF MIND

Language, Thought and Consciousness. An Essay in Philosophical Psychology

By PETER CARRUTHERS

Cambridge University Press, 1996. xvi + 292 pp. £35.00

Carruthers aims to establish the possibility that we often think in the same language that we speak. Not a difficult task, you might think. Doesn't introspection reveal an inner soliloquy conducted in the acoustic imagination? But according to a venerable tradition, natural language is something akin to a code. We translate our thought into natural language only when we want to make it publicly available, or in the case of conscious inner verbalisations, to make a note to ourselves for future reference. A currently influential version of this *communicative conception* of language is due in large part to Fodor: cognition consists primarily in the manipulation of sentences of Mentalese, an innate, universal symbol system. Carruthers thinks this a mistake, and advocates the *cognitive conception* of language: we do a good deal of our thinking in natural language itself.

The first half of the book is largely an extended debate with Fodor. The case for the communicative conception has numerous strands, ranging from everyday occurrences such as the tip-of-the-tongue phenomenon, to theoretical