

THE GENERALITY CONSTRAINT AND CATEGORIAL RESTRICTIONS

BY ELISABETH CAMP

We should not admit categorial restrictions on the significance of syntactically well formed strings. Syntactically well formed but semantically absurd strings, such as 'Life's but a walking shadow' and 'Caesar is a prime number', can express thoughts; and competent thinkers both are able to grasp these and ought to be able to. Gareth Evans' generality constraint, though Evans himself restricted it, should be viewed as a fully general constraint on concept possession and propositional thought. For (a) even well formed but semantically cross-categorial strings often do possess substantive inferential roles; (b) hearers exploit these inferential roles in interpreting such strings metaphorically; (c) there is no good reason to deny truth-conditions to strings with inferential roles.

I. THE GENERALITY CONSTRAINT

This paper concerns the limits of propositional thought, and the requirements on comprehension which are imposed by competence with respect to a given concept. Propositional thoughts are thoughts reportable by that-clauses, for instance, the thought that there is beer in the refrigerator. In the terms I shall use, thoughts are composed out of concepts, and have propositions as their contents. Different thoughts can have the same propositional content, by virtue of being composed out of distinct but co-extensional concepts. (In what follows I can be neutral about just how to understand propositions: as structured sets of objects and properties, as possible worlds, or in some other way.) Concepts and the thoughts they compose are individuated by their possession-conditions; it thus makes sense to ask whether particular thinkers meet those conditions, and so whether they grasp a concept or thought. Thoughts are, in this sense, abstract objects and not just the particular psychological states of individuals at times.

Given that propositional thoughts are composed out of concepts, it follows that such thoughts must be connected to one another in systematic ways, in virtue of their constituent concepts. Gareth Evans illustrates the point thus:

It seems to me that there must be a sense in which thoughts are structured. The thought that John is happy has something in common with the thought that Harry is happy, and ... something in common with the thought that John is sad.... Thus, someone who thinks that John is happy and that Harry is happy exercises on two occasions the conceptual ability which we call 'possessing the concept of happiness'.¹

This fact about propositional thought is so basic that if someone fails to grasp that his thoughts are related in this way, we question whether he really understands them: someone who sees no connection between two thoughts like these cannot really be grasping either of them.

Generalizing from this example (Evans, p. 104, fn. 21), we get a picture of concepts as the articulating strands of thought, the lines which at once connect and distinguish distinct thoughts:

We thus see the thought that *a* is F as lying at the intersection of two series of thoughts: on the one hand, the series of thoughts that *a* is F, that *b* is F, that *c* is F, ..., and, on the other hand, the series of thoughts that *a* is F, that *a* is G, that *a* is H.

This picture in turn suggests a condition for which thoughts we ought to be able to understand: if the structure is truly systematic, it should contain no unexplained gaps. Part of what it is for someone to possess a concept, on this view, is for that concept to be fully caught up in a network of potential thoughts – for it to combine generally with the thinker's other concepts (subject, that is, to a mental analogue of being syntactically well formed). Evans (p. 100) calls this a 'fundamental constraint' on 'the nature of our conceivings', and (p. 104) dubs it 'the generality constraint':

If a subject can be credited with the thought that *a* is F, then he must have the conceptual resources for entertaining the thought that *a* is G, for every property of being G of which he has a conception.

Even if one denies that the generality constraint follows ineluctably from the very nature of thought, something like the requisite generality clearly applies to our thinking, and differentiates it from the mental representations of other animals.² The recombinability of our concepts helps to explain the rich generativity of our conceptual capacities. Further, if we accept (as Evans and many others do) that understanding a thought essentially involves grasping its truth-conditions, then it seems to be an essential feature of propositional thought that we can understand a new thought without knowing whether the world is as it specifies. But it is difficult to see how this could happen unless we employ our previous mastery of the thought's constituent concepts to determine what would make the new thought true.

¹ G. Evans, *The Varieties of Reference* (Oxford: Clarendon Press, 1982), p. 100.

² See Evans, 'Semantic Theory and Tacit Knowledge', repr. in his *Collected Papers* (Oxford: Clarendon Press, 1985), pp. 322–42.

Many thinkers have, however, accepted this basic condition of systematicity while resisting a fully general formulation of the generality constraint. They have maintained that I neither can nor need to understand the supposed thoughts constituted by each node of my conceptual network: I need not be capable of entertaining the thought that *a* is G for ‘every property of being G of which I have a conception’. Some barriers to our actually achieving full generality may perhaps be placed to one side. For instance, many combinations of concepts are too complex to be entertained by any finite thinker. Someone might be barred from entertaining some thoughts because they are too psychologically troubling, or even because a physiological reaction prevents the neural states corresponding to two specific concepts from co-occurring.³ However, these barriers are not inherently conceptual in nature, and so do not limit the generality of our conceptual capacities *per se*.

Instead, the primary objection to full generality is that some combinations of concepts are so wildly heterogeneous that we cannot fit them together to form a complete thought, and so should not be expected to. Thus one might think that although I do understand the words involved, neither I nor anyone else really understands what it would take for the thoughts putatively expressed by sentences like

1. Caesar is a prime number
2. Colourless green ideas sleep furiously

to be true. And if we cannot understand such thoughts, then insisting on a fully general formulation of the generality constraint either entails that we are all incompetent thinkers, or else sets an impoverished standard for what counts as understanding across the board.

Indeed, one might go further, and maintain not only that we cannot grasp the conditions under which such supposed thoughts would be true, but also that we cannot even properly assess them as false. Caesar, one might think, just is not the sort of thing that can either be or fail to be a prime number. Absurd ‘thoughts’ like these might seem to involve such serious category mistakes that the strings ‘expressing’ them ought to be counted as syntactically well formed nonsense.⁴ If this is right, then there is no thought *there* to be understood at the nexus of the constituent concepts. And if so, then failure to grasp such nothingness obviously should not impugn anyone’s competence with respect to the relevant concepts.

³ Peacocke raises the last two possibilities in *A Study of Concepts* (MIT Press, 1992), p. 43.

⁴ Cf. ‘Only expressions can be affirmed or denied to be absurd. Nature provides no absurdities; nor can we even say that thoughts such as beliefs or suppositions or conceptions are or are not absurd. For what is absurd is unthinkable’: Ryle, ‘Categories’, in A. Flew (ed.), *Logic and Language*, 2nd series (Oxford: Blackwell, 1953), pp. 65–81, at p. 76.

In the light of these worries, some philosophers, following Ryle, Russell and Carnap, have proposed a highly restricted form of the generality constraint. Strawson writes ‘The idea of a predicate is correlative with that of a *range* of distinguishable individuals of which the predicate can be significantly, though perhaps not necessarily truly, affirmed’.⁵ Evans adds to his definition of the generality constraint, cited above, the following *caveat* (in a footnote): ‘With a proviso about the *categorial appropriateness* of the predicates to the subjects’.⁶ Peacocke’s version of the constraint stipulates that

If a thinker can entertain the thought Fa and also possesses the singular mode of presentation b , which refers to something in *the range of objects* of which the concept F is true or false, then the thinker has the conceptual capacity for propositional attitudes containing the content Fb .⁷

Concepts, they all agree, have limited ‘ranges of significance’ or ‘categories of appropriate application’. It is only within these ranges that there are thoughts with genuine truth-conditions and truth-values to be understood, and only here that the generality constraint applies. Within the ranges, however, the sense in which we can know what it would take for a thought to be true is quite robust, and substantive standards for competence can accordingly be maintained.

I shall argue, against this consensus, that we should not impose categorial restrictions on either conceptual significance or conceptual competence. My argument proceeds in four steps. (a) The project of delimiting appropriate categories faces serious, though perhaps not insurmountable, difficulties. I adopt the least restrictive plausible categories of significance. (b) Strings that count as cross-categorial on this criterion, such as (1) above, often do possess substantive inferential roles, and should therefore be counted as significant. (c) Normal thinkers do routinely make use of these inferential roles, in particular in the process of construing metaphors. Therefore full competence requires that thinkers must be capable of grasping these inferential roles. Finally, (d) there is no good reason to deny that cross-categorial predications with inferential roles also have truth-conditions.

I conclude that we should abandon the project of delimiting a narrow range within which robust understanding of every thought is both necessary and sufficient for competence with a given concept, but outside which there lies no thought at all. We can still admit that our understanding of wildly cross-categorial thoughts is thinner than, and even dependent on, our understanding of more paradigmatic combinations of concepts. We can also admit that some combinations of concepts that correspond with syntactically

⁵ P.F. Strawson, *Individuals* (London: Methuen, 1959), p. 99n.

⁶ Evans, *The Varieties of Reference*, p. 101; my italics.

⁷ Peacocke, *A Study of Concepts*, p. 42; my italics.

well formed strings are indeed nonsense – albeit for reasons other than the violation of categorial restrictions. We do fuller justice to the competence that we actually demand of thinkers if we reject sharp *a priori* boundaries between the intelligible and the nonsensical, and instead treat significance, understanding and competence as matters of degree.

II. CRITERIA OF SIGNIFICANCE

In this section I take up the question of how to fix ranges of significance for concepts, before turning to the question of whether putative thoughts outside this range really are either nonsensical or incapable of being understood. Throughout this discussion, it will be important to remember that these criteria could also be treated as conditions merely on ranges of competence, rather than on ranges of significance. Entertaining thoughts outside the categorial bounds would then be a remarkable but basically gratuitous feat. Treating the criteria in this way remains a fallback position for now (I argue against this weaker position in §III(b) below). However, all the defenders of category restrictions mentioned above have advocated restrictions on significance rather than merely on competence, and the reasoning behind imposing the restrictions supports treating them in this way. That is, it is supposed to be something about the concepts themselves that prevents our fitting them together, and so it is natural to think that the concepts simply cannot be fitted together.

In either case, we still need to make explicit the criteria for determining which concepts can be combined. First, we seek to understand just how and why concepts should be limited in their application, if indeed they are. Secondly, as theorists, we need a way of deciding whether apparent lapses of generality are limitations of thinkers' capabilities, or genuine limitations on a concept's applicability. Speakers disagree about whether strings in the language express thoughts; we need a way to establish who is right. Indeed, even if we all fail to make sense of an alleged thought, we still need to determine whether this should be regarded as a fact about the concepts involved, or about our collective incompetence as thinkers.

How should we go about fixing the relevant criteria? The leading idea behind imposing categorial restrictions is that the world is divided into importantly different sorts of things, and that concepts are supposed to be suited for application to only certain of those sorts. It seems obvious that our categories of significance should mirror the relevant sorts. But how are we supposed to identify what these sorts are? In his attempt to put some flesh on Ryle's sketchy comments about category mistakes, Strawson says roughly

that they are provided by ‘individuating designations’, terms that ‘embody or imply principles for distinguishing, counting, and identifying individuals’.⁸ An individual may be brought under several individuating designations. Thus a particular car may be variously identified as a Honda, a sedan, a foreign-made car, a vehicle and a hunk of metal, among other things. A predicate *F* is ‘category-mismatched’ for an individual *a*, and the sentence ‘*Fa*’ is thus nonsensical, Strawson claims (‘Categories’, p. 203), if and only if *F* is or implies a predicate that is ‘*a priori* rejectable’ not just for one but for all of *a*’s individuating designations. So the predicate ‘is Secretary-General of the United Nations’ is category-mismatched for the car, because we know *a priori* that the predicate cannot combine with any of the car’s individuating designations to produce a true sentence.

Although this method has some appeal, even Strawson’s quite brief discussion reveals how messy and difficult the project would be to implement. Our categories for linguistic and conceptual significance will depend on which terms count as individuating designations. And this in turn will depend upon our principle of identity for individuals. The linguistic and conceptual project of delimiting the boundaries of significance thus turns out to be intimately bound up with the metaphysical project of limning reality’s basic ontology. This is perhaps no great surprise, but it renders rather less plausible the claim that competence in a language brings with it a firm grip on just which sentences are significant, and so on which concepts can be combined into genuine thoughts.

Perhaps the most systematic attempts to work out a detailed system of sortal distinctions come from lexical semantics, and in particular from the attempt to specify the semantic knowledge that speakers employ not only in deciding whether a sentence is ‘semantically anomalous’ (that is, categorially inappropriate), but also in identifying and resolving ambiguities among readings of a sentence, and in determining relationships of paraphrase and implication among sentences.⁹ Lexical entries for words take the form of specifications of ‘semantic markers’, ‘distinguishers’ and ‘selection restrictions’. Semantic markers indicate the restrictions on the sorts of objects that can fall under a given term; the distinguisher specifies the distinctive feature(s), if any, of things that do fall under it; and the selection restrictions identify the semantic markers that must occur in the surrounding linguistic context in order for the term in question to be inserted into that context without semantic anomaly (for instance, the adjective ‘red’ must modify a

⁸ Strawson, ‘Categories’, in O. Wood and G. Pitcher (eds), *Ryle* (Garden City: Doubleday, 1970), pp. 181–211, at p. 200.

⁹ See, e.g., N. Chomsky, ‘Degrees of Grammaticalness’, and J. Katz and J. Fodor, ‘The Structure of a Semantic Theory’, both in J. Fodor and J. Katz (eds), *The Structure of Language* (Englewood Cliffs: Prentice-Hall, 1964), pp. 384–9, 479–518.

concrete noun). Thus the lexical entry for 'colourful' might be written as follows:¹⁰

- (a) *Colourful* → Adjective → (*Colour*) → [*Abounding in contrast or variety of bright colours*] <(*Physical object*) ∨ (*Social activity*)>
 (b) *Colourful* → Adjective → (*Evaluative*) → [*Having distinctive character, vividness, or picturesqueness*] <(*Aesthetic object*) ∨ (*Social activity*)>

while the lexical entry for 'man' is something like

- (a) *Man* → Noun concrete → Noun masculine → (*Physical object*) → (*Human*)
 → (*Adult*) → (*Male*)
 (b) *Man* → Noun concrete → (*Physical object*) → (*Human*)
 (c) *Man* → Noun abstract → (*Human*).¹¹

(Here markers are represented in parentheses, and distinguishers in square brackets; terms in Roman type specify grammatical categories, and terms in angle brackets specify selection restrictions.) The combinatorial rules then specify that only terms with compatible selection restrictions and semantic markers can combine without anomaly, and only the compatible markers are retained in giving the combined phrase's meaning. The hope is that such an account will isolate a relatively small set of key markers, such as *Physical object*, *Social activity*, *Human* and *Male*, which represent the fundamental categories into which 'things' are sorted.

One point to notice about this approach, and about every attempt to delineate categorial restrictions, is that the project always ends up 'revealing' that natural languages are massively ambiguous in ways we would not otherwise have suspected.¹² That is, in order to draw categorial boundaries that do any real work, one ends up sorting things so finely that many terms turn out to have applications across multiple categories. But then because meaning is by hypothesis defined only on a categorial basis, it must be defined anew for each category.

We should, however, resist this 'revelation' of massive systematic ambiguity unless it is genuinely forced upon us. Postulating ambiguity is, as Kripke says, a 'lazy man's approach'; and in this case the evidence for ambiguity is

¹⁰ Katz and Fodor, 'The Structure of a Semantic Theory', p. 507.

¹¹ The second entry for 'man' is exemplified in 'Every man on board was saved except an elderly couple', the third in 'Man is occasionally rational': 'The Structure of a Semantic Theory', p. 510.

¹² For instance, Katz and Fodor conclude on the basis of their analysis that 'The man hits the colourful ball' exhibits a four-fold ambiguity. Similarly, Ryle concludes that 'existence' has at least two senses, 'somewhat as "rising" has different senses in "The tide is rising", "Hopes are rising" and "The average age of death is rising"': *The Concept of Mind* (London: Hutchinson, 1949), p. 23; and as 'in' exhibits different senses in 'She came home in a flood of tears and a sedan-chair' (p. 22).

quite weak.¹³ First, speakers are not in general aware of much of the postulated ambiguity. Secondly, if we count these terms as ambiguous, then we lose the resources for explaining how speakers extend their understanding of a term's application from one category to another, as they clearly and easily do. Finally, we lack any way to distinguish these cases from paradigmatic instances of ambiguity, such as 'cape', 'bank', and 'mass', where there is little or no projectability from one meaning to the other.

Ambiguity aside, we should worry about whether most terms in our language do admit of the neat analysis that the project requires. However, even if something like the selectional approach did succeed in providing a systematic analysis of the categories 'encoded' in our language, this would only make the present difficulty clearer. We would then be left with two unpalatable alternatives. The first is to confine the range of significance to the narrowest semantic categories marked in each lexical entry. But this would be restrictive, rendering a much broader swathe of our thought and talk nonsensical than one might have hoped. For instance,

3. The man in black is quite a colourful guy

would count as meaningless on the basis of the lexical entries above. This seems the wrong way to go; we initially intended to rule out only the most absurd combinations of concepts, like 'Caesar is a prime number'. The second option is to treat only some categories as delimiting the ranges of significance. But we originally turned to the lexical categories in the hopes that they would isolate the fundamental 'sorts' of thing. If we do not opt for the most restrictive categories, then we need a new criterion for deciding which categories mirror the especially fundamental sorts. And this seems just to throw us back onto our initial question-begging notions about whether it is 'really' possible for the predicate to apply significantly.

These are important difficulties for someone seriously engaged in delimiting the ranges of concepts' significance. However, for my purposes they are mere matters of detail. I shall work with the most permissive categories that could hope to make the restriction on ranges of significance worth imposing. Both the Strawsonian and the lexical-semantic approaches employ these categories, along with other finer ones. Among 'things' broadly construed, then, I shall distinguish abstract from concrete objects, animate concrete objects from inanimate ones, and human animate concrete objects from non-human ones. Just these three coarse divisions turn out to pose too strong a restriction on the generality that our conceptual abilities both do and need to exhibit.

¹³ S. Kripke, 'Speaker's Reference and Semantic Reference', repr. in P. Ludlow (ed.), *Readings in the Philosophy of Language* (MIT Press, 1997), pp. 382–414, at p. 401.

III. INFERENCEAL ROLE, METAPHOR AND LITERAL NONSENSE

In this section, I challenge the idea that categorially inappropriate predications, as fixed by the criteria above, are in general nonsensical. I shall argue that semantically cross-categorial, syntactically well formed strings can be used in a range of ways in which syntactically malformed strings cannot. Specifically, they have inferential roles which can be, and routinely are, exploited in material reasoning and in metaphorical interpretation. If even some cross-categorial strings are significant, then the criterion developed above fails. Because my arguments do not rely on distinctive features of my examples, it seems unlikely that any other criterion could succeed.

I begin with some examples of cross-categorial predications. For each category, I have offered at least one example with the categorial violation running in each direction (the relevant predications are italicized). The narrowest sortal distinction is between humans and non-humans; cross-categorial predications of this sort are

4. *Odysseus was a pig* while on Circe's island
5. *George is a real rooster* of a guy
6. *The lion reigns* over the savanna.

The next sortal distinction is between animate and inanimate objects:

7. *The prison guard was an iron statue*, his arms folded across his chest
8. But soft! What light through yonder window breaks? It is the east, and *Juliet is the sun*
9. A solitary book lay in the driveway, dropped by the movers, *its pages waving adieu* in the breeze.

Presumably violations of the most general distinction, between concrete and abstract objects, will be the most wildly heterogeneous, and therefore the most difficult to construe. Examples here are

10. *Life's but a walking shadow*
11. *Confusion now hath made his masterpiece*
12. *Caesar is a prime number*.

I shall focus my discussion on the last category, precisely because it is the most challenging. I hope that the examples for the other categories already suggest how easy cross-categorial strings can be to generate and comprehend; thus the considerations I adduce apply with even greater force to these cases.

III(a). *Semantic evidence*

Given how much we can do with such cross-categorial predications, it is important to remember that we cannot, or do not, do any of this with syntactically malformed strings. For instance, a Dadaist string like Max Ernst's

13. Price they are yesterday agreeing afterwards paintings

or Kurt Schwitters' Poem #48

14. Staggering / Earthworm / Fishies / Clocks / The cow / The forest leaf
the leaves¹⁴

may be evocative, at least for some people. What is evoked in each case may depend upon the constituent words, and even upon the order in which they appear. Nevertheless what hearers get out of these strings is at most a feeling, or a constellation of images and emotions. They cannot extract any *claims* to which the speaker has committed himself by saying what he does. When listeners talk about the images and emotions associated with these strings, they do not offer paraphrases of what the speaker meant. Rather they describe their own responses, much as they might describe their responses to a sound or a smell. Moreover, syntactically malformed strings like these, made up of real words, are often less comprehensible even than apparently syntactically well formed strings that are partially constituted of meaningless pseudo-words, such as Lewis Carroll's 'Jabberwocky'.

Next, speakers can understand and answer only syntactically well formed 'Yes'-'No' questions:

15. Could staggering earthworm fishies be clocks the cow?

does not afford an answer, while

16. Could Caesar be a prime number?

does. Upon hearing the latter, one might justifiably wonder about the speaker's intentions in asking such a question. But what the question asks is clear. It is equally clear, in ordinary talk at least, that the answer is 'No'. That Caesar is not a prime number is necessarily and obviously true, precisely because Caesar is not a number at all. By contrast

17. Staggering earthworm fishies is not clocks the cow

is no more amenable to truth-evaluation than the original without the 'not'.

The fact that speakers have ideas about appropriate responses to, and the truth-values of, such complex constructions provides provisional evidence that the initial cross-categorial strings, such as (1), are themselves significant:

¹⁴ The slashes indicate line breaks. Both examples are from R. Motherwell (ed.), *The Dada Painters and Poets: an Anthology*, 2nd edn (Cambridge: Belknap, 1979).

in this regard they class together with significant strings, and separately from mere word-salad (and subsentential phrases). The apparent significance of the complex constructions also raises the question of how to provide a principled and unified treatment of both complex and simple ones. But such data are inconclusive. We might have independent grounds (perhaps from the semantics of vague predicates) for believing that natural-language negation is sometimes external. If this is so, then the truth-assessability of

18. Caesar is *not* a prime number

need not imply that the unnegated thought itself has determinate truth-conditions and a truth-value. We might also have independent reason (though it is much less clear what reason) for thinking that questions of the form ‘Could *a* be *F*?’ should in general be analysed meta-linguistically, as requests for information about the types of predicate and subject involved.

However, competent thinkers can do more with these strings: they can and do generate material inferences from them. The fact that they do so shows, I think, that they have indeed succeeded in combining the inferential roles of the constituent concepts together to determine the inferential role of the thought as a whole. (By ‘inferential role’ I mean the core set of inferences a thinker needs to be able to draw so as to be considered competent in the use of a thought’s constituent concepts, and which determine the inferential power of the whole thought. I can be relatively neutral about just how to fix this set, but I tend to think that its boundaries are fuzzy, and largely overlapping rather than fully identical across different thinkers. I also assume that inferences have different degrees of strength, representing how central they are to the concept in question; but not much hangs on this.) Competent thinkers can reason, for instance, from the hypothetical truth of Caesar’s being a prime number to the conclusion that he is not evenly divisible by a number other than one and himself, or (more interestingly) to the conclusion that he is an abstract object, and therefore lacks efficacy. From this latter conclusion they would be entitled to infer that Caesar could not be an effective emperor. These are not merely formal inferences, such as the inference from (1) to either of the following:

19. Caesar is not *not* a prime number

20. There is at least one prime number.

Such inferences are licensed by the initial string as well, but deriving them does not require specific semantic knowledge of non-logical terms. By contrast, material inferences depend by definition on the meanings of their constituent terms; and material inferential reasoning exploits knowledge of this meaning – often along with broader worldly knowledge.

Of course, consisting of meaningful terms is not yet sufficient for expressing a thought. As the Dadaist examples remind us, a well formed syntactic structure is also a necessary condition for possessing an inferential role. And equally obviously, a string's syntax plays an essential role in determining the inferential role it does have. The two examples

11. Confusion now hath made his masterpiece
 21. His masterpiece now hath made confusion

license different inferences, even though they contain all the same words. At a minimum, then, to grasp a string's inferential role, one must understand, and exploit one's understanding of, both the meanings and the mode of combination of the words in that string.

In my view, consisting of meaningful words and being syntactically well formed are sufficient for a string to express a thought with a well defined inferential role – at least for strings with a relatively simple syntax. By the same token, understanding both the meaning and mode of combination of the words in such a string is sufficient for a thinker to grasp the inferential role of the thought it expresses.¹⁵ I simply cannot see in what way understanding both the meaning and the mode of combination of a string's constituent words could fall short of grasping its inferential role. But this is just what proponents of limited ranges of significance must think *is* possible. Even if they succeed in finding room for such a gap in principle, however, it seems clear that competent thinkers often can bridge that gap for simple cross-categorial predications, like (1) and (11), so as to draw material inferences from them. And if we can use these cross-categorial strings in material inferential reasoning, then they are not nonsense – although they are often quite absurd.

The proponent of categorial restrictions is likely to respond to all this by objecting that the seeming generability of material inferences only demonstrates that one can play a kind of empty game with words – a mere parody of understanding. Unless I can show on independent grounds that the initial string really does express a thought, I have not established that the rigmorole I cite as evidence counts as genuine reasoning. Rather, I have simply begged the question by assuming that it does. I think this broad objection might take two more specific forms, one in terms of the supposed thought itself, the other in terms of the thinker's supposed understanding of it.

First, one might worry that the inferences' conclusions suffer from just the same sort of categorial inappropriateness as the initial string, for

22. Caesar is not evenly divisible by a number other than 1 and himself
 is no less inappropriate than (1) in this regard. Moving from one string to

¹⁵ Peacocke, for one, accepts this claim as well: see *A Study of Concepts*, p. 43.

another is useless if we remain always within a closed circle of nonsense. However, as I have already shown, this inappropriateness will not be inherited by all of the initial string's inferential conclusions. (1) also licenses the conclusions that Caesar has no efficacy, and so also that he could not be an effective emperor. These latter thoughts are not cross-categorical.

It is true that the more absurd the initial predication, the more work it will take to generate categorially appropriate (let alone interesting) conclusions, and the weaker those inferential connections will be. Our potential understanding of the initial thought will be correspondingly less rich. Cross-categorical strings containing mathematical and scientific terms are thus particularly challenging, because the domains of concepts to which they are inferentially connected tend to be quite narrow. Sentences like (10) or (11) fare better in this regard, and the activity of drawing inferences from them seems less like a silly schoolbook exercise. However, as my discussion of (1) shows, even in the cross-categorical mathematical case there will be some available inferences that do not cross categorial boundaries.

The second form which the worry about emptiness might take is to object that one could play this pseudo-inferential word game without really understanding any of the concepts involved, just by following formal transformation rules. Generating inferences from a sentence like (1) in particular, whose subject term consists entirely of a name, might seem to require no more than treating the name as a free variable, and running inferences from the open sentence ' x is a prime number'.

With respect to this particular sort of case, two points should be made. First, if 'Caesar' really is functioning like a free variable, as many contemporary theories of names suggest, then it is not clear why there should be any obstacle to construing the whole sentence after all. It is only because the category 'man' is supposed to be built into the meaning of the name itself that the sentence counts as nonsensical. Secondly, and more importantly, when we consider sentences of the form 'The $F \phi s$,' then the suggestion of inferential one-sidedness is much weaker. Thus from a sentence like

23. The shadow struts and frets its hour upon the stage

one can conclude that something with no real substance expends energy on fruitless activity, and that something dark and derivative walks proudly. Both of these inferences require exploiting the inferential power of the entire sentence. Similarly, inferring from (1) that Caesar could not be an effective emperor involves bringing in specific information about the name's referent. So the defender of categorial restrictions needs to identify a specific way in which the understanding that is required to draw these inferences is more one-sided than that required for inferential reasoning more generally.

I do think that the worry in its more general form – the worry that the apparent activity of drawing material inferences in fact requires no more than following formal transformation rules – raises a serious problem for a ‘pure’ inferentialist view of concept possession. This is the view that grasping the contribution a concept F makes to the inferential role of thoughts is by itself sufficient for mastery of that concept. It does seem that some referential component is also essential for full understanding. (Such a referential requirement needs to be cashed with some care for concepts referring to abstract and otherwise causally distant entities. Perhaps in such cases it amounts to no more than the disposition to make certain judgements.) For instance, there is something seriously wrong with a thinker who grasps all the inferential implications of *being a car*, but who cannot recognize cars even in the bright lights of a car showroom. But I need not hold the view that grasping inferential role is all there is to concept possession. By hypothesis, the thinker under consideration, because supposedly otherwise competent with respect to the constituent concepts, does meet any such additional requirements for concept possession. The question before us is rather whether, if one really does understand F-thoughts *within* F’s normal range of application – whatever that may require – then one’s ability to do something that looks like drawing F-inferences outside that range ought to be counted as a capacity for genuine reasoning. I claim that it ought to be. At least, given that this activity has the *prima facie* appearance of reasoning, the proponent of categorial restrictions needs an independent argument to show it not to be so in fact. Further, given that normal thinkers can make inferences from cross-categorial strings which exploit not just analytic truths but also a wide range of worldly knowledge, it seems especially unlikely that performing merely formal transformations could enable someone to mimic the full range of a normal thinker’s inferential ability.

III(b). *Pragmatic evidence*

So far, I have argued that syntactically well formed, semantically cross-categorial strings do in general have inferential roles, and that thinkers otherwise competent can grasp them. Therefore we should count them as genuine sentences, and reject categorial restrictions on significance. But it is still natural to wonder why thinkers should *need* to grasp these inferential roles: why should this be a condition of competence with the constituent concepts? My answer is that we employ the inferential roles of such cross-categorial sentences in practical communicative contexts. Given this, a thinker who could not grasp those inferential roles would not manifest the sorts of conceptual capacities that we ordinarily do require thinkers to possess. One could still fix a minimal standard of competence for being able

to entertain the thought that *a* is F at all, one which required that thinkers must be able to combine *a* and F generally with their other concepts only within a certain range. But this minimal standard would not reflect the demands we ordinarily set for full conceptual competence.

How do we use such cross-categorical sentences? All the examples given above exemplify standard rhetorical devices – most prominently, metaphor. Of course, the fact that *we* mean something by these sentences does not show that the sentences themselves have that meaning, or any meaning at all. But the fact that they are used in this way does imply that the proponent of categorial restrictions needs a satisfactory account of how we manage to use them thus. The use of cross-categorical strings is not limited to metaphor: for instance, many metonymic sentences like

24. The front desk is getting anxious
25. The ham sandwich left without paying

are also cross-categorical. However, the most common use of cross-categorical strings is metaphorical. I shall therefore focus my attention here on showing that an adequate theory of metaphor will need to exploit the inferential role of the thought literally expressed by the sentence uttered. I shall explore what form a theory of metaphor that refused to exploit this inferential role would need to take.

No one who accepts that metaphors can communicate thoughts would deny that the hearer exploits the meanings of the words uttered in determining the thought(s) the speaker intended to communicate. To deny this would be to reduce metaphorical utterances to complex grunts. Rather, it seems, the proponent of categorial restrictions must insist that metaphorical interpretation relies on the meanings of the words alone, without the hearer's combining those meanings into a complete thought. That is, the theory must take the following general form: a hearer, confronted by a cross-categorical sentence 'Fa', is prevented by its nonsensicality from construing it any further, decides that it must be intended metaphorically, and begins straight away casting around for another related concept G to apply to *a* in lieu of F (or perhaps another concept *b* to apply F to), without applying F itself to *a*.

What is this theory to say about how the hearer arrives at the replacement concept G? The theory cannot appeal solely to the hearer's knowledge of what is involved in being F, because appropriate interpretation depends heavily upon what F is being applied *to*. Thus 'is the sun' gets interpreted very differently when it is applied metaphorically to Juliet from when it is applied to Achilles, or Louis XIV, or an atomic bomb. The theory must maintain that this constraint on F's replacement is produced simply through

the juxtaposition of *a* and F, because by hypothesis *a* and F cannot be combined. But this then makes it quite difficult for the theory to accommodate the role that syntactic structure also plays in interpretation.

First, only syntactically well formed strings can be used metaphorically. While Dadaist strings may be evocative, they are not metaphors; and while (10) is metaphorically interpretable,

26. But is shadow life a walking

is not. If juxtaposition were all that is required for metaphorical interpretation, these non-syntactic strings should be just as effective as their well formed counterparts.

Secondly, different syntactic structures determine different metaphorical interpretations for sentences that contain the same terms, as emerges from the above contrast between these two:

11. Confusion now hath made his masterpiece

21. His masterpiece now hath made confusion.

Thirdly, when the literal interpretation of a term is constrained by that term's role in the sentence's overall syntactic structure, then its metaphorical interpretation is constrained in the same way as well. So, for instance, the same weak crossover effects as constrain literal interpretation also prevent

27. He slew the dragon of Peter's greed

from being interpreted metaphorically to mean that Peter conquered his own vicious tendency.¹⁶

The proponent of categorial restrictions might admit that metaphorical interpretation exploits both the meanings and mode of combination of a string's constituent words, but insist that doing so still falls short of grasping the entire string's inferential role. However, again there seems to be little room in which to locate such a gap, especially for simple subject-predicate strings like (1), (8) or (10). And in more complex sentences, the syntax and semantics mutually constrain one another, and so in turn constrain both literal and metaphorical interpretation, in a way that makes it difficult to treat them in full isolation from one another. But unless there is something more to construing inferential role than grasping the sentence's constituent words' meanings and mode of combination together, then the account has already allowed that the hearer often does employ the inferential role of the thought expressed by 'Fa' in arriving at the replacement thought Ga.

Another difficulty is that not all metaphors are categorially inappropriate:

28. The rock is becoming brittle with age

¹⁶ Cf. also J. Stern, *Metaphor in Context* (MIT Press, 2000), ch. 2.

said of an eminent but doddering professor emeritus shows that ‘whole-sentence’ metaphors are often semantically unimpeachable.¹⁷ Some metaphors are even literally true, like

29. No man is an island
30. Anchorage is a cold city.¹⁸

Others, like

31. Sam is a gorilla

fall somewhere between the extremes of categorial propriety and absurdity. Thus even if proponents of limited ranges of significance do find a way to exclude inferential role from the theory of metaphorical interpretation for cross-categorial strings, they still face a difficult choice. On the one hand, they might simply deny that the inferential role of categorially impeccable sentences plays any role in construing them metaphorically. If so, they thereby abjure obvious, and apparently relevant, explanatory resources. On the other hand, they might deny that metaphors form a unitary kind of utterance, generated and understood along the same general principles. If so, they need to explain why the process of construing seems to be so similar across categorially appropriate and inappropriate metaphors.

A third option, and the one usually taken by proponents of categorial restrictions, is to adopt a non-cognitivist theory of metaphor across the board. On such a theory, no distinctive thought is communicated by a metaphorical utterance: there is just the arousal of more or less delicate and nuanced feelings and insights. Thus Davidson claims that a metaphor is like ‘a bump on the head’, or a drug. All three ‘nudge us into noting’ surprising aspects of the world, by *causing* us to ‘see’ one thing ‘as’ another.¹⁹ However, the non-cognitivist must deny the essential fact that in speaking metaphorically we do undertake speech acts, such as assertions and requests, which commit us to determinate cognitive contents that are distinct from but bear systematic relations to what is literally said.²⁰ This may not be all there is to metaphor. I think Davidson is right that metaphor often also involves a richer non-propositional understanding which we might well describe in terms of

¹⁷ M. Reddy, ‘A Semantic Approach to Metaphor’, in R.I. Binnick *et al.* (eds), *Papers from the Fifth Regional Meeting of the Chicago Linguistics Society* (Dept of Linguistics, Univ. of Chicago, 1969), pp. 240–51.

¹⁸ T. Cohen, ‘Figurative Speech and Figurative Acts’, *Journal of Philosophy*, 71 (1975), pp. 669–84, at p. 679.

¹⁹ Davidson, ‘What Metaphors Mean’, repr. in his *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984), pp. 245–64, at p. 253.

²⁰ See, e.g., J. Searle, ‘Metaphor’, in A. Ortony (ed.), *Metaphor and Thought* (Cambridge UP, 1979), pp. 92–123; M. Bergmann, ‘Metaphorical Assertions’, in S. Davis (ed.), *Pragmatics* (Oxford UP, 1991), pp. 485–94.

‘seeing as’. But the communication of content is one important part of why we use sentences metaphorically.

To insist that grasping inferential role is a necessary condition for interpreting utterances metaphorically is not to maintain that it is sufficient for doing so. Indeed, grasping inferential role is only the starting-point for construing metaphor, as for pragmatics more generally. Successful interpretation also exploits heavily context-dependent and affective associations, and requires imagination and ingenuity. For this reason, failure to interpret a particular metaphorical utterance as the speaker intended does not itself indicate a lack of semantic or conceptual competence. Someone could even be deaf to metaphor across the board without being conceptually impoverished, so long as he did grasp the uttered sentences’ literal meanings. Rather the point is that hearers normally do succeed in arriving at the intended metaphorical interpretations of cross-categorical sentences, and this requires grasping the sentences’ inferential roles. Further, given that as speakers we do routinely use such sentences metaphorically, it follows that we expect hearers to have this ability. In our ordinary practice, we do impose an unrestricted version of the generality constraint on our interlocutors.

IV. THOUGHTS AND TRUTH-CONDITIONS

Suppose you agree with me so far, first, that cross-categorical strings often do possess inferential roles, so that there is often something there at the intersection of heterogeneous concepts to be grasped; secondly, that speakers and hearers do routinely use such strings’ inferential roles in communication, so that someone who could not grasp such inferential roles would be conceptually impoverished. Still, you might wonder, why should having an inferential role be sufficient for a string to express a genuine thought, and why should what speakers and hearers manage to do with a string’s inferential role count as grasping a thought? Proponents of categorial restrictions are likely to insist that having propositional content or truth-conditions is the real criterion of significance; and so, *a fortiori*, that grasping truth-conditions, not just inferential role, is the real criterion of understanding. And this is just what they maintain we cannot do for cross-categorical strings.

But they need to insist on more than just our inability to fix truth-conditions for cross-categorical strings. They need to maintain that these strings cannot be assessed as either true or false: that is the *sine qua non* of nonsensicality. In general, though, cross-categorical strings appear to have all too obvious truth-conditions and truth-values. It seems obvious that the condition for being a prime number is being a number that is divisible only

by 1 and itself, and it seems equally obvious that Caesar fails to satisfy this condition, by virtue of not being a number at all. Unless we insist on specifying terms' meanings by distinct truth- and falsity-conditions, as Carnap for instance does,²¹ many cross-categorical strings will turn out to have clear truth-conditions and truth-values after all.

We have as yet no independent reason for specifying meanings in this way. It is almost always possible, and more straightforward, to follow the Fregean model, fixing necessary and sufficient conditions for truth alone, and stipulating that the sentence is false otherwise.²² Likewise, if one is operating with the lexical semantics model, one can stipulate that incompatibility of semantic markers guarantees falsity, rather than depriving the sentence of a truth-value altogether.²³ So lack of truth-conditions and truth-values for syntactically well formed semantically cross-categorical strings is not forced upon us by the semantics of natural language itself.

Instead, the real barrier to assigning full-blown truth-conditions to such strings seems to be the presupposition that what little sense we can muster of the strings' truth-conditions is too thin to count as genuine understanding. We are unable to imagine or conceive of any scenario in which the supposed thought could be true, and so we cannot even get started on investigating whether the actual world instantiates such a scenario. This is certainly true for many cross-categorical strings. But if the criterion for possessing truth-conditions is our ability to imagine a verifying situation, then quite a wide variety of sentences will end up counting as nonsensical: tautologies, mathematical conjectures, and at least certain counterfactuals, statements describing situations that violate physical laws, and even hypotheses of current scientific theory – none of which need involve cross-categorical predication.²⁴ If there is a problem about truth-conditional meaning, then all these sorts of sentences share it as well. Thus unless the proponents of limited ranges of significance can articulate a more restricted difficulty with cross-categorical sentences, the range of nonsense will mount higher than most people would now be willing to accept, and in unexpected locales.

²¹ Carnap claims that '*a* is a prime number' is false iff *a* is divisible by a natural number different from *a* and 1: 'The Elimination of Metaphysics through the Logical Analysis of Language', in A.J. Ayer (ed.), *Logical Positivism* (New York: Free Press, 1959), pp. 60–81, at p. 68.

²² See, e.g., Frege, 'Function and Concept', repr. in M. Beaney (ed.), *The Frege Reader* (Oxford: Blackwell, 1997), pp. 130–48, at p. 141; see also D. Lewis, 'General Semantics', in D. Davidson and G. Harman (eds), *Semantics of Natural Language* (Dordrecht: Reidel, 1972), pp. 169–218, at p. 179; and W.V. Quine, *Word and Object* (MIT Press, 1960), p. 229.

²³ The truth-conditions can be derived from the constituent lexical entries in several ways: see, e.g., U. Weinreich, 'Explorations in Semantic Theory', in T.A. Sebeok (ed.), *Current Trends in Linguistics*, Vol. III, *Theoretical Foundations* (The Hague: Mouton, 1966), pp. 395–477.

²⁴ See S. Yablo, 'Is Conceivability a Guide to Possibility?', *Philosophy and Phenomenological Research*, 53 (1993), pp. 1–42, for discussion of thoughts whose propositional contents may be believable (and even true) without being conceivable.

One might also object that because thoughts as discussed here are abstract objects rather than particular psychological states, therefore all considerations about our understanding must be irrelevant: questions about whether a string expresses a thought must be settled through metaphysical investigation instead. I think this objection misconstrues the sense in which thoughts and concepts are abstract. They are individuated by their possession-conditions, and in this sense they are abstractions from particular psychological states. But those possession-conditions are themselves informed by and responsive to our understanding and practices. An ideal marriage is also an abstract type, but it would be absurd to hypothesize about what constituted an ideal marriage without considering how people actually live. Likewise, the English alphabet is a set of abstract types, but its individuation depends on our actual language.

In the absence of a difficulty specific to cross-categorical strings, I conclude that there is no good reason to impose categorial restrictions on significance. Because speakers do regularly employ syntactically well formed semantically cross-categorical strings in order to communicate, hearers need to be capable of comprehending the thoughts these strings express. Otherwise those hearers will fail to be full participants in the game of thinking and talking. Therefore we have good reason to require that a fully competent thinker must be able to make at least some sense of even wildly heterogeneous combinations of concepts.

It is indeed true that people who seem to be generally competent thinkers (philosophers especially) do sometimes say that they can make no sense of a certain thought, that they find it unintelligible or incoherent. Often, however, when people say that they find a thought Fa unintelligible, they mean that they cannot believe that a speaker might really believe Fa , or that they cannot figure out what might reasonably be communicated by saying ' Fa '. But the pragmatic absurdity of saying something does not itself imply the semantic nonsensicality of what is said. Some speakers who say that a thought is unintelligible may mean only that it is pragmatically absurd. Others may really mean that it is nonsensical, but have concluded this – inappropriately, I maintain – on the basis of pragmatic evidence.

There is also plenty of genuine nonsense, although from sources other than semantic cross-categoricity. Some expressions, like 'divided by', are only partially defined, but for quite special reasons. Some expressions having to do with semantics, like 'is true', generate paradoxes when combined with first-order phrases in the language; this may be good reason to restrict their range of application. Demonstratives and names can fail to secure referents, and thereby generate the mere illusion of thought. And even some syntactically well formed strings containing only fully defined terms can still

fail to express thoughts. Syntactic complexity combined with pervasive but unsystematic semantic cross-categoricity will produce strings from which few non-formal inferences can be drawn, and which cannot be put to pragmatic use.²⁵ However, the nonsensicality in all of these cases stems from something more than just the crossing of semantic categories. Thus these special considerations do not support general categorial restrictions on the generality constraint of the sort that Evans, Strawson and others impose.

V. NORMAL RANGES OF APPLICATION AND THE IDEAL OF COMPLETE UNDERSTANDING

The impetus to impose categorial restrictions on the generality constraint comes in large part, I think, from a desire to maintain a robust model of the sort of understanding involved in genuine thought, without making the constraint so onerous that we all end up counting as incompetent thinkers. The hope is that by requiring such a robust understanding only within a certain realm, we can guarantee that most of us actually meet its high standard.²⁶ By contrast, I think this standard is unattainable in any case, so that it should be treated as an ideal for understanding across the board. We often fall short of this ideal, in a variety of ways and especially for cross-categorial predications. But I also think that we can acknowledge these failures as such without concluding that they undermine our capacity to think the relevant thoughts altogether.

Our understanding of cross-categorial thoughts is indeed both thinner than, and dependent upon, our understanding of their intra-categorial cousins. We usually first learn a concept by having it applied for us within some paradigmatic range; and it would be difficult to acquire new concepts from anomalous applications of them. It will not be surprising if we fail, at least at first, to recognize instances of a concept outside its paradigmatic range. The seriousness with which we regard someone's failure to exhibit the appropriate (inferential and judgemental) dispositions in a given case will thus depend on how far removed that thought is from the normal ranges of its concepts' applications.

But at the same time, our competence with concepts is in general a matter of degree, even within their paradigmatic ranges. Any given thinker

²⁵ An example might be 'The orbited candle would have been imposing a sharpened carpet's fourteenth copper gesture, but insignificance elects the first folder time'. Some sentences may also fail to express *complete* thoughts: for example, 'Steel is strong enough' – for what? See K. Bach, 'Speaking Loosely', *Midwest Studies in Philosophy*, 25 (2001), pp. 249–63.

²⁶ However, Evans does himself admit that the generality constraint is an unattainable ideal: see *The Varieties of Reference*, p. 105.

is likely to lack the dispositions to make all and only the appropriate inferences associated with being F, or to pick out all and only the right objects as instantiations of it. Usually, if he possesses both the conceptual resources that are necessary to think at all, and a sufficiently rich body of dispositions for applying the particular concept F, then we treat him as capable of grasping the thought *Fa*, despite his lack of full competence with the constituent concepts. So too, given the appropriate background, we usually accept a comparatively thin understanding of some particular sentence '*Fa*' as sufficient for putting a thinker 'in touch with' the thought it expresses. Thus while full competence with a concept F does require the ability to combine F with one's other concepts even outside F's paradigmatic range, it is also true both that a relatively thin understanding of those thoughts may suffice for competence, and that failure to grasp any one of them need not undermine one's ability to think F-thoughts altogether.

In the light of these ways in which we often, even usually, fall short of the ideal of complete understanding, the following line of objection is possible. One might accept my arguments for a fully unrestricted generality constraint as a condition for full competence, but also insist on delimiting a restricted generality constraint as a condition for minimal conceptual competence. On this view, getting into the business of thinking F-thoughts at all requires just the ability to apply the concept F generally within its paradigmatic range, but an inability to apply F outside that range still counts against full mastery of the concept. This move would capture the structured systematicity of genuine propositional thought without making the requirement too restrictive. The discussion of §II, as well as independent syntactic constraints on the individuation of semantic kinds, suggests that the relevant categories of application will be rather messy. In addition, evidence about 'prototype effects' suggests that being paradigmatic is a matter of degree rather than of kind.²⁷ Thus I am suspicious that the relevant categories could be carved out in a useful, clearly delimited way. However, in principle such a move is available.

The important point is that we should not carve off a sharply defined area within which we can insist on a rich understanding as the mark of minimal competence, but outside which lies no thought at all. Scientific and mathematical progress, for instance, consists at least sometimes and in part in the formulation of hypotheses which are only minimally understood, and which sometimes seem nonsensical from the perspective of current theory. The thoughts that light is both particle and wave, that unconscious thoughts can cause actions, and that mental states are brain states have all counted as

²⁷ See, e.g., E. Rosch, 'Principles of Categorization', in E. Rosch and B.B. Lloyd (eds), *Cognition and Categorization* (Hillsdale: Erlbaum, 1978), pp. 27–48.

cross-categorical nonsense in someone's book. If we count such hypotheses as nonsensical, or scientists' and mathematicians' groping understanding as no understanding at all, then we fail to account for how those investigators could have proceeded with their enquiry, except by viewing them as filled with mystical inspiration. We also thereby commit ourselves to the view that if their hypotheses are eventually accepted, the nonsensical is suddenly transformed into the necessarily true. But this seems absurd. Kripke and Putnam have shown us, if nothing else, that our *a priori* grip on necessity and possibility is considerably more slippery than we once thought.²⁸

The most plausible examples of syntactically simple cross-categorical nonsense do involve mathematical and technical scientific concepts. This is because the range of further concepts to which they are inferentially connected, and the domain of objects to which they apply truly, are both well defined and discrete. Wittgenstein calls the symbolism of chemistry and the notation of the infinitesimal calculus 'the suburbs of our language', and offers us this image:

Our language can be seen as an ancient city: a maze of little streets and squares, of old and new houses, and of houses with additions from various periods; and this surrounded by a multitude of new boroughs with straight regular streets and uniform houses.²⁹

I agree that our concepts have natural homes, that is, normal ranges of significance: in these neighbourhoods they find their richest application. What I am concerned to resist is the claim that these neighbourhoods and suburbs are, in general, gated communities or ghettos, so that a concept from one area is in principle barred from commerce in another, or can travel to it only in disguise. It can still be admitted that when concepts do travel far from home, they become more tentative. And it can also be admitted that some suburbs are more disconnected from the rest, and thus that some concepts have more difficulty than others in travelling far from home.³⁰

Society of Fellows, Harvard University

²⁸ S. Kripke, *Naming and Necessity* (Harvard UP, 1980); H. Putnam, 'The Meaning of "Meaning"', repr. in his *Mind, Language, and Reality* (Cambridge UP, 1975), pp. 215–71.

²⁹ Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1958), §18.

³⁰ I have received extremely generous and substantive assistance in developing this paper. I especially want to thank Kent Bach, Bill Brewer, Cheryl Chen, Jeff King, John MacFarlane, Axel Mueller, Chris Pincock, John Searle, Barry Stroud and Dmitri Tymoczko for suggestions which led to significant alterations in the argument. Anonymous comments from referees were also helpful.

Copyright of Philosophical Quarterly is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.