

### **The Generality Constraint, Nonsense, and Categorical Restrictions**

I argue that we should not adopt categorical restrictions on the significance of syntactically well-formed strings. Even syntactically well-formed but semantically absurd strings, such as ‘Life is but a walking shadow’ and ‘Caesar is a prime number’, can express thoughts; and competent thinkers both can and ought to be able to grasp such thoughts. A more specific way of putting this claim is that Gareth Evans’ Generality Constraint should be viewed as a fully general constraint on concept possession and propositional thought, even though Evans himself accepted only a categorially-restricted version of the Constraint. I establish this by arguing, first, that even well-formed but semantically cross-categorical strings often do possess substantive inferential roles; second, that hearers exploit these inferential roles in interpreting such strings metaphorically; and third, that there is no good reason to deny truth-conditions to strings with inferential roles.

#### **1. *The Generality Constraint***

This paper concerns the limits of propositional thought, and the requirements on comprehension that are imposed by competence with respect to a given concept. Propositional thoughts are thoughts reportable by ‘that’-clauses: for instance, the thought that there’s beer in the refrigerator. In the terms I will 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 a particular thinker meets those conditions, and so whether she grasps 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 each other 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’.<sup>1</sup>

---

<sup>1</sup> Gareth Evans, *The Varieties of Reference*, ed. John McDowell (Oxford: Clarendon Press, 1982), p. 100.

This fact about propositional thought is so basic that if someone fails to grasp that the thoughts he entertains are related in this way, we will question whether he really does understand them after all: someone who can see no connection between his thoughts that John is happy and that Harry is happy can't really be grasping either of those thoughts.

Generalizing from this example, we get a picture of concepts as the articulating strands of thought, the lines that 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*,....<sup>2</sup>

This picture in turn suggests a condition on which thoughts we ought to be able to understand: if the structure is truly systematic, it should contain no unmotivated 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 syntactic well-formedness). Evans calls this a “fundamental constraint” on “the nature of our conceivings,”<sup>3</sup> and 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.<sup>4</sup>

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 representings of other animals.<sup>5</sup> The recombining of our concepts helps to explain the rich generativity of our conceptual capacities. Further, if we accept that understanding a thought essentially involves grasping its truth-conditions (as Evans and many others do), 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's difficult to see how this could happen unless we employed our previous mastery of the thought's constituent concepts to determine what would make the new thought true.

---

<sup>2</sup> *ibid*, p. 104, fn. 21.

<sup>3</sup> *ibid*, p. 100.

<sup>4</sup> *ibid*, p. 104.

<sup>5</sup> Cf. Evans, “Semantic Theory and Tacit Knowledge,” reprinted in his *Collected Papers* (Oxford: Clarendon Press, 1985), pp. 322-342.

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 purported 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.<sup>6</sup> However, these barriers are not inherently conceptual in nature, and so don’t 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 can’t, and so shouldn’t be expected to, fit them together to form a complete thought. 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 purportedly expressed by

(1) Caesar is a prime number

or

(2) Colorless green ideas sleep furiously

to be true. And if we can’t 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 can’t grasp the conditions under which such purported thoughts would be true, but also that we can’t even properly assess them as false. Caesar, one might think, just isn’t the sort of thing to either succeed *or fail* at being prime number. Absurd ‘thoughts’ like these might seem to involve such serious category mistakes that the strings

---

<sup>6</sup> Peacocke raises the last two possibilities in *A Study of Concepts* (Cambridge, MA: MIT Press, 1992), p. 43.

‘expressing’ them ought to be counted as syntactically well-formed nonsense.<sup>7</sup> If this is right, then there is no thought *there* at the nexus of the constituent concepts to be understood. And if so, then the failure to grasp such nothingness obviously should not impugn anyone’s competence with respect to the relevant concepts.

In the light of these worries, philosophers like Strawson, Evans, and Peacocke, 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.”<sup>8</sup> 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.”<sup>9</sup> 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*.<sup>10</sup>

Concepts, they all agree, have limited ‘ranges of significance’ or ‘categories of appropriate application’. It is only within this range that there are indeed thoughts with genuine truth-conditions and truth-values to be understood, and only here that the Generality Constraint applies. Within this range, 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 will argue, against this consensus, that we should not impose categorial restrictions on either conceptual significance or conceptual competence. My argument proceeds in four steps. First, I show that the project of delimiting appropriate categories faces serious, though perhaps not insurmountable, difficulties. I adopt the least restrictive plausible categories of significance. I then argue that strings that

---

<sup>7</sup> Cf. Ryle: “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 supposals or conceptions are or are not absurd. For what is absurd is unthinkable.” (“Categories,” in *Logic and Language* (2<sup>nd</sup> series), ed. A. Flew (Oxford: Blackwell, 1953), pp. 65-81, p. 76.)

<sup>8</sup> Strawson, *Individuals: An Essay in Descriptive Metaphysics* (Garden City NJ: Anchor Books, 1963), p. 95; emphasis added.

<sup>9</sup> *The Varieties of Reference*, p. 101; emphasis added.

<sup>10</sup> *A Study of Concepts*, p. 42; emphasis added.

count as cross-categorial on this criterion, such as (1), often do possess substantive inferential roles. Such strings should therefore be counted as significant. Third, I argue that normal thinkers do routinely make use of these inferential roles, in particular in the process of metaphorical construal. Therefore, full competence requires that thinkers be capable of grasping these inferential roles. Finally, I claim that there is no good reason to deny that cross-categorial predications with inferential roles do also have truth-conditions.

In light of this argument, I conclude that we should abandon the project of delimiting a narrow range within which a robust understanding of every thought is both necessary and sufficient for competence with a given concept, but outside of which 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 to syntactically well-formed strings are indeed nonsense — albeit for reasons other than the violation of categorial restrictions. We will 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.

## **2. *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 purported thoughts outside this range really are nonsensical. 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 nice but basically gratuitous feat. Treating the criteria in this way remains a fallback position for now. (I argue against this weaker position in §3b.) However, all of the defenders of category restrictions mentioned above have advocated restrictions on significance rather than merely on competence; and the reasoning behind imposing such restrictions supports treating

them in this way. That is, it's supposed to be something about the concepts themselves that prevents our fitting them together, and so it's natural to think that the concepts simply *can't* be fitted together.

In either case, we still need to make explicit the criteria for determining which concepts can be combined. First, we want to understand just how and why concepts should be limited in their application, if indeed they are. Second, as theorists, we need a way of deciding whether apparent lapses of generality are limitations of thinkers' capabilities, or rather genuine limitations in 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 a purported 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."<sup>11</sup> An individual may be brought under multiple 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, just in case *F* is or implies a predicate that is "a priori rejectable" not just for one, but for all of *a*'s individuating designations.<sup>12</sup> So, the predicate 'is president 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.

---

<sup>11</sup> Strawson, "Categories," in *Ryle: A Collection of Critical Essays*, ed. O. Wood and G. Pitcher (Garden City, NY: Doubleday Anchor Books, 1970), pp. 181-211, p. 200.

<sup>12</sup> *ibid*, p. 203.

Although this method has some intuitive 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.<sup>13</sup> 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 give 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 concrete noun). Thus, the lexical entry for 'colorful' might be written as follows:<sup>14</sup>

- (1) *Colorful* → Adjective → (*Color*) → [*Abounding in contrast or variety of bright colors*] <(Physical Object) v (Social activity)>
- (2) *Colorful* → Adjective → (*Evaluative*) → [*Having distinctive character, vividness, or picturesqueness*] <(Aesthetic Object) v (Social activity)>;

while the lexical entry for 'man' is something like:

- (1) *Man* → Noun concrete → Noun masculine → (*Physical object*) → (*Human*) → (*Adult*) → (*Male*)
- (2) *Man* → Noun concrete → (*Physical object*) → (*Human*)

---

<sup>13</sup> See, e.g., Jerrold Katz and Jerry Fodor, "The Structure of a Semantic Theory," in *The Structure of Language*, eds. J. Fodor and J. Katz (Englewood Cliffs, NJ: Prentice Hall, 1964), pp. 479-518; see also Noam Chomsky, "Degrees of Grammaticalness," in *The Structure of Language*, pp. 384-389.

<sup>14</sup> "The Structure of a Semantic Theory," p. 507.

(3) *Man* → Noun abstract → (*Human*).<sup>15</sup>

(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.<sup>16</sup> That is, in order to draw categorial boundaries that do any real work, one ends up sorting things finely enough 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"<sup>17</sup>; and in this case the evidence for ambiguity is quite weak. First, speakers are not in general aware of much of the postulated ambiguity. Second, 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

---

<sup>15</sup> 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.

<sup>16</sup> For instance, Katz and Fodor conclude on the basis of their analysis that 'The man hits the colorful ball' exhibits a four-fold ambiguity. Similarly, Ryle concludes that 'existence' has at least two senses, "somewhat as 'rising' indicates two 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).

<sup>17</sup> Saul Kripke, "Speaker's Reference and Semantic Reference," reprinted in *Readings in the Philosophy of Language*, ed. Peter Ludlow (Cambridge, MA: MIT Press, 1997), pp. 382-414, p. 401.

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 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 our 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 overly restrictive, rendering a much broader swath of our thought and talk nonsensical than we had hoped. For instance,

(3) The man in black is quite a colorful guy

would count as meaningless on the basis of the lexical entries above. This seems like 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 for us. If we don’t opt for the most restrictive categories, then we will need a new criterion for deciding which categories mirror the especially fundamental sorts. And this would seem just to throw us back onto our initial, question-begging intuitions about whether it’s ‘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 our purposes they are mere matters of detail. I will 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 semantical approaches employ these categories, along with other, finer ones. Among ‘things’ broadly construed, I will 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.

### 3. *Inferential 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 argue that semantically cross-categorial, syntactically well-formed strings can be used in a range of ways that 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 don't 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've 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!*<sup>18</sup>
- (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.*<sup>19</sup>
- (11) *Confusion now hath made his masterpiece!*<sup>20</sup>
- (12) *Caesar is a prime number.*

---

<sup>18</sup> *Romeo and Juliet*, II.ii.1, reprinted in *The Riverside Shakespeare*, 2<sup>nd</sup> edition, ed. J. J. M. Tobin et al. (New York: Houghton Mifflin Co., 1997).

<sup>19</sup> *Macbeth*, V.v.26, reprinted in *The Riverside Shakespeare*.

<sup>20</sup> *Macbeth*, II.iii.66, reprinted in *The Riverside Shakespeare*.

I will 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-categorical strings can be to generate and comprehend; the considerations I adduce thus apply with even greater force to these cases.

### 3a. Semantic Evidence

Given how much we can do with such cross-categorical predications, it's important to remember that we can't, or won't, do any of this with syntactically malformed strings. Take 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 leafs the leaves.<sup>21</sup>

Each of these strings 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, however, 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 smell. Notice too that 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, notice that speakers can understand and answer only syntactically well-formed 'yes-no' questions:

(15) Could staggering earthworm fishies be clocks the cow?

---

<sup>21</sup> The slashes indicate line breaks. Both examples are from *The Dada Painters and Poets: An Anthology*, 2<sup>nd</sup> edition, ed. R. Motherwell (Belknap Press, Cambridge MA, 1979).

doesn't afford any answer, while

(16) Could Caesar be a prime number?

does. Upon hearing the latter, one might justifiably wonder about the speaker's intentions in posing such a question. But what the question asks, even if not why the speaker asks it, is clear. And it is equally clear, in ordinary talk at least, that the answer is 'no'. Caesar is not a prime number — as we usually construe it, *that* thought is necessarily and obviously true, precisely because Caesar isn't a number at all.

By contrast,

(17) Staggering earthworm fishies is not clocks the cow

is no more amenable to truth-evaluation than the initial, 'not'-less string was.

The fact that speakers have intuitions about the appropriate responses to and truth-values of such complex constructions provides provisional evidence that the initial cross-categorical strings, such as (1), are themselves significant: they class together with significant strings, and apart from mere word salad (and subsentential phrases) in this regard. It also raises the important question of how to provide a principled, and hopefully unified, treatment of both the complex and simple constructions. 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 un-negated thought itself has determinate truth-conditions and a truth-value. We might also have independent reason (though here the potential grounds are much less clear) for thinking that questions of the form 'Could *a* be *F*?' should in general be analyzed 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 in fact succeeded in combining the inferential roles of the constituent concepts together to determine the inferential role of the

thought as a whole.<sup>22</sup> A competent thinker can reason, for instance, from the hypothetical truth of Caesar's being a prime number to the conclusion that he's not evenly divisible by a number other than one and himself, or (more interestingly) to the conclusion that he's an abstract object, and therefore lacks causal efficacy. From this latter conclusion she 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

(19) Caesar is not *not* a prime number,

or to

(20) There is at least one prime number.

Such inferences are licensed by the initial string as well, but deriving them doesn't require specific semantic knowledge. 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 on possessing an inferential role. And equally obviously, a string's syntax plays an essential role in determining the inferential role it does have:

(11) Confusion now hath made his masterpiece!

and

(21) His masterpiece now hath made confusion!

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

---

<sup>22</sup> By 'inferential role' I mean the core set of inferences a thinker needs to be able to draw 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.

In my view, consisting of meaningful words and being syntactically well-formed is 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.<sup>23</sup> I simply can't 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 the proponent of limited ranges of significance must think *is* possible. Even if he succeeds in finding room for such a gap in principle, however, it seems clear that competent thinkers often can bridge that gap for simple cross-categorical predications, like (1) and (11), to draw material inferences from them. And if we can use these cross-categorical strings in material inferential reasoning, then they are not nonsense — although they are often quite absurd.

The proponent of categorical 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 haven't established that the rigamarole I cite as evidence counts as genuine reasoning. Rather, I've 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 purported thought itself, the other in terms of the thinker's purported understanding of it.

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

(22) Caesar is not evenly divisible by a number other than 1 and himself

is no better than (1) in this regard. Moving from one string to another is useless if we remain always within a closed circle of nonsense. However, as we've already seen, this inappropriateness will not be inherited by all of the initial string's inferential conclusions. (1) also licenses the conclusions that Caesar

---

<sup>23</sup> Peacocke, for one, accepts this claim as well; see *A Study of Concepts*, p. 43.

has no causal efficacy, and so 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 far better in this regard, and the activity of drawing inferences from them seems less like a silly schoolbook exercise. However, as our 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 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 in a manner akin to a free variable, as most contemporary theories of names suggest, then it's not clear why there should be any obstacle to construing the whole sentence after all. It's only because the category 'man' is supposed to be built into the meaning of the name itself that the sentence counts as nonsensical. Second, and more importantly, when we consider sentences of the form 'The F  $\Phi$ s,' then the intuition of inferential one-sidedness is much weaker. Thus, from a sentence like

(23) The shadow struts and frets

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.<sup>24</sup> For instance, there is something seriously wrong with a thinker who grasps all the inferential implications of *being a car*, but who can’t recognize cars even in the bright light of an automotive showroom. But I needn’t hold the view that grasping inferential role is all there is to concept possession. By hypothesis, the thinker under consideration does meet any such additional requirements on concept possession, because she is supposed to be otherwise competent with respect to the constituent concepts. 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 of that range ought to be counted as a capacity for genuine reasoning. I claim that it ought to. 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. 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.

### **3b. Pragmatic Evidence**

So far, I’ve argued that syntactically well-formed, semantically cross-categorial strings do in general have inferential roles, and that otherwise competent thinkers can grasp them. Therefore, we

---

<sup>24</sup> Such a referential requirement needs to be cashed out 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 judgments.

should count such strings as genuine sentences, and reject categorial restrictions on significance. But it's 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 do ordinarily 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 a thinker be able to combine *a* and *F* generally with her 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-categorial sentences? All of the examples given above exemplify standard rhetorical techniques — most prominently, metaphor. Of course, the fact that *we* mean something by these sentences doesn't show that the sentences themselves have that meaning, or any meaning at all. But the fact that they are intuitively used in this way does imply that the proponent of categorial restrictions needs a satisfactory account of how we manage to so use them. The use of cross-categorial strings is not limited to metaphor; for instance, many metonymic sentences, like

(24) The front desk is getting anxious

and

(25) The ham sandwich left without paying,

are also cross-categorial. However, the most common use of cross-categorial strings is metaphorical. I will therefore focus my attention here on showing that an adequate theory of metaphor will need to exploit the inferential role of the thought that is literally expressed by the sentence uttered. To show this, I'll 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 combining those meanings into a complete thought. That

is, his theory must take the following general form: when the hearer is confronted by a cross-categorical sentence ‘ $Fa$ ’, she 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  $F$ ,<sup>25</sup> 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’s applied metaphorically to Juliet from when it’s so applied to Achilles, or Louis XIV, or an atom 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, notice that only syntactically well-formed strings can be used metaphorically. While our 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 were required for metaphorical interpretation, these non-syntactic strings should be just as effective as their well-formed counterparts. Second, different syntactic structures determine different metaphorical interpretations for sentences that contain the same terms: contrast, once again,

(11) Confusion now hath made his masterpiece!

with

(21) His masterpiece now hath made confusion!

Third, 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 that constrain literal interpretation also prevent

---

<sup>25</sup> Or perhaps, another concept  $b$  to apply  $F$  to. The process is presumably similar in the relevant respects.

(27) He slew the dragon of Peter's greed

from being interpreted metaphorically to mean that Peter conquered his own vicious tendency.<sup>26</sup>

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, once 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*.

A further difficulty for the theory is that not all metaphors are categorially inappropriate. Whole-sentence metaphors are often semantically unimpeachable, such as

(28) The rock is becoming brittle with age,

said of an eminent but dottering professor emeritus.<sup>27</sup> Some metaphors are even literally true, like

(29) No man is an island,

and

(30) Anchorage is a cold city.<sup>28</sup>

Others, like

(31) Sam is a gorilla,

fall somewhere between the extremes of categorial propriety and absurdity. Thus, even if the proponent of limited ranges of significance does find a way to exclude inferential role from his theory of

---

<sup>26</sup> Cf. also Josef Stern, *Metaphor in Context*, (Cambridge, MA: MIT Press, 2000), ch. 2.

<sup>27</sup> Michael Reddy, "A Semantic Approach to Metaphor," in *Papers from the Fifth Regional Meeting of the Chicago Linguistics Society*, ed. R. I. Binnick et al (Department of Linguistics, University of Chicago, 1969), pp. 240-251.

<sup>28</sup> Ted Cohen, "Figurative Speech and Figurative Acts," *Journal of Philosophy* 71 (1975), pp. 669-684, p. 679.

metaphorical interpretation for cross-categorical strings, he still faces a difficult choice. On the one hand, he might simply deny that the inferential role of categorially impeccable sentences plays any role in construing them metaphorically. If so, he thereby abjures obvious, and apparently relevant, explanatory resources. On the other hand, he might deny that metaphors form a unitary kind of utterance, generated and understood along the same general principles. If so, he needs to explain why the process of construal seems to be so similar across categorially appropriate and inappropriate metaphors.

A third option, and the one classically 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's just the arousal of more or less delicate and nuanced feelings and insights. Thus, Donald Davidson claims that a metaphor is like "a bump on the head," or a drug. All three "nudge us into noting"<sup>29</sup> surprising aspects of the world, by *causing* us to 'see' one thing 'as' another. 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.<sup>30</sup> This may not be all there is to metaphor. I think Davidson is right that metaphor often also involves a richer, non-propositional understanding that we might well describe in terms of '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 metaphorical construal, as for pragmatic construal 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

---

<sup>29</sup> Davidson, "What Metaphors Mean," reprinted in his *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984), pp. 245-264, p. 253.

<sup>30</sup> See, e.g. John Searle, "Metaphor," in *Metaphor and Thought*, ed. A. Ortony (Cambridge: Cambridge University Press, 1979), pp. 92-123; Merrie Bergmann, "Metaphorical Assertions," in *Pragmatics*, ed. S. Davis (Oxford: Oxford University Press, 1991), pp. 485-494.

deaf to metaphor across the board without being conceptually impoverished, so long as she 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 that this requires grasping the uttered 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.

#### **4. *Thoughts and Truth-Conditions***

Now, suppose that you agree with me so far: first, that cross-categorical strings often do possess inferential roles, and so that there is often something there at the intersection of heterogeneous concepts to be grasped; and second, that speakers and hearers do routinely use such strings' inferential roles in communication, and 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? The proponent of categorical restrictions is 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 he maintains we cannot do for cross-categorical strings.

Notice, though, that the proponent of categorical restrictions needs to insist on more than just our inability to fix truth-conditions for cross-categorical strings. He needs to maintain that these strings cannot be assessed as either true *or* false: that's the *sine qua non* of nonsensicality. But in general, cross-categorical strings appear to have all too obvious truth-conditions and truth-values. It seems intuitively obvious that the condition for being a prime number is being a number that's evenly 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,<sup>31</sup> many cross-categorial strings will turn out to have clear truth-conditions and -values after all.

We don't yet have an independent motivation for specifying meanings in this way, though. It is quite possible, and more straightforward, to follow the Fregean model, fixing necessary and sufficient conditions on truth alone and stipulating that the sentence is false otherwise.<sup>32</sup> Likewise, if one is operating with the lexical semantics model, one can stipulate that an incompatibility of semantic markers guarantees falsity, rather than depriving the sentence of a truth-value altogether.<sup>33</sup> So a lack of truth-conditions and -values for syntactically well-formed, semantically cross-categorial 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 antecedent intuition 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 a scenario in which the purported thought could be true, and so we cannot even get started investigating whether the actual world instantiates this scenario. This is certainly true for cross-categorial strings. But if the criterion for possessing truth-conditions is our ability to generate a rich imaginative conception of 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-

---

<sup>31</sup> Thus, Carnap claims that “‘ $a$  is a prime number’ is false if and only if  $a$  is divisible by a natural number different from  $a$  and 1.” “The Elimination of Metaphysics Through the Logical Analysis of Language,” in *Logical Positivism*, ed. A. J. Ayer (New York: Free Press, 1959), pp. 60-81, p. 68.

<sup>32</sup> See, e.g. Frege, “Function and Concept,” reprinted in *The Frege Reader*, ed. M. Beaney (Oxford: Blackwell, 1997), pp. 130-148, p. 141; see also David Lewis, “General Semantics,” in *Semantics of Natural Language*, eds. D. Davidson and G. Harman (Dordrecht, Holland: D. Reidel, 1972), pp. 169-218, p. 179, and W.V.O. Quine, *Word and Object*, (Cambridge, MA: MIT Press, 1960), p. 229.

<sup>33</sup> The truth-conditions can be derived from the constituent lexical entries in several ways. The simplest method (Uriel Weinreich, “Explorations in Semantic Theory,” in *Current Trends in Linguistics*, vol. 3, *Theoretical Foundations*, ed. T. A. Sebeok (The Hague: Mouton, 1966), pp. 395-477) is to combine terms by transferring all markers and distinguishers from the entries for the subsidiary terms to the higher-order phrasal entry, and then eliminating redundant features in the new entry.

categorial predication.<sup>34</sup> If there is a problem about truth-conditional meaning, then all these sorts of sentences share it as well. Thus, unless the proponent of limited ranges of significance can articulate a more restricted difficulty with cross-categorial sentences, the range of nonsense will mount higher, and in different locales, than most people would now be willing to accept.

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. Thoughts and concepts are individuated by their possession conditions, and in this sense they are abstractions from individuals' 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 appealing in some way to considerations about 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-categorial strings, I conclude that there is no good reason to impose categorial restrictions on significance. And because speakers do regularly employ syntactically well-formed, semantically cross-categorial strings 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 will 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* to be

---

<sup>34</sup> See Steve Yablo, "Is Conceivability a Guide to Possibility?" (*Philosophy and Phenomenological Research* 53:1 (March 1993), pp. 1-42) for discussion of thoughts whose propositional contents may be believable (and even true) without being conceivable.

unintelligible, they mean that they can't believe that a speaker might really believe *Fa*, or that they can't figure out what he might reasonably want to communicate 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 indeed only partially defined, but for quite special reasons. Specifically, some semantic expressions, like 'is true', generate paradoxes when combined with first-order phrases in the language; this may be good reason 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.<sup>35</sup> 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.

### ***5. 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 the 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.

---

<sup>35</sup> 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 relatively simple, well-formed sentences also fail to express *complete* thoughts: for example, 'Steel is strong enough' — for what? See Kent Bach, "Speaking Loosely," *Midwest Studies in Philosophy* 25 (2001), pp. 249-263.

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.<sup>36</sup> By contrast, I think this standard is unattainable in any case, and 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 awfully 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 of its paradigmatic range. The seriousness with which we regard someone's failure to exhibit the appropriate (inferential and judgmental) 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 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 that thinker possesses both the rich 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 her as capable of grasping the thought  $Fa$ , despite her 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

---

<sup>36</sup> Notice, however, that Evans does himself admit that the Generality Constraint is an unattainable ideal: see *Varieties of Reference*, p. 105.

understanding of those thoughts may suffice for competence, and that failure to grasp any one of these thoughts needn't undermine one's ability to think *F*-thoughts altogether.

In light of these ways in which we often, even usually, fall short of the ideal of complete understanding, someone might make the following move. He might accept my arguments for a fully general Generality Constraint as a condition on full competence, but also insist on delimiting a restricted Generality Constraint as a condition on 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 of 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 overly restrictive. The discussion of §2, as well as independent syntactic constraints on the individuation of semantic kinds,<sup>37</sup> suggests that the relevant categories of application will be rather messy. In addition, evidence about 'prototype effects' suggests that paradigmaticity is a matter of degree rather than of kind.<sup>38</sup> 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 of which lies no thought at all. Scientific and mathematical progress, for instance, consist at least sometimes and in part in the formulation of hypotheses that are only minimally understood, and that 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 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 inquiry except by viewing them as filled with mystical inspiration. We also thereby commit ourselves to the view that if their hypotheses are eventually

---

<sup>37</sup> See Kent Johnson, "How Systematic are Language and Thought?" (manuscript).

<sup>38</sup> See e.g. Eleanor Rosch, "Principles of Categorization," in *Cognition and Categorization*, ed. E. Rosch and B. B. Lloyd (Hillsdale, NJ: Erlbaum, 1978), pp. 27-48.

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.<sup>39</sup>

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 different periods; and this surrounded by a multitude of new boroughs with straight regular streets and uniform houses.<sup>40</sup>

I agree that our concepts have natural homes, that is, normal ranges of significance; in those neighborhoods they find their richest application. What I am concerned to resist is the claim that these neighborhoods 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 traveling far from home.<sup>41</sup>

---

<sup>39</sup> Saul Kripke, *Naming and Necessity* (Cambridge, MA: Harvard University Press, 1980); Hilary Putnam, “The Meaning of ‘Meaning’,” in *Mind, Language, and Reality* (Cambridge: Cambridge University Press, 1975), pp. 215-271.

<sup>40</sup> Ludwig Wittgenstein, *Philosophical Investigations*, tr. G. E. M. Anscombe (Oxford: Basil Blackwell, 1972), §18.

<sup>41</sup> I have received extremely generous and substantive assistance in developing this paper. I especially want to thank Kent Bach, Bill Brewer, Jeff King, John MacFarlane, Chris Pincock, John Searle, Barry Stroud, and Dmitri Tymoczko for suggestions that led to significant alterations in the argument. Anonymous comments from two referees for *Philosophical Quarterly* were also helpful.

## Bibliography

- Bach, Kent: "Speaking Loosely," *Midwest Studies in Philosophy* 25 (2001), pp. 249-263.
- Bergmann, Merrie: "Metaphorical Assertions," in *Pragmatics*, ed. S. Davis (Oxford: Oxford University Press, 1991), pp. 485-494.
- Carnap, Rudolf: "The Elimination of Metaphysics Through the Logical Analysis of Language," in *Logical Positivism*, ed. A.J. Ayer (New York: Free Press, 1959), pp. 60-81.
- Chomsky, Noam: "Degrees of Grammaticalness," in *The Structure of Language*, eds. J. Fodor and J. Katz (Englewood Cliffs, NJ: Prentice Hall, 1964), pp. 384-389.
- Cohen, Ted: "Figurative Speech and Figurative Acts," *Journal of Philosophy* 71 (1975), pp. 669-684.
- Davidson, Donald: "What Metaphors Mean," in *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984), pp. 245-264.
- Evans, Gareth: *The Varieties of Reference*, ed. John McDowell (Oxford: Clarendon Press, 1982).
- : "Semantic Theory and Tacit Knowledge," reprinted in his *Collected Papers* (Oxford: Clarendon Press, 1985), pp. 322-342.
- Frege, Gottlob: "Function and Concept," reprinted in *The Frege Reader*, ed. Michael Beaney (Oxford: Blackwell, 1997), pp. 130-148.
- Johnson, Kent: "How Systematic are Language and Thought?" (manuscript).
- Katz, Jerrold and Fodor, Jerry: "The Structure of a Semantic Theory," in *The Structure of Language*, eds. Jerry Fodor and Jerrold Katz (Englewood Cliffs, NJ: Prentice Hall, 1964), pp. 479-518.
- Kripke, Saul: *Naming and Necessity* (Cambridge, MA: Harvard University Press, 1980).
- : "Speaker's Reference and Semantic Reference," reprinted in *Readings in the Philosophy of Language*, ed. Peter Ludlow (Cambridge, MA: MIT Press, 1997), pp. 382-414.
- Lewis, David: "General Semantics," in *Semantics of Natural Language*, eds. D. Davidson and G. Harman (Dordrecht, Holland: D. Reidel, 1972), pp. 169-218.
- Motherwell, Robert, editor: *The Dada Painters and Poets: An Anthology*, 2<sup>nd</sup> edition (Belknap Press, Cambridge MA, 1979).
- Peacocke, Christopher: *A Study of Concepts* (Cambridge, MA: MIT Press, 1992).
- Putnam, Hilary: "The Meaning of 'Meaning'," in *Mind, Language, and Reality* (Cambridge: Cambridge University Press, 1975), pp. 215-271.
- Quine, Willard V.O.: *Word and Object* (Cambridge, MA: MIT Press, 1960).
- Reddy, Michael: "A Semantic Approach to Metaphor," in *Papers from the Fifth Regional Meeting of the Chicago Linguistics Society*, ed. R. I. Binnick et al (Department of Linguistics, University of Chicago, 1969), pp. 240-251.
- Rosch, Eleanor: "Principles of Categorization," in *Cognition and Categorization*, ed. E. Rosch and B. B. Lloyd (Hillsdale, NJ: Erlbaum, 1978), pp. 27-48.
- Ryle, Gilbert: "Categories," in *Logic and Language* (2<sup>nd</sup> series), ed. A. Flew (Oxford: Blackwell, 1953), pp. 65-81.
- : *The Concept of Mind* (London: Hutchinson, 1949).
- Searle, John: "Metaphor," in *Metaphor and Thought*, ed. A. Ortony (Cambridge: Cambridge University Press, 1979), pp. 92-123.
- Shakespeare, William: *The Riverside Shakespeare*, 2<sup>nd</sup> edition, ed. J. J. M. Tobin et al. (New York: Houghton Mifflin Co., 1997).
- Stern, Josef: *Metaphor in Context* (Cambridge, MA: MIT Press, 2000).
- Strawson, P. F: *Individuals: An Essay in Descriptive Metaphysics* (Garden City NJ: Anchor Books, 1963).
- : "Categories," in *Ryle: A Collection of Critical Essays*, ed. O. Wood and G. Pitcher (Garden City, NY: Doubleday Anchor Books, 1970), pp. 181-211.
- Weinreich, Uriel: "Explorations in Semantic Theory," in *Current Trends in Linguistics*, vol. 3, *Theoretical Foundations*, ed. T. A. Sebeok (The Hague: Mouton, 1966), pp. 395-477.
- Wittgenstein, Ludwig: *Philosophical Investigations*, tr. G. E. M. Anscombe (Oxford: Basil Blackwell, 1972).

Yablo, Stephen: "Is Conceivability a Guide to Possibility?", *Philosophy and Phenomenological Research* 53:1 (March 1993), pp. 1-42.