Varieties of the Generality Constraint

Abstract

Since its introduction by Evans (1982), the generality constraint (GC) has been invoked by various philosophers for different purposes. Our purpose here is, first, to clarify what precisely the GC states by way of an interpretive framework, *the GC Schema*, and second, to demonstrate in terms of this framework some problems that arise if one invokes the GC (or systematicity) without clearly specifying an appropriate interpretation. By utilizing the GC Schema these sorts of problems can be avoided, and we thus propose it as a tool to facilitate argumentation that appeals to the GC.

Since its introduction by Evans (1982), the generality constraint (GC) has been invoked by various philosophers for different purposes. Evans first formulated the principle to serve as a sort of constraint on theories of human thought, and it is put to the same use by Peacoke (1992). More recently, Carruthers (2004, 2009) and Camp (2004, 2009) have invoked the GC not as a constraint on *theories* of human thought, but rather as a criterion for a *creature* to have conceptual thought. Another use of it is found in Heck (2007), who invokes the GC to argue that conceptual states (e.g. beliefs) and non-conceptual states (e.g. perceptions) have different kinds of *content*. Moreover, the general idea expressed by the GC is also expressed by what is known as "systematicity," and Fodor (and others) argue that the systematicity of (human) language and thought supports the conclusion that they both are truth-conditionally compositional. Fodor has also invoked systematicity to argue in support of "classical" architectures of the mind.[1]

---

[1] That systematicity and the GC are (at least in some forms) equivalent is noted, for example, by Johnson (2004, p. 111): "Systematicity (sometimes under the name "generality" or "recombinability") has … been discussed frequently in the philosophy of mind" and he cites Evans (1982) Peacocke (1992) and Travis (1994) as examples of such discussion.

Given the different uses to which the GC (and/or systematicity) is put, it is neither surprising nor controversial that there is considerable variation concerning what precisely the GC is; there is even some explicit debate concerning what it should be.[2] Our purpose here is not to defend or criticize a particular argument which utilizes the GC as a premise, nor to advance a particular interpretation of it. Rather our purpose is, first, to clarify the question of what precisely the GC states by proposing a framework for distinguishing between the various possible interpretations, and second, to demonstrate in terms of this framework some problems that arise if one invokes the GC (or systematicity) without clearly specifying an appropriate interpretation. We believe that by utilizing our framework to clarify the GC these sorts of problems can be avoided. We thus propose our interpretive framework as a tool to facilitate argumentation that appeals to the GC.

The paper proceeds as follows: In section I we examine Evans' original presentation of the GC in order to illustrate the interpretative options that are available. On the basis of this examination, we propose the *GC Schema* (GCS) which provides a sort of three-dimensional space of possible interpretations. In section II we apply the framework provided by the GCS to demonstrate some difficulties that arise for appeals to the GC which do not appropriately specify interpretations along the three dimensions.

I. Critical Discussion of Evans' Statement of the Generality Constraint, and Introduction of the Generality Constraint Schema.

The GC was first formulated by Evans (1982) as a constraint on theories of thought. As Evans puts it:

---

[2] There is considerable debate concerning what precisely the GC should say, *given* that it is a necessary condition for the capacity of conceptual thought. See Camp (2004, 2009) and Carruthers (2009).

> In discussing the nature of our conceivings we have little enough to go on, but there is one fundamental constraint that must be observed in our reflections: I shall call it 'The Generality Constraint'. (1982, p. 100).

Evans then formulates the GC as follows:

> … 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. (1982, p. 104)

Now of course what Evans' statement of the GC says depends upon how it is interpreted. In particular, what it says depends upon the interpretation of three notions: First, how is one to understand the notion of *ability* that is invoked? Or, to refer to Evans' formulation, what is it to "have the conceptual resources" for doing something? If we construe the ability as a disposition, and we agree that dispositions can be characterized with subjunctive conditionals, we can phrase the question this way: to say of a subject that he is *able to* entertain thought X is to say that *if* certain conditions *were* met, he *would* entertain X. The question is, what are those "certain conditions" and how far from reality is one to go in evaluating the truth of the conditional? (For example, would we say of a brilliant philosopher who is having a bad day that she is not *able* to entertain thought X because, *if* she were presented with stimuli that on a normal day would lead her to entertain X, then she would not, on this bad day, entertain X?) We will refer to the requirement of providing an interpretation of the notion of ability invoked in the GC as the requirement to *specify an interpretation along the ability dimension.*[3]

Second, how is one to interpret the notion of *entertaining* a thought? One is not likely to get far clarifying the notion of *entertaining* a thought without at the same time providing some conception of *thoughts*. (Concurrently with his presentation of the GC, Evans proposes that thoughts be understood as complexes of abilities, 1982, p. 101.)

---

[3] Camp (2009) ought to be understood as arguing in support of a particular interpretation along the ability dimension. Her proposal is that to be *able* think a thought X a creature must be able to think X *independent of any particular stimuli.*

Evans' presentation of the GC in terms of "*a* is *F*" and "*a* is *G*" might suggest that entertaining a thought is equivalent to understanding a sentence of natural language. This is, to be sure, one possible interpretation of *entertaining a thought*, and it should be noted that Evans introduces the general idea of the GC in terms of understanding sentences (1982, p. 101). But it seems that many who invoke the GC, Evans included, have not conceived of *entertaining a thought* as being dependent on language in this way. But then what is it to entertain a thought? Surely, *entertaining* X does not require *believing* that X is true, (nor *desiring* that X be true, etc.) but then what does it require? Assuming that thoughts are the sorts of things that can be true, does entertaining a thought X require a grasp of *truth conditions* of X?[4] Or the ability to verify, and/or falsify, X? Or perhaps one merely has to recognize that X is *apt for truth*, without necessarily having a grasp of what it would be for X to be true? E.g. to borrow an example from Searle (1980), even if one does not know what it would be for John to cut the sun, can one nonetheless entertain the thought that *John cut the sun*? And what about thoughts expressed by highly context-dependent sentences? Right now, out of the blue, can you *entertain* the thought expressed by 'Everybody else was there too.'? We will refer to the requirement of providing an interpretation of the notion of *entertaining a thought* invoked in the GC as the requirement to *specify an interpretation along the entertainment dimension*.[5]

And finally, third, how is one to understand Evans' claim that the putative entertainer of *a is F* must be able to entertain *a is G* for "every property of being *G* of which he has a conception"? The basic idea of the GC is that a putative entertainer of

---

[4] For Evans *entertaining* a thought X requires *knowing what it would be for X to be true*. One of Evans' goals in *The Varieties of Reference* is to clarify what *knowing what it would be for X to be true* amounts to in the case where X is a singular thought.

[5] Travis (1994) ought to be understood as objecting to a relatively strong interpretation along the entertainment dimension. That is, Travis argues that if *entertaining* a thought X requires the ability to *grasp the truth conditions* of its relevant substitution instances, then normal humans do not entertain many thoughts – i.e. normal humans do not satisfy the GC.

X must be able to entertain some range of *relevant substitution instances* of X. In the citation above Evans states that the range of relevant substitution instances for a thought of the form *a is F* is "*a is G* for every property G of which [the subject] has a conception." The first thing we should note here is that this phrase does not accurately capture even Evans' view, for in other passages Evans clearly indicates that he holds that the range of relevant substitution instances includes substitutions for the subject '*a*', and not merely substitutions for the predicate term '*G*'. For example, Evans writes,

> 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 *a* is *F*, that *a* is *G*, that *a* is *H*, … (1982, p. 104, fn. 21)

The second thing that should be noted about Evans' characterization of the range of relevant substitution instances is that it involves a sort of use/mention, or rather concept/referent, confusion. Evans writes as if '*a*' and '*F*' (etc.) are used to refer to (or quantify over) objects and properties: he says a putative thinker of *a is F* must be able to entertain *a is G* "for every *property G* of which he has a conception" (our emphasis). But this cannot be correct. Rather when Evans uses '*a*', '*b*', '*F*', '*G*' he must be referring to (or quantifying over) *thought constituents* – which Evans identifies with *abilities* to represent things. (More precisely but less perspicuously, Evans must be using '*a*', '*F*' etc. as either schematic letters, proxies for terms referring to thought constituents, or as variables ranging over thought constituents.) So, *a, b, F, G,* etc. are thought constituents, or concepts, and not properties and objects represented by concepts. Now Evans maintains that concepts are abilities rather than, as the LOT hypothesis would have it, symbols (Evans 1982, p. 101), but that view of concepts, of what '*a*', '*F*', etc., refer to, lies outside what is explicitly stated by the GC. To remain as neutral possible on the nature of concepts and thoughts, the things that subjects

entertain, we will refer to what subjects entertain as *molecular representations* (e.g. *a is F*) composed of *atomic representations* (*a*, *F*, and *is*.).

So, Evans must be using '*a*' and '*F*' etc. as schematic letters serving as proxies for terms referring to *atomic representations*. And thus both the *source* (e.g. *a is F*) and the resulting *substitution instances* (e.g. *a is G*, *b is F*, etc.) are molecular representations composed of such atomic representations. Moreover, the substitution instances differ from the source in that, to put it somewhat metaphorically, there is some position common to both the source and the substitution instance such that where the source has one atomic representation in that position, the substitution instance has some other atomic representation. (For reasons that will appear later, it is useful to allow that every molecular representation is a substitution instance of itself; i.e. there is no reason to require that a substitution instance of a source involve some "other" atomic representation.) But once it is appreciated that the GC says something concerning substitutions of atomic representations in various *positions* of molecular representations, we must acknowledge that *the relevant substitution instances* – the ones the putative entertainer of *a is F* must be able to entertain – must be restricted by some sort of *grammaticality constraint*. Indeed, that there is some sort of grammaticality constraint on the range of relevant substitution instances is displayed – though not explicitly stated – by Evans' presentation of the GC. According to Evans, the putative entertainer of *a is F* must be able to entertain every molecular representation in the sequence *b is F*, *c is F*, *d is F* …., and she must be able to entertain every molecular representation in the sequence *a is G*, *a is H*, *a is I*, ... But note what the GC does *not* require her to be able to entertain: she need not be able to entertain, e.g., *a F is*, nor *a b a*, nor *is is F*, and so on. The point is that Evans simply assumes that the relevant substitution instances are well-formed, or grammatical. And this assumption in turn presupposes that there is

some *grammar* for the representational system of which *a* and *b* and *F* and *G* and *is*, etc. are atomic representations, a grammar which distinguishes well-formed molecular representations (including sentence-like molecular representations) from mal-formed, ungrammatical, molecular representations.

Though Evans neither explicitly states such a *grammaticality constraint* on the range of relevant substitution instances nor specifies a representational system with an attendant grammar relative to which such a constraint could be determined, his assumption that such a constraint is required is well motivated. Again in terms of Evans' statement of the GC, we clearly do not want to require the putative entertainer of *a is F* that she be able to entertain ungrammatical representations such as *a F F*. For instance, we would certainly want to say that most normal adult citizens of the US are able to entertain the molecular representation *Obama smokes.* But it seems it would not be correct to say that most such subjects can entertain complex representations such as, e.g., *cycles smokes*, and *Obama Bush*, or, going further, *of smokes*, *Obama of*, etc. So, the GC must be interpreted as imposing some sort of grammaticality constraint on the range of relevant substitution instances, and this presupposes that there is some representational system (RS) with an attendant grammar that determines the range of *grammatical* substitution instances. Typically philosophers who invoke the GC (or systematicity) do not specify such a grammar, but only indicate implicitly the RS that is relevant to their use of the GC. For example, Evans never specifies a grammar for the RS to which '*a is F*', '*b is F*', etc. belong, but it is implicit in what he says that the relevant RS is the conceptual system of a normal adult human (whatever that system turns out to be). The general point is that providing an interpretation of the GC (or systematicity) requires at least an implicit indication of an RS. A statement that

something does, or does not, satisfy the GC (or is, or is not, systematic) can be evaluated only relative to such an RS.

Though some sort of grammaticality constraint must be invoked as providing a necessary condition for a candidate molecular representation to be in the set of relevant substitution instances, it is not obligatory to consider such a constraint as providing a sufficient condition. In fact, Evans himself seems to have viewed grammaticality as only a necessary condition. In a footnote to the passage where Evans introduces the GC as applying to understanding sentences of natural language, Evans says that the range of relevant substitution instances will have to be restricted by a "proviso about the categorical appropriateness of the predicates to the subjects" (p. 101, fn. 17). What Evans seems to have in mind here by "categorical appropriateness" goes beyond the mere *grammaticality* constraint displayed by his uses of lower-case and upper-case letters; rather the categories he has in mind seem to be something like those involved in Ryle's (1949) notion of a *category mistake*. Thus, to use a familiar example, we could say that even though 'Caesar is prime' is *grammatical*, it, we may suppose, involves a category mistake. And thus, if an appropriate sort of category restriction on the range of relevant substitution instances were imposed in addition to the minimal grammaticality constraint, putative entertainers of the molecular representation *Two is prime* would not be required to be able to entertain *Caesar is prime*.

In the preceding paragraphs we have been concerned to illustrate that one of the things one must do in specifying an interpretation of the GC is to specify a *range of relevant substitution instances*. But once it is appreciated that such a range, or perhaps we should now say *domain*, must be specified, another question of interpretation presents itself: *How many* (and perhaps what *sorts*) of such substitution instances in the range must the putative entertainer be able to entertain? Evans' official view on this

8

point seems to be that the putative thinker of *a is F* must be able to entertain *every* relevant substitution instance (where the range of relevant instances is limited by both the grammaticality constraint and by categorical restrictions). But, as several philosophers have noted (XXX forthcoming, Carruthers 2009), under this strong interpretation it is less plausible that even normal humans satisfy the GC. And even Evans seems to have thought that requiring that the putative thinker be able to entertain *every* relevant substitution instance is too strong. For after stating the GC, as cited above, Evans writes, "perhaps it ought to be conceded that the Generality Constraint is an ideal, to which our actual system of thoughts only approximately conforms" (1982, p. 105). Though it is not clear what exactly is the cause of Evans' reservation in this passage, it is very plausible to read it as suggesting that *really* thinkers can be credited with the thought that *a is F* even though they are able to entertain only *some* of the relevant substitution instances of *a is F*. And this of course raises the question of *how many* relevant substitution instances the putative entertainer of *a is F* must (really) be able to entertain. Would the ability to entertain *one* relevant substitution instance be sufficient? Perhaps a plausible interpretation would involve *sorts* of substitution instances.[6] For example, it could be said that a putative entertainer of *a is F* must be able to entertain one admissible substitution instance for every atomic representation in *a is F*. (So, the putative entertainer would have to be able to entertain *one* of *b is F, c is F*, … etc., *and* she would have to be able to entertain *one* of *a is G, a is H*, … etc..) At any rate, the general point is that interpreting the GC requires not only specification of

---

[6] We could construe such *sorts* of substitution instances as resulting from the addition of restricting modifiers *M* that modify the noun 'substitution instance' as it appears in the consequent of the GC: "If a subject is able to entertain a molecular representation *a is F*, then she is able to entertain *Q* admissible substitution instance(s) *M* of the molecular representation *a is F*," where *Q* is some quantifier. For example, *Q* could be "seven" and *M* could be "that result from a substitution for *a*." This would yield, "If a subject is able to entertain a molecular representation *a is F*, then she is able to entertain *seven* admissible substitution instances *that result from substitution for a* in the molecular representation *a is F*."

the range of relevant substitution instances, but moreover specification of *how many* (and perhaps what *sort*) of such relevant substitution instances the putative entertainer must be able to entertain.

So, to specify the relevant substitution instances an interpreter of the GC must complete two related interpretive tasks: First, she must specify the range of *relevant substitution instances* for a given molecular representation. This requires appeal to a *grammaticality constraint*, which in turn requires specification of, or at least allusion to, a pertinent RS. Completion of this task may also involve specification of further "categorical restrictions." Second, the interpreter must specify *how many* (and perhaps *what sort*) of these relevant substitution instances the putative thinker of a given molecular representation must be able to entertain. Though these interpretive tasks are independent in the sense that completion of one of them does not necessarily imply completion of the other, they are closely related; indeed, there is a clear sense in which the tasks are complementary: If, as in Evans' official statement of the GC, one requires that a putative entertainer be able to entertain *all* relevant substitution instances, then one will probably be compelled to restrict the domain of quantification in some way. Whereas, on the opposite extreme, if one requires a putative entertainer to be able to entertain only *one* relevant substitution instance, then the motivation for restricting the domain, i.e. the range of relevant substitution instances, is undermined.[7] Because of this complimentary relation between the two tasks, we will treat them both as being aspects of the same interpretive dimension. Hence, we will refer to the requirement of specifying the range of relevant substitution instances, together with the related task of

---

[7] An interpretation of the GC that requires a putative entertainer to be able to entertain all relevant substitution instances has the general form, *If P, then all Fs are Gs*. The *fewer Fs* there are, the more plausible *All Fs are Gs* will be, and thus the more plausible the entire conditional will be. In contrast, an interpretation of the GC that requires a putative entertainer to be able to entertain only *some* relevant substitution instances has the general form *If P, then some F is G*. The *more Fs* there are, the more plausible *some F is Gs* will be, and thus the more plausible will be the entire conditional. (Carruthers, 2009, endorses a version of the latter, existential, interpretation)

specifying how many (and perhaps what sort) of such relevant substitution instances a putative entertainer must be able to entertain, as the requirement to *specify an interpretation along the substitution dimension.*[8]

The purpose of the preceding examination of Evans' original presentation of the GC is not to criticize Evans for not stating the GC with sufficient clarity. Indeed, though we maintain that Evans does not really specify a theoretical principle of cognitive psychology so much as provide a sort of schema for a range of such principles, it is far from clear that such a schema is insufficient for Evans' rather general purposes. But, as was noted in the introduction, other philosophers have used the GC for different purposes, and, as will be made clear below, for these other purposes more specification is required. So, we propose to extract from Evans' presentation of the GC a *GC Schema* which serves as a sort of template for more specific interpretations of the GC. There are various ways that such a schema could be formulated (using free-variables, or schematic letters), but we have chosen to present it with a sentence that triggers many presuppositions, all of which must be resolved in order for the sentence to be interpreted. Where *RS* is the pertinent representational system, the *GC Schema* (GCS) is as follows:

> GCS: if a subject has *the requisite sort of ability* to *entertain in the requisite way* a molecular representation X (of *RS*) then she has *the requisite sort of ability* to *entertain in the requisite way the requisite number* (perhaps of *the requisite sort*) of *the relevant grammatical (in RS) substitution instances* of X.

Obviously, in order to specify an interpretation of the GCS, in addition to indicating the pertinent RS, one must specify all of the following: (1) what *the requisite sort of ability* is; (2) what *the requisite way of entertaining* is; (3a) what *the relevant substitution instances of a given molecular representation* are; and finally (3b) what *the requisite number of relevant substitution instances* is (and perhaps also what *the requisite sort of*

---

[8] Carruthers (2009) ought to be understood as objecting to a relatively strong interpretation along the substitution dimension, and proposing a weaker one in its place.

11

relevant substitution instance is). To satisfy the requirement stated in (1) is to specify an interpretation along the *ability dimension,* and to satisfy the requirement in (2) is to specify an interpretation along the *entertainment dimension.* And satisfying the requirements stated in (3a) and (3b) is to specify different aspects of an interpretation along the *substitution dimension.*

Before explaining some problems that arise for philosophers in using the GC without specifying an appropriate interpretation along the three dimensions given by the GCS, several general remarks are in order. Note that the three dimensions are mutually dependent in the following sense: One can arrive at a plausibly true interpretation of the GC *and* specify relatively strong (i.e. difficult to satisfy) interpretations along *some* of the dimensions as long as one specifies relatively weak (i.e. easy to satisfy) interpretations along the *other* dimension(s). For example, a relatively strong interpretation along the substitution dimension would be one that required a putative thinker to be able entertain *every* grammatical (according to some pertinent grammar) substitution instance of, e.g., *two is prime*; this strong interpretation would require a putative thinker of *two is prime* to be able to entertain categorically mismatched thoughts such as *Caesar is prime,* and *Two is dead.* But is a putative thinker of *two is prime* able to entertain such thoughts? It depends, of course, on the interpretations that are specified along the *ability* and *entertainment* dimensions. If one complements the above strong interpretation along the substitution dimension, with weak interpretations along the ability and entertainment dimensions, then it is more plausible that normal humans satisfy the GC, so interpreted.[9]

_____

[9] For example, suppose *able* is interpreted in such a way that the putative thinker would be *able* to entertain *Caesar is prime* even if it were nomologically impossible for her to do so, and further suppose that *entertain* is interpreted in such a way that the putative thinker would be able to *entertain Caesar is prime* even if she had no grasp whatsoever of conditions under which this molecular representation would be *true.* If the strong interpretation along the substitution dimension is combined with these weaker

A related point concerns what happens if strong interpretations are specified along all three dimensions, or if weak interpretations are specified along all three dimensions. On the one hand, if all dimensions are interpreted strongly, then the resulting overall interpretation of the GC would be very implausible – perhaps even incoherent. But, on the other hand, if all dimensions are interpreted weakly, then the resulting overall interpretation of the GC would be very plausible, but insignificant – perhaps even trivial. And finally, as one might expect, there is an inverse correlation between *plausibility* and *significance*: there are overall interpretations of the GC that render it very plausible, but such readings are less likely to have the sorts of significant consequences that philosophers have supposed the GC to have. And, conversely, there are overall interpretations of the GC that render it very significant, yet on these readings the GC is much less plausible than philosophers have supposed it to be.

II. Problems that Arise as a Result of Failure to Specify an Interpretation of the GC

In the preceding section we have been concerned to demonstrate that interpreting the GC requires specifying interpretations along three dimensions of the GCS. In this section we will present four arguments that utilize the GC (or "systematicity") as a premise and we will demonstrate how failure to appropriately specify interpretations along the three dimensions undermines the soundness of each argument.

---

interpretations along the ability and entertain dimensions, then the resulting overall interpretation of the GC is plausibly true (applied to normal adult humans).

Our first example is Fodor's argument in support of the *compositionality* of thought and (at least at one time[10]) language. This argument appeals to "systematicity," which is clearly equivalent to some version of the GC. For instance, Fodor (1987) introduces the notion of systematicity with regard to linguistic competence as follows:

> The property of linguistic capacities that I have in mind is one that inheres in the ability to understand and produce sentences. That ability is – as I shall say – *systematic*: by which I mean that the ability to produce/understand some of the sentences is *intrinsically* connected to the ability to produce/understand many of the others (1987, p. 149).

This citation makes it apparent that *systematicity* is equivalent to some unspecified version of the GC, though Fodor and Evans are clearly using the principle for different purposes: whereas Evans invokes the GC as a constraint that putative theories of thought must satisfy – thus in effect imposing the *apriori* restriction on thought that it satisfy the GC – Fodor invokes systematicity as an indisputable empirical fact about thought that is in need of explanation.

The following passage presents a paradigmatic instance of Fodor's systematicity argument in support of the *compositionality* of thought:

> 'Systematicity' is a cover term for a cluster of properties that quite a variety of cognitive capacities exhibit … Here are some typical examples. If a mind can grasp the thought that $P \rightarrow Q$, it can grasp the thought $Q \rightarrow P$; if a mind can grasp the thought ~$(P\&Q)$, it can grasp the thought that ~P and the thought that ~Q; if a mind can grasp the thought that Mary loves John, it can grasp the thought that John loves Mary, etc. Whereas it is by no means obvious that a mind that can grasp the thought $P \rightarrow Q$ can also grasp the thought that $R \rightarrow Q$ … . That will depend on whether it is the kind of mind that's able to grasp the thought that R. Correspondingly, a mind that can think *Mary loves John* and *John loves Mary* may none the less be unable to think *Peter loves Mary*. That will depend on whether it is able to think about Peter.

---

[10] Fodor has appealed to *systematicity* to support the *compositionality* of thought and language many times (e.g. 1987, 1988, 1998, 2002). But, more recently (2001, 2008) he has defended the view that only thought is compositional. Thus Fodor's current position with regard to the systematicity arguments in support of compositionality must be that they succeed for thought, but not for language. To our knowledge Fodor has yet to explain why the systematicity argument goes through for thought, but not for language. (For an argument against Fodor's "mixed view" that thought, but not language, is compositional, see XXX under review.)

It seems pretty clear why the facts about systematicity fall out the way they do: mental representations are compositional, and compositionality explains systematicity. (1998, p. 97)

The argument thus has the following general form:

1. Thought (of normal adult humans) is *systematic*.

2. If thought is *systematic*, then thought is (probably[11]) *compositional*. So,

3. Thought is (probably) compositional.

Fodor is not very clear as to what systematicity is. In particular, he says nothing with regard to what it is to "grasp" a thought; in terms of the GCS, he does not specify an interpretation along the *entertainment* dimension. Now, the failure to specify an interpretation of the GC (or systematicity) may or may not be problematic, depending upon the use to which the GC (or systematicity) is being put. But in the case of the above argument, the lack of specificity is problematic, for, as we will now explain, the argument seems to trade on an equivocation on 'systematicity' as it is invoked in the two premises.

Let us consider premise 2 first. In order for this premise to be plausible, 'sytematicity' as it appears here must be interpreted relatively strongly. Or, to put it in terms of the GCS, premise 2 is plausible only if a relatively strong (discriminating) interpretation is specified along the entertainment dimension; "grasping a thought" must be interpreted as requiring *representation of truth conditions*. This is because the sort of compositionality Fodor has in mind concerns *propositions*, entities which possess *truth conditions* essentially. This is illustrated in the following passage, wherein Fodor clarifies what he has in mind by *compositionality*:

3. Compositionality: concepts are the constituents of thoughts and, in indefinitely many cases, of one another. Mental representations inherit their contents from the contents of their constituents.

---

[11] Fodor maintains that the support for this premise involves an abductive inference: "inferences from systematicity to compositionality are 'arguments to the best explanation'" (1987, p. 98, fn. 11)

> *Some terminology*: I'll use 'thoughts' as my cover term for the mental representations … which express the *propositions* that are the objects of propositional attitudes. Thus, a belief that it will rain and a hope that it will rain share a thought as well as a *proposition* which that thought expresses. For present purposes, it will do to think of thoughts as mental representations analogous to closed sentences (1987, p. 25, our emphasis on both occurrences of 'proposition').

So, the sort of *compositionality of thoughts* Fodor has in mind is the familiar propositional, truth-conditional, sort of compositionality: the *proposition* (or truth *conditions*) that is the content of a thought is a function of the contents of the atomic representations within the thought, and its structure.   If compositionality, as it is invoked in premise 2, is this sort of truth-conditional compositionality, what must systematicity be, if premise 2 is to be true? That is, what must the *systematicity of thought* be, if it is to warrant the inference to the *truth-conditional compositionality of thought*? Clearly the requisite sort of systematicity must invoke representation of *truth-conditions*. Or, in terms of the GCS, in order for premise 2 to be plausibly true, we must specify a relatively strong interpretation along the entertainment dimension, an interpretation which requires of a putative entertainer of, e.g. *John loves Mary*, that she represent what the world would have to be like in order for this thought to be true. For if "grasping" a thought had nothing to do with representing *truth conditions*, then the nomological connection between the ability to grasp *John loves Mary* and the ability to grasp *Mary loves John* would not imply anything about how the *truth conditions* of such thoughts are determined.

Thus the plausibility of premise 2 depends upon specification of a relatively strong interpretation along the entertainment dimension; premise 2 is plausible only if "grasping a thought" requires representation of how the world would have to be in order for the thought to be true. But, it is far from clear that premise 1 is true under this relatively strong interpretation; i.e. it is not at all obvious that we (normal adult humans)

would satisfy the GC (or systematicity) if such a strong reading were specified along the entertainment dimension. For example, to modify a classic example from Searle (1980), it is plausible that the thought corresponding to '*Obama passed the health care bill*' represents truth conditions, for we have a relatively good idea as to what it would for a use of this sentence, and thus the corresponding thought, to be true. But the same cannot be said for the grammatical substitution instance '*The health care bill passed Obama*'. Of course we the have the ability to recognize that this sentence is *grammatical*, and thus we recognize that, under appropriate circumstances, one might use the sentence to assert something we would judge to be true (or false). But, taken out of the blue, we associate no truth-conditions with the sentence itself, and thus there is no reason to suppose that the corresponding thought (assuming there is one) represents truth-conditions either.[12] Indeed, it is telling that in the passage above wherein Fodor presents examples which are supposed to illustrate the *systematicity* of thoughts, he does not provide examples of sentence – which presume are to correspond with thoughts – that we (his readers) know the *truth conditions* of. For instance, we have no idea what it would be for the thoughts corresponding to '*John loves Mary*' and '*Mary loves John*' to be true, since we have no idea who the relevant John and Mary are (nor *when* the alleged loving is said to occur).

So, first, if 'systematicity' in premise 1 is to be interpreted so as to render the premise obviously – or even plausibly – true, then *systematicity* must be interpreted in a way that does not require in order to "grasp a thought" that a subject represent truth-conditions. And second, to the extent that the examples Fodor provides in the above passage lend support to premise 1, 'systematicity' as it appears in this premise must be

---

[12] Perhaps Fodor could invoke some sort of *categorical restriction* to block such apparent counterexamples to systematicity.   Though we will not pursue the issue here, this suggestion faces an immediate problem: the notion of a *category violation* will have to be defined independently of the alleged systematicity of thought, and it is not clear how this could be done without invoking an analytic-synthetic distinction, and of course this option is not available to Fodor.

interpreted as not requiring for "grasping a thought" such representation of truth-conditions. Now, it is of course possible to interpret 'systematicity' as it appears in premise 1 in a weaker, less discriminating way; it is of course possible to interpret "grasping a thought" so that it does not require representation of truth-conditions. For instance, if "grasping a thought" required only a sort of *syntactic* competence, then premise 1 would be both plausibly true *and* supported by the examples Fodor provides. But then the argument equivocates, for, as we explained above, premise 2 requires the stronger, truth-conditional, interpretation of "grasping a thought."

Our second example is a related argument advanced by Fodor and Pylyshyn. They appeal to *systematicity* to argue against connectionist, non-classical, architectures of the mind (Fodor and Pylyshyn, 1988). The essence of their argument is as follows[13]:

1. Our linguistic capacity is "a paradigm of systematic cognition" (1988, p. 37).

2. The best explanation for the systematicity of our linguistic capacity involves "the postulation of constituent structure" (1988, p. 38) in the sentences of natural language.

3. Our capacity for thought is systematic in the same way that our linguistic capacity is systematic.[14]

4. Just as the systematicity of language shows that there must be structural relations between the constituents of sentences of natural language, so the systematicity of thought shows that there must be structural relations between the constituents of thoughts (mental representations). And therefore,

---

[13] Fodor and Pylyshyn state that "the argument from the systematicity of linguistic capacities to constituent structure in sentences is quite clear. *But thought is systematic too*, so there is a precisely parallel argument from the sytematicity of thought to syntactic and semantic structure in mental representations" (1988, p. 39).

[14] At one point Fodor and Pylyshyn claim, as a sort of aside, that "in the case of verbal organisms the systematicity of thought *follows from* the systematicity of language" (1988, p. 39). They give the following argument: "… understanding a sentence involves entertaining the thought it expresses; on that assumption, nobody could understand both the sentence about John and the girl unless he were able to think both the thoughts about John and the girl" (1988, p. 39). Two points are worth making here: (i) the argument does provide a representationalist – one who thinks thought/understanding involves mental representations – with a good reason to think that *some* thought involves representations with constituent structure. But (ii) if it is allowed that some thoughts are not expressed by sentences of natural language, then the argument fails to provide a reason for thinking that *all* thought is *systematic*.

5. "mental representations have an internal structure" and thus "the architecture of the mind is not a connectionist network" (1988, pp. 39-40).

Our point concerns the third premise. We will demonstrate, by utilizing the interpretive framework provided by the GCS, that the sort of *systematicity* that Fodor and Pylyshyn claim is possessed by our linguistic capacities cannot be applied to the case of thought; or more precisely, we will demonstrate that application of this interpretation of *systematicity* to the case of *thought* renders the third premise trivially true. Our first task then is to make clear the sort of systematicity that Fodor and Pylyshyn claim is possessed by our linguistic capacities. After this sort of systematicity is made clear, i.e. after this interpretation of the GC is specified, we will explain why the claim that thought is *systematic* under this interpretation is trivially true.

Fodor and Pylyshyn define *systematicity* as applied to linguistic capacities in the following passage:

> What we mean when we say that linguistic capacities are *systematic* is that the ability to produce/understand some sentences is *intrinsically* connected to the ability to produce/understand certain others (1988, p. 37).

But how exactly should this rather vague definition is to be interpreted? Let us apply the framework provided by GCS, and determine interpretations along the three dimensions. Let us begin by considering the *entertainment* dimension. It is clear that the pertinent *representational system* is a natural language, and thus that the sort of *molecular representations* that are relevant here are sentences of a natural language. (Fodor's examples of course involve sentences of contemporary English.)   What then is the relevant interpretation of *entertain*? That is, what according to Fodor and Pylyshyn is it for a subject to "produce/understand" a sentence?   It is clear that Fodor and Pylyshyn do *not* require for such "production/understanding" of sentences any sort of

19

knowledge of *content*, or *truth conditions*. Rather, the ability to "produce/understand a sentence" requires only a sort of *syntactic* competence[15]:

> Suppose for example that it's a fact about English that formulas with the constituent analysis 'NP Vt NP' are well formed; and suppose that 'John' and 'the girl' are NPs and 'loves' is a Vt. It follows from these assumptions that 'John loves the girl,' 'John loves John,' 'the girl loves the girl,' and 'the girl loves John' must all be sentences. It follows too that anybody who has mastered the grammar of English must have linguistic capacities that are systematic in respect of these sentences; he *can't but* assume that all of them are sentences if he assumes that any of them are. (1988, p. 38)

This passage makes it clear that Fodor and Pylyshyn are specifying an interpretation along the entertainment dimension which requires a putative entertainer of a molecular representation, i.e. a putative producer/understander of a sentence, only that she be able to recognize that the sentence *is a sentence*. That is, the ability to produce/understand a sentence just is the ability to determine the grammaticality of a sentence. So, for Fodor and Pylyshyn to *entertain a molecular representation* is simply to recognize that a grammatical sentence *is* a grammatical sentence.

Let us now consider what interpretation Fodor and Pylyshyn specify along the *ability* dimension. That is, how are we to understand the notion of *ability* invoked in Fodor and Pylyshyn's above definition of systematicity? Though Fodor and Pylyshyn do not directly specify an interpretation along the ability dimension, what they say with regard to the *productivity* of our linguistic competence suggests that they intend *ability* to be understood as the sort of linguistic *competence* that Chomsky (1968) famously distinguished from linguistic *performance*:

> …Chomsky … has claimed (convincingly, in our view) that the knowledge underlying linguistic competence is generative – i.e., that it allows us *in principle* to generate … an unbounded number of sentences. It goes without saying that no one does, or could, in fact utter … tokens of more than a finite number of sentence types … But there are a number of considerations which

---

[15] This is also implied by Fodor and Pylyshyn's claim that "systematicity is a property of the mastery of the syntax of a language" (1988, p. 37).

suggest that, despite de facto constraints on performance, one's knowledge of one's language supports an unbounded productive capacity … (1988, p. 34)

So, given the acknowledged similarity between Chomsky's *productivity* argument and Fodor and Pylyshyn's *systematicity* argument, we should understand the notion of ability invoked in Fodor and Pylyshyn's definition of systematicity as deriving from Chomsky's notion of linguistic *competence* (as opposed to *performance*). Thus, to say that a subject is "*able* to produce/understand" a sentence is to say that she possesses (perhaps unconsciously) knowledge that *in principle* enables her to generate the sentence; this knowledge may not *in fact* enable her to generate the sentence, as such things as "the transient state of the speaker's memory and attention" (1988, p. 34) may prevent this competence from manifesting itself in performance.

Finally, let us consider what interpretation Fodor and Pylyshyn specify along the *substitution* dimension. That is, how are we to understand the phrase 'certain others' as it is used in Fodor and Pylyshyn's definition of systematicity as "the ability to produce/understand some sentences is *intrinsically* connected to the ability to produce/understand *certain others*" (1988, p. 37, our emphasis of 'certain others')? Though again, Fodor and Pylyshyn do not explicitly specify an interpretation along the substitution dimension, a previously cited passage clearly indicates what they have in mind. In this passage Fodor and Pylyshyn claim that "anybody who has mastered the grammar of English must have linguistic capacities that are systematic in respect of ['John loves the girl,' 'John loves John,' 'the girl loves the girl,' and 'the girl loves John']" (1988, p. 38). And then they state what such systematicity with respect to these sentences amounts to: "[the subject] *can't but* assume that all of them are sentences if he assumes than any of them are" (1988, p. 38). So, for example, the subject's linguistic capacities are systematic with respect to 'John loves the girl' *iff* if the subject has the competence to recognize that 'John loves the girl' is a grammatical sentence,

21

then she has the competence to recognize that each of 'John loves John,' 'the girl loves the girl,' and 'the girl loves John' is a grammatical sentence. Let us note some features of this illustrative example.

First, the relevant range of relevant *substitution instances* of 'John loves the girl' is constrained to *permutations* of 'John loves the girl'. That is, the relevant substitution instances of a sentence *S* involve only words occurring in *S*; so, for example, the systematicity of the subject's linguistic capacities with regard to 'John loves the girl' does not require that she have the competence to recognize that, e.g. 'Sam loves the girl' is a grammatical sentence.[16] Thus, in the case of Fodor and Pylyshyn, it is perhaps more suggestive to speak of a class of relevant *permutations* instead of a class of *relevant substitution instances*.[17]

Second, note that Fodor and Pylyshyn clearly specify what *the requisite number* of relevant permutations is. Given that the set of relevant permutations for a given sentence is relatively small, it is no surprise that Fodor and Pylyshyn require that the putative entertainer of *S* be able to entertain *every* relevant permutation of *S*; or as Fodor and Pylyshyn put it, the putative entertainer must "assume that *all of them* are sentences if he assumes than any of them are" (1988, p. 38, our emphasis).

And finally, third, note that the range of relevant permutations of 'John loves the girl' is constrained by the grammar of English. That is, the range of relevant

---

[16] Thus, Fodor and Pylyshyn's way of characterizing the relevant range of substitution instances is different than Evans' way. If we adopt Evans' version of the GC to the case at hand, we get something like this: a putative entertainer of 'John loves the girl' must be able to entertain every sentence of the form "NP loves the girl" *for every NP of which she has a conception.* In contrast, Fodor and Pylyshyn say that a putative entertainer of 'John loves the girl' requires only the ability to entertain every sentence of the form "NP_1 loves NP_2" where NP_1 and NP_2 are NPs that occur in 'John loves the girl'.

[17] This change in terminology, however, is not forced upon us: We could define *substitution instance* in such a way that the set of *substitution instances* of *aRb* is identical to the set of *permutations* of *aRb*: First, restrict the range of substitutable words to those occurring within *aRb*. (So, *cRb* is not a *substitution instance* of *aRb*, but *bRb* is.) Second, make the *substitution instance* relation transitive. (So, if *bRb* is a substitution instance of *aRb*, and *bRa* is a substitution instance of *bRb*, then *bRa* is a substitution instance of *aRb*.)

permutations includes 'John loves John', 'the girl loves the girl' and 'the girl loves John', but it does not include, e.g., 'John John the girl' or 'loves loves the John'. So, Fodor and Pylyshyn clearly recognize a *grammaticality constraint* on the range of relevant permutations. The interpretation along the substitution dimension specified by Fodor and Pylyshyn then comes down to this: the range of relevant substitution instances for a sentence S are the *grammatical permutations of S*, and to be credited with entertaining S, the subject must be able to entertain *every* such grammatical permutation.[18]

Now it remains only to combine the interpretations that Fodor and Pylyshyn specify along the three dimensions into a complete interpretation of the GC (or "systematicity"). But before doing this let us appreciate that Fodor and Pylyshyn's notion of "the systematicity of our linguistic capacity" must be taken to be language relative. Obviously, one's linguistic capacities might be systematic with regard to one's native language, and not systematic with regard to one in which one is only partially proficient.[19] In terms of the GCS, the pertinent RSs for Fodor and Pylyshyn's interpretation are *natural languages*. So, finally, we can state with precision Fodor and Pylyshyn's interpretation of what it would be for one's *linguistic capacities* to be systematic:

---

[18] Note that under this interpretation *many* sentences of English will not have *any* relevant permutations. For example, the set of relevant permutations of the English sentence 'Sally smokes' is the empty set, since neither 'Sally Sally' nor 'smokes Sally' nor 'smokes smokes' is a grammatical sentence. Similar remarks apply to, e.g., 'Every dog barks' and 'Nobody came' and many many other sentences. This is probably an unwelcome consequence of Fodor and Pylyshyn's interpretation along the substitution dimension, but it is not directly relevant to our main argument, and so we will not pursue the issue further. (If we allow that every sentence is a permutation of itself, then many sentences have only themselves as permutations. This would not resolve the problem; the result is still unwelcome.)

[19] This is evident in Fodor and Pylyshyn's comparison between our actual linguistic capacities and what they call "the phrase book model" (1988, p. 37): "on the phrase book model, it would be perfectly possible to learn that uttering the form of words 'Granny's cat is on Uncle Arthur's mat' is the way to say (in English) that Granny's cat in on Uncle Arthur's mat, and yet have no idea at all how to say … that Uncle Arthur's cat is on Granny's mat …" (1988, p. 37). The clear implication is of course that one might have mere phrase book linguistic capacity *with regard to English*, but have systematic linguistic capacity *with regard to some other language*.

A subject's linguistic capacity (with regard to natural language *L*) is *systematic iff*, for every sentence *S* (of *L*), if she has the competence to recognize that *S* is a grammatical sentence (of *L*), then, for every *grammatical* (in *L*) permutation *P* of *S*, she has the competence to recognize that *P* is a grammatical sentence (of *L*).

Let us now turn to the case of thought. If the above specifies Fodor and Pylyshyn's interpretation of systematicity, what sense can be made of Fodor and Pylyshyn's claim that "*thought is systematic too*" (1988, p. 39)? A straightforward application of the above definition of systematicity to the case of thought would yield the following:

A subject's *thought* capacity (with regard to language *L*) is systematic *iff*, for every sentence *S* (of *L*), if she has the competence to recognize that *S* is a grammatical sentence (of *L*), then, for every *grammatical* (in *L*) permutation *P* of *S*, she has the competence to recognize that *P* is a grammatical sentence (of *L*).

But what is one to make of the relativization to *language L* in the above application to the case of thought? That is, in the case of thought, what is the pertinent RS? The only plausible reply is that the pertinent RS is the subject's LOT.[20] So we must understand the relativization to *language L* as relativization to the subject's *language of thought*. This yields the following:

A subject's thought capacity is *systematic iff*, for every sentence *S* (of the subject's *LOT*), if she has the competence to recognize that *S* is a grammatical sentence (of her *LOT*), then, for every *grammatical* (as determined by her *LOT*)

---

[20] Why the *subject's LOT*? Why not some *foreign* LOT, a LOT in which the subject does not think? We have two points against this suggestion. First, the suggestion is obviously out-of-step with Fodor and Pylyshyn's argument in support of the conclusion that "mental representations have internal structure" (1988, p. 40). The idea is supposed to be that *our thought* is systematic, and this implies something about the mental representations *we utilize in thought*. If the implication is to be at all plausible, the *systematicity* of our thought must in some way concern the representations we utilize in thought, i.e. systematicity must concern *our LOT*, not some *foreign LOT*. Second, the suggestion does not seem to avoid the charge of triviality, developed below: If all the relevant sentences *S* are sentences of a *foreign LOT*, then there is *no S* such that the subject possesses the requisite competence to recognize that *S is* a sentence of the foreign LOT. (If a subject had such competence with regard to a sentence *S*, then *S* would be a sentence in the subject's LOT.) In which case it is trivially true that for *all* sentences *S* of the foreign LOT that *if* the subject has the competence to recognize that *S* is a grammatical sentence (of the *foreign LOT*), then she has the competence to recognize that *P* is a grammatical sentence (of the foreign LOT), for every relevant permutation *P* of *S*.

permutation *P* of *S*, she has the competence to recognize that *P* is a grammatical sentence (of her *LOT*).[21]

But the above interpretation of the claim that thought is systematic is trivially true. Assuming that your thought does take place in a LOT, *of course* you possess the requisite sort of syntactic competence with respect to *your LOT,* for such competence is requisite for its being *your LOT* – if your thought capacity were not systematic with respect to your LOT, *it would not be your LOT. Systematicity of thought* amounts to *syntactic competence with regard to one's own LOT.* But it is trivially true that, if one has a LOT, one is syntactically competent with regard one's LOT. [22]

The triviality can be brought into clearer focus by again considering systematicity with regard to natural languages. What prevents the collapse to triviality in these cases is that we can specify a language L (and thus a *grammar* that will determine the *grammatical sentences of L*) independently of a particular subject's linguistic competence. So, for example, the claim that a subject's linguistic capacity *with regard to Spanish* is systematic is non-trivial because (or at least to the extent that) the grammar of Spanish, and thus the sentences of Spanish and their grammatical permutations, can be determined independently of her syntactic competence. If we allow the subject's syntactic competence – her unconscious knowledge of her own idiolect – to determine what the relevant sentences and permutations are, then it follows trivially that she possesses syntactic competence with respect to *those* sentences and permutations. The problem with applying to *thought* the interpretation of *systematicity* that Fodor and Pylyshyn apply to *linguistic capacity* is then this: In the case of thought

---

[21] Note that, on pain of begging the question, it cannot be assumed that the subject's LOT has a combinatorial grammar; i.e. it cannot be assumed that the sentences of the subject's LOT have internal structure. For the claim that "mental representations have an internal structure" is the conclusion, which is supposed to follow from the premise that the subject's thought capacity is *systematic*.

[22] And note that if thought does not take place in a LOT, then the claim that one's thought is *systematic* is trivially true as a result of vacuous quantification: if one has no *LOT*, then there are no sentences of one's LOT, in which case it is trivially true that for every sentence *S* (of one's *LOT*), … .

there is no way to specify the pertinent LOT – and thus the relevant sentences and permutations – independently of the subject's own syntactic competence. And as a result the claim that her thought capacity is *systematic* with respect to this *LOT* is trivially true; if her thought capacity were not systematic with respect to this LOT, it would not be her *LOT*.[23]

We thus conclude that the third premise in Fodor's and Pylyshyn's argument is trivially true; under Fodor and Pylyshyn's interpretation of *systematicity*, thought cannot but be systematic. This is, to be sure, a rather odd objection: How could it be problematic that one of the crucial premises is not only true, but moreover *trivially* true? Recall, however, that our purpose here is not so much to criticize Fodor and Pylyshyn's argument, but to illustrate the sorts of problems that arise when the GC, or systematicity, is invoked without adequately specifying an interpretation. Moreover, the triviality of the third premise clearly casts doubt on the validity of the inference from the third premise to the fourth premise, since the fourth premise does *not* seem to be a triviality. Indeed, once the triviality of the third premise is appreciated, the fallaciousness of this inference becomes apparent.

---

[23] Johnson (2004) makes a similar point in the following passage:

> Assume … we have a connectionist device *C* which is being trained to learn English. Now suppose that *C* returns the following judgments:
>
> (28)   a.   John shot Sue.
>   b.   Sue shot John.
>   c.   John loves Sue.
>   d.   *Sue loves John.
>   e.   Shot loves John.
>   f.   *John Sue loves.
>
> In such a case, *C* has not learned English. It does not follow, though, that *C* has not learned a systematic language. (Johnson, 2004, p. 136)

Johnson's point here, approximately anyway, is that it is the capacity – or *competence* – of *C* that determines the language that *C* has learned. So if according to *C* 'Sue loves John' is ungrammatical, then it is ungrammatical in the language *C* has learned. And thus it is no threat to the systematicity of *C*'s linguistic capacity with regard to the language it has learned that it recognizes 'John loves Sue' as a grammatical sentence, but not 'Sue loves John'.

Suppose a certain creature thinks in a "phrase-bookish" LOT. The "grammar"

simply consists of a list of "sentences" which contain no constituents, and/or no internal

structure. The creature's thought capacity is none-the-less *systematic* in Fodor and

Pylyshyn's sense: Since none of the sentences S in her LOT have internal structure, for

every such S the set of relevant grammatical permutations is the empty set. So, trivially,

for every such sentence S of her LOT, if she has the competence to recognize that S is a

grammatical sentence, then, for every *grammatical* permutation P of S, she has the

competence to recognize that P is a grammatical sentence – for there are *no* such

permutations P for *any S*. (We could allow that every sentence is a grammatical

permutation of itself. But this would not affect the point; triviality still ensues, though

for slightly different reasons.) The claim that a creature's thought is *systematic* is

trivially true, yet it is at least conceptually possible that a creature have a phrase-

bookish LOT. So the claim made in premise four, i.e. "the systematicity of thought

shows that there must be structural relations between the constituents of thoughts," is

false.[24]

The triviality of Fodor and Pylyshyn's interpretation of systematicity in the case

of thought arises from a combination of factors: First, Fodor and Pylyshyn specify a

very weak interpretation along the *entertainment* and *ability* dimensions: *entertaining* a

molecular representation S is simply recognizing that S is a grammatical sentence of the

---

[24] The inadequacy of the inference from the systematicity of thought to the internal structure of *mental representations* raises doubts about the parallel inference from the systematicity of our linguistic capacities to the internal structure of *sentences*. Obviously, a systematic linguistic capacity with regard to a phrase-bookish language would not imply that the sentences of the phrase-bookish language had internal structure. (Since, by definition, sentences of a phrase-bookish language have no such internal structure.) But, no less obviously, systematic linguistic capacity with regard to a generative, combinatorial, natural language, a language whose sentences do have internal structure, would imply that the sentences of the language have internal structure. So, depending upon what sort of language L is, the inference from the systematicity of a subject's linguistic capacity with regard to L is either invalid (if L is phrase-bookish), or question-begging (if L is combinatorial). This resonates with Johnson's conclusion regarding claims concerning the systematicity of our linguistic capacities: "Recursive syntax (and semantics, and so forth) and a highly structured lexicon are well-known and well-studied components of human language, whereas systematicity is at best a by-product of the former, and at worst little more than a will-o'-the-wisp" (2004, p. 130).

pertinent RS, and being able to recognize that *S* is a grammatical sentence is merely

having the requisite sort of grammatical *competence* to do so. Despite the weakness of

these interpretations along the *entertainment* and *ability* dimensions, the claim that a

subject is competent to entertain all grammatical permutations of all molecular

representations of, e.g., English, is not trivial, because the grammaticality of sentences

of English can be determined independently of the subject's grammatical competence

with regard to English. But when the pertinent RS is taken to be the subject's own

LOT, such independent determination of grammaticality is not possible.[25] The problem

thus arises because Fodor and Pylyshyn fail to appreciate that systematicity (or the GC)

must be relativized to an RS. This failure to appreciate the relativity of the GC (or

systematicity) to an RS also undermines, though in different ways, arguments advanced

by Carruthers (2009) and Heck (2007).

Carruthers (2009, but originally in Carruthers 2006) argues that taking a relatively

strong interpretation of the GC as a necessary condition for conceptual thought will

"very likely require us to deny thoughts and concepts to most if not all non-human

animals" (2009, p. 94). The example Carruthers considers concerns bees:

> A bee that is capable of thinking that there is nectar 200 meters north of the hive
> … and that is capable of thinking that the brood chamber is now above it, might
> nonetheless be *incapable* of thinking that there is nectar 200 meters north of the
> brood chamber. This is because the bee's spatial navigation and mental map-
> building outside the hive are based on solar bearings, whereas bees navigate inside
> the hive in the dark, where they employ quite other (gravity based and olfactory)
> ways of representing spatial relationships. And bees might very well lack any

---

[25] Note that the triviality could be avoided if Fodor and Pylyshyn specified a stronger, more discriminating, interpretation along the entertainment dimension. So, for example, if Fodor and Pylyshyn, following Evans, required of a putative entertainer of sentence of mentalese *S* that she know what the world would have to be like in order for *S* to be true, then the triviality result would be avoided: For clearly the subject might possess a *syntactic competence* that determines *S* to be a sentence and also determines what the grammatical permutations of *S* are, and yet lack a *semantic competence* providing knowledge of the truth conditions of all such grammatical permutations. But, obviously, the price of non-triviality might be falsity; it is just such a separation of syntactic and semantic abilities that Travis (2004) argues in support of.

means of integrating the two sets of spatial representations into a single thought. (2009, p. 94-95)[26]

The above passage is supposed to describe why bees, because they contain two distinct cognitive systems, would fail to satisfy what Carruthers (2009, p.96) refers to as the "Strong GC." The "Strong GC" specifies a relatively strong interpretation along the substitution dimension:

Strong Generality Constraint

If a creature possesses the concepts *F* and *a* (and is capable of thinking *Fa*), then, for *all* other concepts *G* and *b* that the creature could posses, it is metaphysically possible for the creature to think *Ga*, and in the same sense possible for it to think *Fb*.

Let us use Roman letters in boldface to represent a bee's "outside the hive" (solar-based) thoughts, and Greek letters to represent its "inside the hive" (olfactory/gravity-based) thoughts. Thus, let **Fa** represent the bee's thought that *there is nectar 200 meters north of the hive* and Φα represent the bee's thought that *the brood chamber is now above it*. Carruthers' point in the above passage is that if, as empirical evidence suggests, the bee is unable to think **F**α (and/or Φ**a**), the bee fails to satisfy the above "Strong GC."

But note that the above statement of the "Strong GC" omits any explicit mention of a grammaticality constraint, and no indication is given as to what the pertinent RS is supposed to be. This is significant because the above passage provides a reason for thinking bees fail to satisfy the "Strong GC" only if hybrid molecular representations such as **F**α (and/or Φ**a**) are grammatical sentences of the pertinent RS. But Carruthers nowhere presents a reason for thinking that there is a RS that would include such hybrid

---

[26] Carruthers does not endorse the conclusion of this argument. The principal premise of the argument is what Carruthers calls the "Strong GC," but Carruthers rejects the "Strong GC" in favor of what he calls the "Weak GC." Thus, though Carruthers rejects the argument as unsound, he does endorse the conditional claim that *if* the "Strong GC" were taken as a necessary condition for conceptual thought, then bees would not be capable of conceptual thought. We take issue with this conditional claim.

molecular representations as grammatical sentences. Indeed, Carruthers' description of the bee's cognitive architecture seems to imply that the outside-the-hive and inside-the-hive representations belong to *different* RSs, with distinct grammars. And this implies that hybrid representations such as, e.g., **F**α, are not grammatical representations in either RS. So, contrary to what Carruthers suggests, the inability of bees to entertain such hybrid molecular representations does not constitute a failure to satisfy the "Strong GC."

Similar problems undermine an argument recently presented by Heck (2007) in support of the conclusion that conceptual *states* and non-conceptual *states* have different *kinds of content*. The principal premise Heck invokes in support of this conclusion is the conjunctive claim that (i) conceptual *states* satisfy the GC, whereas (ii) non-conceptual *states* do not satisfy the GC. Our focus here will not be on the inference from the premise to the conclusion, but rather on the principal premise; we will demonstrate that conjunct (ii) is either incoherent, or at best true by stipulation. The first step in appreciating the problem is noting that Heck invokes the GC not to distinguish between *theories of thought* (as in Evans 1982), nor to distinguish between *creatures* (as in Carruthers 2009 and Camp 2009), but rather to distinguish between types of *representational states*. But what is it to say of a *type of representational state* that it satisfies, or fails to satisfy, the GC?

In an attempt to answer this question, let us examine what Heck says in support of the principal premise. Heck considers beliefs to be paradigmatic of conceptual states, or *thoughts*, and he claims that such states satisfy the generality constraint:

> … beliefs arguably satisfy what Evans called the 'generality constraint': A thinker who is capable of entertaining the thought that *a* is *F* and is also capable of entertaining the thought that *b* is *G* will typically also be capable of entertaining the thoughts that *a* is *G* and that *b* is *F*. (2007, pp. 121-122)

Visual perceptions (of colored objects) are considered by Heck to be paradigmatic

examples of non-conceptual states, and he argues that, though it is an empirical

question, it is at least plausible to suppose that such states do not satisfy the GC:

> I am now going to assume that it can be represented *explicitly* in one's visual
> perception of a surface that it is uniform in color and, further, that only a small,
> connected surface can be explicitly represented as uniform in color – that is,
> neither a large surface occupying, say, half of my visual field nor two small
> surfaces that make no contact can be explicitly represented as being of the same
> color. … if something along these lines were true, the generality constraint
> would fail: "*x* is the same color as *y*" would be explicitly representable only
> under certain circumstances, for example, when *x* and *y* were points on a small,
> connected, surface. (2007, p. 132).

Let us note that in the above applications of the GC, there is, again, no explicit

mention of a grammaticality constraint, nor any indication as to what the pertinent

RS(s) is (are). But, just as the reason Carruthers' offers for thinking that bees do not

satisfy the GC is an even better reason for thinking that a bee's in-the-hive states and its

outside-the-hive states utilize distinct representational systems, so the reasons Heck

offers for thinking that *thought* satisfies the GC while *visual perceptions* do not are even

better reasons for thinking that thought and visual perception utilize different

representational systems, with different grammars. Heck maintains that if a subject is

able to entertain the *thoughts aRb* and *cRd*, then she is able to entertain all the relevant

permutations of these molecular representations (viz. *aRc*, *aRd*, *cRb*, *dRb*, *aRa*, *bRb*,

cRc and *dRd*); but if she is able to entertain *visual perceptions aRb* and *cRd,* then she is

not necessarily able to entertain all these permutations. Heck claims that in the case of

visual perception the reason that a subject will not be able to represent all the relevant

permutations is that some of these permutations will not be "explicitly representable" in

the visual system. Now, that all relevant permutations are "explicitly representable" in

the case of thought, yet not all such permutations are "explicitly representable" in the

case of visual perception, is a compelling reason to think that that the RSs of thought and visual perception are distinct, and have distinct grammars.

Let us then use boldfaced Roman letters to stand for the representations utilized in the representational system of *thought*, and Greek letters to stand for the representations utilized in the RS of *visual perception*. Simplifying somewhat, we can then express the two conjuncts of Heck's principal premise as follows:

(i) *Conceptual states* satisfy the GC: For all molecular representations in the RS for *thought* of the form Fa and Gb, if a subject is able to entertain Fa and Gb, then she is also able to entertain Fb and Ga.

(ii) *Non-conceptual states* do not satisfy the GC: For some molecular representations in the RS for *visual perception* of the form $\Phi\alpha$ and $\psi\beta$, a subject is able entertain $\Phi\alpha$ and $\psi\beta$, yet not able to entertain $\Phi\beta$ and $\psi\alpha$.

Now let us ask why Heck thinks (ii) is true. That is, what explains why the subject is able to entertain perceptions $\Phi\alpha$ and $\psi\beta$, yet not able to entertain perceptions $\Phi\beta$ and $\psi\alpha$? The explanation is clearly *not* that while $\Phi\beta$ and $\psi\alpha$ are perfectly grammatical sentences of the RS for visual perception, the subject for some reason lacks the *ability* to *entertain* these sentences. For, first, Heck says absolutely nothing in the way of specifying interpretations along the *ability* and *entertainment* dimensions – such issues are apparently irrelevant to his argument. And second, if the explanation concerned *the subject's* inability to entertain allegedly grammatical sentences $\Phi\beta$ and $\psi\alpha$, then the failure to satisfy the GC would indicate something about *the subject*, and not about the *type of representational state* $\Phi\beta$ and $\psi\alpha$ exemplify. Therefore the explanation of the subject's inability to entertain $\Phi\beta$ and $\psi\alpha$ must concern $\Phi\beta$ and $\psi\alpha$ themselves. The idea must be then that the subject is unable to entertain $\Phi\beta$ and $\psi\alpha$ because $\Phi\beta$ and $\psi\alpha$ are not "explicitly representable," and that must mean that the subject is unable to entertain $\Phi\beta$ and $\psi\alpha$ because $\Phi\beta$ and $\psi\alpha$ *are not grammatical sentences of the RS for visual representations*.

We are now in a position to answer our initial question, viz. "What is it for a type of *representational state* to either satisfy, or fail to satisfy, the GC?" The idea is that a type of representational state *r* satisfies the GC if and only if the RS for *r*, and its attendant grammar, allows for a certain sort of *syntactic transformation*. In particular, *thought* satisfies the GC if and only if, if **Fa** and **Gb** are grammatical representations in the RS for thought, then **Fb** and **Ga** are also grammatical representations in the RS for thought. Thus the particular syntactic transformation invoked in the GC, which is displayed in Heck's statement of the GC applied to thought, is what we will call *subject-predicate permutation*:

> An RS allows for *subject-predicate permutation* iff for every pair of grammatical subject-predicate sentences *Fa* and *Gb* in the RS, *Fb* and *Ga* are also grammatical subject-predicate sentences in the RS.

And now the two conjuncts of the principal premise can be more explicitly stated as follows:

> (i) *Conceptual states* satisfy the GC: The RS for *thought* is such that grammaticality is preserved under *subject-predicate permutation*.

> (ii) *Non-conceptual states* do not satisfy the GC: The RS for *visual perception* is such that grammaticality is not preserved under *subject-predicate permutation*.

Our objection concerns (ii). We claim that it is incoherent, or at best true only by stipulation. Consider a simple, yet analogous, example: Suppose we have two RSs that are used to represent quadrants of a two-by-two grid. One RS, the Cartesian RS, utilizes molecular representations such as '<1,1>' and '<2,1>'. The latter designates the right-lower quadrant, and first designates the left-lower quadrant. The other RS, the "order-free" RS, utilizes molecular representations such as '(a, 2)', '(2, a)' and '(1, b)', where the first and second both designate the left-upper quadrant, and the third designates the right-lower. The grammars of these RSs are different in several respects: 'a' and 'b' are atomic representations in the order-free RS, but not in the Cartesian RS.

And the Cartesian RS is sensitive to *order* in a way the order-free RS is not: e.g. '<1,2>' and '<2, 1>' are different sentences of the Cartesian RS, but '(a, 1)' and '(1, a)' are mere stylistic variants of the same sentence in the order-free RS.[27]

The Cartesian RS permits the following syntactic transformation: If *<x, y>* is grammatical, then so is *<y, x>*. Let us refer to this transformation as *order-switching*. Now, does the order-free RS also permit *order-switching*? The question has no answer, for the syntactic notion of *order* that is essential to the transformation applies to the Cartesian RS, but not to the order-free RS: one cannot *switch the order* of the constituents in an order-free sentence because the constituents are never in any *order* to begin with. Of course one could *stipulate* that this incoherence implies that the order-free RS does not permit order-switching. Under such a stipulation, we could say things such as "The Cartesian RS satisfies the order-switching constraint, whereas the order-free RS fails to satisfy the order-switching constraint." But this would be a stipulation; it would me more accurate to say that it does not make sense to ask whether or not the order-free RS permits order-switching. Moreover, even if one did stipulate that the order-free RS does not permit order-switching, the claim that "the Cartesian RS, but not the order-free RS, satisfies the order-switching constraint" would not imply any difference in *kinds of content*; the content of the "sentences" of both RSs are quadrants on the same two-by-two grid.

The same points apply, *mutatis mutandis*, to conjunct (ii) of Heck's principal premise. If Heck's empirical supposition is correct, and thus we have reason to think that the RS of visual perception is substantially different than the RS for *thought*, then conjunct (ii) is incoherent, or at best true only by stipulation. Just as in the above example the syntactic transformation of *order-switching* makes essential appeal to the

---

[27] The commas, brackets, and parentheses could be omitted from these RSs without loss of expressive power.

syntactic notion of *order* that applies in the Cartesian RS but not in the order-free RS, so the syntactic transformation of *subject-predicate permutation* makes essential appeal to the syntactic notions of *subject* and *predicate* that apply in the thought RS, but not in the visual perception RS. If the empirical suppositions Heck describes in the above passage are correct, then the syntactic notions of *subject* and *predicate* are not applicable to the RS of visual perception; for if the notions of *subject* and *predicate* were applicable to the RS of visual perception, then, obviously, the RS of visual perception would permit *subject-predicate permutation*. And thus the claim made in conjunct (ii), that the RS for visual perception does not permit *subject-predicate permutation*, is incoherent, or at best true by mere stipulation. And even if one were to stipulate its truth, it would be difficult to see how it could have any significant implications with regard to the *kinds of contents* of thoughts and perceptions have. For, as the above example illustrates, syntactic differences between RSs need not correspond to differences in the *kinds of content* represented by molecular representations of the RSs.

If one accepts Heck's empirical suppositions pertaining to visual perception, then one has a good reason to think that thought and visual perception utilize different RSs, each with a different grammar. That one RS permits for a certain sort of transformation that cannot be applied to the other RS is what one would expect – after all, they are different RSs with different grammars, and the GC is *intentionally designed to be satisfied by the RS for thought*.[28] Once it is appreciated that the GC must be interpreted relative to an RS with an attendant grammar, Heck's claim that thought, but not visual perception, satisfies the GC is revealed to be tantamount to the very plausible claim that there is a particular sort of grammatical transformation that applies to some of

---

[28] Note that if one did not appeal to the RS of *thought* to determine the grammatical transformation relevant to the GC, but instead appealed to the RS of *visual perception* to determine the grammatical transformation relevant to an analogous *GC**, then one could derive the analogous, and opposite, conclusion that visual representation, but not thought, satisfies the *GC**.

the representations in the RS for thought, but does not apply to any of the representations in the RS for vision. But it is rather dubious that the significant philosophical conclusion Heck is attempting to establish – viz. that thoughts and perceptions have significantly different *kinds of content* – follows from this plausible premise.[29]

Conclusion

In the first section we showed that what is called the Generality Constraint (GC) is actually better viewed as a template for a range of principles of cognitive psychology, and we proposed the Generality Constraint *Schema* (GCS) as a clear presentation of this template. The GCS requires a specification along three dimensions of interpretation: the ability, the entertainment and the substitution dimensions. In the second section we presented examples of problems that can arise if, in utilizing the GC in argumentation, one does not carefully specify an interpretation. We suggest that if interpretations along the three dimensions of the GCS are specified, arguments that invoke the GC can avoid these sorts of problems.

---

[29] Consider the following principle: *For all representation S and S' , if there is a grammatical transformation T s.t. S admits T, and S' does not admit T, then S and S' have different kinds of content.* This principle will entail the existence of many different "kinds of content." For there are transformations admitted by subject-predicate sentences that are not admitted by quantified sentences, and transformations admitted by conditionals that are not admitted by negations, and, probably, transformations admitted by some sentences of Russian that are not admitted by no sentences of Japanese, etc.

References

Camp, E. 2004. "The Generality Constraint, Nonsense and Categorial Restrictions".

Philosophical Quarterly 54: 209-231.

Camp, E. 2009. "Putting Thoughts to Work: Concepts, Systematicity, and Stimulus-

Independence" *Philosophy and Phenomenological Research* 78(2): 275-311.

Carruthers, P. 2004. "On Being Simple Minded". *American Philosophical Quarterly* 41

(3): 205-220.

Carruthers, P. 2006. *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Oxford University Press.

Carruthers, P. 2009. "Invertebrate Concepts Confront the Generality Constraint (and Win)". In R. Lurz (ed.), *The Philosophy of Animal Minds*. Cambridge University Press.

Chomsky, N. 1968. *Language and Mind.* New York: Harcourt, Brace and World.

Evans, G. 1982. *The Varieties of Reference.* Oxford University Press.

Fodor, J. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind.* The MIT Press.

Fodor, J. 1998. *Concepts: Where Cognitive science Went Wrong.* Oxford University Press.

Fodor, J. 2001. "Language, Thought and Compositionality" *Mind and Language* 16(1): 1-15.

Fodor, J. 2002. *The Compositionality Papers.* Oxford University Press.

Fodor, J. 2008. *LOT 2: The Language of Thought Revisited.* Oxford University Press.

Fodor, J. and Pylyshyn, Z. 1988. "Connectionism and Cognitive Architecture: A Critical Analysis". *Cognition* 28: 3-71.

Heck, R. 2007. "Are There Different Kinds of Content?". In J. Cohen and B. McLaughlin (eds.) *Contemporary Debates in Philosophy of Mind*. Blackwells.

Johnson, K. 2004. "On the Systematicity of Language and Thought". *Journal of Philosophy* CI (3): 111-139.

Peacocke, C. 1992. *A Study of Concepts,* MIT Press.

Ryle. G. 1949. *The Concept of Mind.* University of Chicago Press.

Searle, J. 1980. "The Background of Meaning" In Searle J. R., Kiefer F., and Bierwisch, M. (eds.), *Speech Act Theory and Pragmatics*. Dordrecht, Reidel.

Travis, C. (1994) "On Constraints of Generality". *Proceedings of the Aristotelian Society* 94: 165-188.

XXX. under review. "Is Even Thought Compositional?"

XXX. forthcoming. "Conceptuality and Generality: a Criticism of an Argument for Content Dualism".