

On the systematicity of language and thought

Kent Johnson,
Department of Logic and Philosophy of Science, UCI
johnsonk@uci.edu
<http://hypatia.ss.uci.edu/lps/home/fac-staff/faculty/johnson/>
Draft: July, 2002

The last twenty years or so have seen a lot of discussion about systematicity – roughly the idea that if a person’s repertoire contains the sentence/thought that John loves Mary, then her repertoire also contains the sentence/thought that Mary loves John (e.g., Fodor and Pylyshyn 1988, Fodor and McLaughlin 1990, Smolensky 1987, 1988, Hadley 1994, 1997, Niklasson and van Gelder 1997, Cummins 1996; the bibliographies of these papers contain many more citations). Much of the discussion of systematicity has focused on the technical question of whether and how connectionists can explain the systematicity of cognitive capacities without simply implementing a classical architecture. That is, the debate has centered on whether connectionist systems offer a genuinely novel way to understand the empirical phenomenon of systematicity. Discussions of systematicity have not been confined to the connectionism/classicism debate. Systematicity (sometimes under the name “generality” or “recombinability”) has also been discussed frequently in other areas of the philosophy of mind (e.g., Evans 1982, ch. 4, Travis 1994, Peacocke 1992, ch. 2). Two common threads unite these discussions of systematicity. First, claims about systematicity – especially the systematicity of language – are typically taken to be trivially true. For instance, Matthews writes of “our capacity for natural language, which is by all accounts paradigmatically systematic” (Matthews 1997, p. 175). Cummins tells us that “the systematicity of language ... is relatively unproblematic” (Cummins 1996, p. 591). Fodor and Pylyshyn assure us that “linguistic capacity is a paradigm of systematic cognition” (Fodor and Pylyshyn 1988, p. 37). In the case of thought, Peacocke writes, “The recombinability of concepts to form new thoughts has been largely unquestioned in the published literature” (Peacocke 1992, p. 42). (He then goes on to assume this form of systematicity, and to develop an explanation for it.) Second, the nature of systematicity is rarely more than cursorily sketched: an example or two is typically thought to adequately characterize the phenomenon. In short, the prevailing optimism of the literature suggests that everybody

knows what systematicity is, and everybody agrees that language (and probably thought, too) has it.

Despite this widespread enthusiasm, I am skeptical about both of these properties of systematicity. I don't think it is very clear what systematicity is, and I'm even more skeptical that language (or thought) has it. Since I find it easier to study overtly realized phenomena like language, I will focus on the systematicity of language. I will argue that according to way systematicity is conceived of in the literature, natural languages are not systematic. More importantly, I will contend that the way systematicity is conceived of in the literature rests on a crucial but inaccurate assumption about the nature of linguistic natural kinds. When this assumption is corrected, it becomes utterly mysterious what it even means to say that natural language is systematic. Although my primary goal is to explore the nature of the systematicity of language as such, I will also suggest that this exploration has consequences for the systematicity of thought and for the role that systematicity plays in the connectionism/classicism debate. The systematicity of thought, I argue, suffers from problems analogous to the systematicity of language, and it does not appear to play any interesting role in the debates about cognitive architecture. In sum, it would only be a slight exaggeration to say that my view of systematicity is maximally pessimistic: it doesn't exist; but if it did, languages wouldn't have it; and even if systematicity existed and languages had it, it still wouldn't be all that interesting.

The paper is structured as follows. I begin by developing a natural and intuitive characterization of the claim that language is systematic. In section 2, I show that this claim is false. In section 3, I raise and discuss an objection to the argument against the systematicity of language. I show that the argument rests on an assumption that is either false or undefended with no obvious defense available. In section 4, I show how my reply to this objection also applies to natural and intuitive formulations of systematicity, leaving these formulations false or undefended as well. In section 5, I consider the systematicity of thought in light of the current discussion. In section 6, I briefly explore what role systematicity plays in the connectionism/classicism debate. I conclude in section 7.

1 An overview of systematicity

In order to evaluate the claim that language is systematic, we need to get a better understanding of what the claim is. Unfortunately, most characterizations of systematicity are less than perfectly clear (cf. Niklasson and van Gelder 1994, Cummins 1996, Hadley 1997 for comments to this effect). Of the very few explicit characterizations of systematicity, Robert Cummins offers one that gets really close to the heart of the matter:

A system is said to exhibit systematicity if, whenever it can process a sentence *s*, it can process systematic variants of *s*, where systematic variation is understood in terms of permuting constituents or (more strongly) substituting constituents of the same grammatical category (Cummins 1996, p. 591).

I will adopt a version of Cummins' characterization, with a couple minor qualifications. First, instead of speaking about a system "processing" sentences, I will speak about whether a sentence is grammatical or not. I do this primarily because I want to focus on whether language is systematic, and our linguistic abilities may be distinct from our abilities to process language. For example, although speakers reject **the child seems sleeping* as ungrammatical, there is also a sense in which it has a clear meaning, and might be thought to be processible, in some extended sense of the word.¹ Indeed, utterances like **him give she cookie now!* and *take and slab!* might also be processible, too. I wish to treat all these examples as falling outside of natural language – or at least outside of English, which is the natural language I will focus on. (I suspect that I haven't really adjusted Cummins' idea, since the quotation above might be referring to the system that constitutes our linguistic abilities.)

Secondly, for expository convenience, I will adopt what Cummins calls the "stronger" version of systematicity. According to this version, what is relevant to systematicity is not whether we can interchange *John* and *Mary* in *John saw Mary* and still have a grammatical sentence, but whether we can substitute in *any* words of the same grammatical category as *John* and *Mary*, and still preserve grammaticality. Since there are only finitely many primitives in our language, and sentences can be arbitrarily finitely long, it is not clear that this version is really

¹ I will employ a number of standard conventions. (i) For ease of readability, I use italics instead of quotation marks when mentioning an expression. (ii) The asterisk denotes ungrammaticality. (iii) The term *grammar* denotes all the components of a natural language, including its syntax, semantics, and phonology. (iv) I will use such terms as *kind* and *category* interchangeably.

stronger than the other.² In any case, nothing essential rides on adopting this version. All the points could be made with the other version simply by using sentences with e.g., some additional relative clauses, conjunctions, or the like.

With these qualifications made, I can now offer an initial characterization of the systematicity of language. My characterization is primarily based around the following condition on something's being a constituent of language.

- (S) A is a constituent of language L only if for all B of the same linguistic kind as A, and all things C, C can compose with A (in a certain way) to form a sentence iff C can compose with B (in that way) to form a sentence.

We then say that the language L is systematic iff (S) holds for all expressions A. I will often refer to (S) as the statement of systematicity, but it should be remembered that systematicity really involves (S) holding for all expressions. (S) also captures a type of systematicity that has been discussed outside of the connectionism/classicism debate. Gareth Evans, for instance, discusses what he calls the “Generality Constraint”:

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 the thought that John is happy has something in common with the thought that John is sad (Evans 1982, p. 100).

The Generality Constraint requires us to see the thought that a is F as lying at the intersection of two series of thoughts: the thoughts that a is F, that a is G, that a is H, ..., on the one hand, and the thoughts that a is F, that b is F, that c is F, ..., on the other (Evans 1982, p. 209; cf. p. 104, fn. 21).

Although Evans is primarily concerned with thought, he illustrates the Generality Constraint with language (Evans 1982, pp. 101-03). Others have discussed it in a linguistic context (e.g., Travis 1994). Thus, (S) appears to capture a property that is often believed to hold of language.

Before moving on to a criticism of (S), let me note a feature of it that will loom large in the discussion to follow. (S) is essentially a condition on the individuation of linguistic kinds. (I will use such terms as *linguistic kind*, *grammatical category*, etc. interchangeably.) (S) is in fact equivalent to (vb1), which gives a necessary condition on two things being members of the same linguistic kind:

² There are, however, infinitely many words in the typical speaker's repertoire, because there are recursive word-forming

- (vb1) If A but not B can compose with C in a certain way to form a sentence, then A and B do not belong to the same linguistic kind.

Although I am going to argue that there is a serious flaw in statements of systematicity like (S) – and the versions like Cummins’ from which (S) is extracted – there is nevertheless something very right about (S). (S) is correct to emphasize that any statement of systematicity is essentially relativized to a theory of linguistic kinds. After all, what is supposed to make language systematic is that since there are sentences like *John sings*, there are also sentences like *Mary sings* and *John jumps*. We do not want to say that language is not systematic simply because we do not also have **Jumps sings*, where a verb is substituted for a noun. The importance of a theory of kinds of constituents to systematicity would seem obvious. However, sometimes it goes unmentioned (cf. e.g., Fodor and McLaughlin’s (1990) discussion of the systematicity of thought, where it is never mentioned that the elements that systematicity predicts can be legitimately switched around must be of the same kind). Although systematicity’s dependence on a theory of linguistic (or mental) kinds may be too obvious to bear mention, I will argue that it is this dependence that renders standard formulations of systematicity like (S) both false and otherwise unacceptable. My first goal is to show that there is no obvious sense in which (S) is true.

2 The argument against the systematicity of language

We have just seen that (S) – and indeed any version of the systematicity of natural language – presupposes some theory of linguistic kinds of expressions. So before we can explore (S) we need to fix on some theory of kinds of expressions. Following common trends in philosophy and linguistics, we might take the relevant kinds to include verbs (transitive and intransitive), quantifiers, connectives, adjectives, nouns, and singular terms. I will work with these kinds of expressions simply because they are the kinds one finds discussed in this literature, when kinds are mentioned at all. In the long run, the particular kinds chosen will not matter. In the next section, I will argue that the lack of systematicity exhibited with these kinds is actually a weak

devices, e.g. *great-grandmother*, *great-great-grandmother*, *great-great-great-grandmother*.

instance of a more serious problem that holds for any realistic choice of kinds of expressions. In any case, given some collection of linguistic kinds, (S) predicts that we cannot find a pair of terms x and y belonging to one of these kinds and a linguistic environment C such that the composition $C(x)$ is grammatical but $C(y)$ is not. In all cases, this prediction is false, and I will now give counterexamples for each of the categories listed above. It bears emphasis, though, that the non-systematicity of natural language is a *pervasive* phenomenon: many more instances of the sorts of examples I offer can be found throughout the linguistics literature.³

I begin with verbs. I will dwell a bit on this example, in order to display some of the issues involved, and I will move more quickly through the other kinds of words. Consider the verbs *put* and *stow*. On the one hand, *put* appears in the structure given in (9a), but not in the structure given in (9b). In contrast, the distribution of *stow* is precisely the opposite: it appears in the second structure but not the first, as (9c-d) show:

- (9) a. John put his gear down.
 b. * John put his gear.
 c. * John stowed his gear down.
 d. John stowed his gear.

Even though you can neither put nor stow something without putting or stowing it somewhere, only the verb *put* demands that some indication be given as to where the item was placed. The distribution in (9) appears to be due, in large part at least, to syntactic properties of the word. However, the grammatical behavior of a given word is influenced not only by its (morpho-) syntactic properties, but also by its semantic and phonological properties. For instance, consider the well-known “locative alternation” (e.g., Higginbotham 1989, Pinker 1989, Gropen et al. 1991). Words like *drench* and *throw* exhibit complimentary distribution over two grammatical forms:

- (s3) a. I drenched the flowers with water.
 b. *I drenched water onto the flowers.
- (s4) a. *I threw the flowers with water.
 b. I threw water onto the flowers.

³ It's also worth observing that many languages make many other distinctions that are not found in English.

(The interpretations of (s3b) and (s4a) on which the sentences are grammatical are not relevant here.) It is common to maintain that *drench* does not fit into the second form because the meaning of *drench* does not specify how the water got onto the flowers, but only specifies the end state (being drenched) of the flowers. In contrast, *throw* does say something about how the water got onto the flowers, but it does not specify what happened to the flowers as a result of throwing water on them (e.g., they may or may not have been drenched). Here, then, is a case where the semantics appears to affect the grammatical properties of some words. In the case of phonology, consider the two verbs *show* and *describe*. They appear to be syntactically and semantically very similar. However, if you wanted to say that Alice showed or described the book to Martha, only in the case of *show* would you use the double-object construction:

- (10) a. Alice showed Martha the book;
b. * Alice described Martha the book.

(10a-b) differ because English places phonological restrictions on what verbs can enter into the double-object construction. Roughly speaking, the differing behavior in (10) is due to the fact that *show* but not *describe* is typically stressed on its first syllable. This is seen elsewhere with e.g., *I told Martha a story* vs. * *I reported the police the accident*. (cf. e.g. Pinker 1989, pp. 45-7 for more discussion).

I turn now to a briefer survey of the remaining categories. Quantifiers differ widely in their grammatical behavior, even when they are apparently similar in meaning. For instance, *each* and *every* both seem to express universal quantification, but they are not systematic in distribution (cf. e.g., Beghelli and Stowell 1997):

- (6) a. * The men every read the book carefully
b. The men each read the book carefully
c. One boy ate almost every apple in the basket
d. * One boy ate almost each apple in the basket

Similarly, the English words *and* and *but* are often thought to be extremely similar. In fact, it is commonly supposed that the two differ only in their pragmatic implicatures (e.g., Grice 1989, p. 46). However, their differences are linguistically deeper.

- (4) a. John and Mary carried the piano upstairs together.
 b. * John but Mary carried the piano upstairs together.
 c. But for Kristen's help, I would have failed the test
 d. * And for Kristen's help, I would have failed the test

Other predicates like adjectives can also have restrictions on where they appear. Suppose the basket is full of apples, and John is fond of pie. Our language will allow you to say that the basket is full without specifying what it is full of. Not so for fondness:

- (8) a. The basket is full.
 b. * John is fond.

As for nouns, we have, e.g., the dissimilar distribution given in (7).

- (7) a. McDonald's destruction by Burger King had a major impact on beef production.
 b. *McDonald's resemblance by Burger King had a major impact on beef production. (cf. *McDonald's resemblance to Burger King*)

(Cf. e.g., Tenny 1994 for discussion.) Our final category is the singular terms, which encompass such things as proper names, pronouns, and demonstratives. Singular terms are often heralded as the paradigms of systematicity: since *John loves Mary* is a sentence, so is *Mary loves John*. Strictly speaking, the two occurrences of *John* (and *Mary*) are different in some respects. Although English typically does not give phonological expression to its case assignment, a great many languages do overtly distinguish between nominative and accusative case. In Attic Greek, for instance, *Socrates loves Xanthippe* is expressed by (gr1a), and *Xanthippe loves Socrates* is expressed by (gr1b):

- (gr1) a. ho Socratês philei tèn Xanthippên.
 b. hê Xanthippê philei ton Socratên

Moreover, there is ample reason for supposing that case assignment is a fundamental property of all languages. In English, case assignment is overt in pronouns: *He loves her* is a sentence and so is *She loves him*, but neither **Him loves she* nor **Her loves he* are. I will not dwell on the

facts about case assignment, although such facts are surely relevant to a precise formulation of the systematicity of language. A further look at the behavior of singular terms shows that they exhibit other different properties. For instance, it is possible to treat names as common nouns and quantify over them. However, the same cannot be done with other singular terms. If you wanted to announce that two women in your class, both of whom were named *Mary*, needed to report to the principal's office, you might use a sentence like (2a). However, even if the two women were standing near you so that you could easily demonstrate them, you would never use (2b) or (2c).

- (2) a. The two Marys should report to the principal's office immediately.
 b. *The two shes should report to the principal's office immediately.
 c. *The two these should report to the principal's office immediately. (Observe that the subject is distinct from the partitive construction *the two of these*; cf. **the two of Marys*, **the two of shes*.)⁴

Furthermore, pronouns and demonstratives differ from one another in that only the former has a possessive form. Oddly enough though, complex demonstratives can take the possessive (cf. Johnson and Lepore forthcoming):

- (3) a. Katie's leg is broken / Bob and Susie's legs are broken.
 b. * His leg is broken / Their legs are broken.
 c. *This's leg is broken / Those's legs are broken.
 d. This dog's leg is broken / These dogs' legs are broken.

As an aside, it is worth noting that the failure of systematicity extends beyond individual words to complex constituents of sentences. It has been suggested (Schiffer, forthcoming) for instance, that the direct object of *believes* is the same in both of (11a-b):

- (11) a. John believes that snow is white
 b. John believes the proposition that snow is white

There are deep and difficult issues concerning the identity of the meaning of the content clause of propositional attitude reports. However, regardless of these issues, human languages do not

treat the clause *that S* and the noun phrase *the proposition that S* alike. For instance, although both expressions are fine in (12), only *that S* can be moved to the end of the sentence in (13):

- (12) a. That snow is white is commonly known.
 b. The proposition that snow is white is commonly known.
- (13) a. It is commonly known that snow is white.
 b. * It is commonly known the proposition that snow is white.

(Cf. Chomsky 1986, ch. 3 for more discussion.) There are numerous other such differences between the two claims. Although the town gossip might speak of *the guy who Jane believes that Jill likes*, nobody speaks of **the guy who Jane believes the proposition that Jill likes*. Indeed, language's resistance to systematicity between these two forms of content clauses is so strong that even the presence or absence of *that* in a propositional attitude report can matter. It was observed long ago (Ross 1967) that if the subject of a *that*-clause is pulled out of its normal place, as happens when we form questions, the complementizer *that* is better left out of the sentence. For instance, **Which politicians do you think that voted against it?* reliably jars speakers when they encounter the word *that*, but remove the *that* and the sentence is fine. These differences between expressions of the form *that S* and *the proposition that S* (and even *S* alone) underscore the complex differences between individual elements of language, a difference that (S) appears to deny.

The failure of systematic behavior of expressions of natural language was noted early on in the linguistics literature, even before linguistics was considered a branch of cognitive science. According to Zelig Harris, words should be individuated solely by their distributional properties (Harris 1951). That is, Harris held that natural language lacks systematicity so much that *x* and *y* are distinct words if and only if there is a *C* such that *C(x)* is grammatical but *C(y)* is not. It is probably an exaggeration to say no two words pattern in exactly the same way. However, the fact that some linguists have found this view plausible shows just how differently seemingly similar words behave. In the next section, I will consider an attempt to restore the systematicity of language despite this high degree of variability across words.

⁴ Note that even though the latter two sentences are ungrammatical, they still have a clear meaning. But many sentences have a

3 The possibility of very fine-grained linguistic categories

We have just observed that if we take seriously the kinds of grammatical categories that are often appealed to, then natural language is not systematic. At this point, though, it would be natural to object that the evidence just reviewed shows only that we have not found the right linguistic categories to use in a statement of systematicity. Let me put this objection in the strongest form I can.

Look, there is an a priori reason why your argument can't work. You have presented numerous linguistic differences between ostensibly similar words, but it does not follow from this that you have undermined the systematicity of language. We are always free to interpret these differences as showing only that the relevant linguistic types must be more fine-grained than appears at first glance. Here's how this strategy works. Whenever you come up with some new linguistic details, we will simply create more fine-grained taxonomies of linguistic types in such a way that systematicity is preserved. The idea is that if you show that x and y in category A behave differently, then we will just break A into two types: the A s that behave like x and the A s that behave like y . This strategy is of course completely general. So the fact that two words are grammatically different is evidence that the words are in different categories, and that language is systematic over these fine-grained categories.

Since the objection just given is in essence a proposal to use very fine-grained linguistic categories, I will refer to it as FGC. If FGC is right, then nothing I have said so far is relevant to the issue of the systematicity of language. Moreover, FGC is driven only by a straightforward interpretation of systematicity. In particular, (S) says that if two words behave grammatically differently, then they belong to different linguistic kinds. FGC operates by assuming that the appropriate taxonomy of semantic natural kinds is a *partition* over the set of words: each word belongs to one and only one type. (A partition of a set X , recall, is a set of subsets of X such that each element of X belongs to exactly one of these subsets.) Since each set in this partition is by definition a set of grammatically equivalent expressions, FGC guarantees systematicity-preserving categories. Furthermore, FGC's strategy appears to be completely innocent. Surely the grammatical types are what are responsible for differing grammatical behaviors. So if two words have different linguistic properties, then they must be of distinct linguistic types. What could be wrong with a proposal like FGC?

FGC appears very compelling at first glance. However, I will argue that it contains an important and relevant methodological flaw. My basic strategy is as follows. I will begin by arguing that the only relevant kinds of words for a theory of systematicity are natural kinds. This

meaning even though there is something structurally wrong with them: **The child seems sleeping* is one famous case of this.

looks like a small point, but it will be crucial to the remaining argument. FGC assumes that if A and B are legitimate natural kinds, there is also a third natural kind A-and-B, whose extension is the intersection of the extensions of A and B. But it is a principle of the philosophy of science that natural kinds do not in general distribute over Boolean combinations of natural kinds. For instance, an object that is both gold and spherical will have metallurgical properties in virtue of being gold and mechanical properties in virtue of being spherical. But there is no natural kind gold-sphere that is employed in any of the sciences to explain the behavior of the object. A precisely similar case holds in the organization of language. If linguistic natural kinds are not closed under Boolean combinations, then the fine-grained categories that FGC relies on will not be available, and FGC is therefore untenable.

Let me begin by saying why it is critical that the kinds referred to in (S) – or any other statement of systematicity – are linguistic natural kinds of expressions. The reason is, of course, that some kinds of words are not relevant, and if we attended to all possible kinds of words, putatively empirical linguistic generalizations involving reference to kinds of words become trivial and unexplanatory. In the cases of (S) and FGC, the claims become trivially false or vacuous, because the assumption that each word belongs to a unique natural kind is trivially false. So obviously, only some kinds of words are relevant. But as Goodman showed, unless we can say which kinds are relevant, the basic problem remains. There is, for instance, a class of words that Sally likes, and a class of words that she does not like (Sally is neutral about all possible words that she doesn't know or that have not yet entered the language, so she doesn't like them.) Does the existence of these classes guarantee that English is “Sally-systematic” in the following sense?

- (SS) A is a constituent of language only if for all B of the same Sally-kind as A, and all things C, C can compose with A (in a certain way) to form a sentence that contains a word that Sally likes iff C can compose with B (in that way) to form a sentence that contains a word that Sally likes.

Given how Sally-kinds are defined, (SS) is trivially true, but uninteresting. It would be interesting only if there were some independent characterization of the Sally-kinds, that is, if the Sally-kinds did some theoretical work beyond securing the truth of (SS). For analogous reasons, the relevant kinds to be used in a statement of systematicity cannot just be stipulated to be those

whose extensions are always sets of grammatically equivalent words. In such a case, the claim that language is systematic is a tautology, epistemically trivial, and scientifically useless. So in forming a statement of systematicity, we need to appeal to those kinds of words that really do have some independently useful status. To that end, we should look to linguistics to see what kinds of words are actually used in theorizing. If language turns out to be systematic over these kinds – or over some subset of them that we can argue are the relevant ones – then language is systematic; if it is not systematic over these kinds, language is not systematic. Demanding that we use the kinds actually posited by linguistics restores systematicity to its rightful status as an empirical claim. It also follows that FGC now makes empirical predictions of a sort. FGC predicts that linguistic theorizing will produce the fine-grained natural kinds that FGC needs. I will argue that it does not.

To start the argument against the fine-grained categories that FGC needs, let me introduce two kinds of verbs.⁵ The first kind of verb is the causatives. These are verbs like *break*, *snap*, *soften*, *warp*, etc. Verbs like these all express that the subject of the verb brought it about that the object of the verb entered into some state. For example, *X break Y* – as in *Susan broke the vase* – has roughly the meaning *X cause Y to be broken*. There are numerous reasons for thinking that the causative verbs are a natural kind of verb, and that verbs of this kind are distinctive in possessing a bit of structure in the verb’s linguistic representation that means roughly “cause” (e.g., Pinker 1989, Parsons 1990, Travis 2000; this bit of structure is overtly realized as a morpheme in hundreds of languages; cf. Hale and Keyser 1993, p. 102). One reason for distinguishing the causative verbs is that (*ceteris paribus*) they have an intransitive form in which the subject of the transitive verb disappears and the object assumes the subject position, as in *the vase broke*. The second type of verb is the object-deletion verbs, which are in a way the converse of the causatives. Object-deletion verbs are verbs like *eat*, *sew*, and *write*, which although typically transitive, can also appear as intransitive verbs with their object understood, as in *While Mary was eating/sewing/writing, the phone rang*. Importantly, some verbs (e.g., *boil*) are causatives but not object-deletion verbs, other verbs (e.g., *eat*) are not causatives but are object-deletion verbs, and yet further verbs (e.g., *cook*) are both. These facts are witnessed in (sy1) – (sy3):

⁵ The two kinds of verbs I use are for illustrative purposes only. You may, if you like, replace them with other interacting types of verbs.

(sy1) a. Mary boiled the meat.
 b. The meat boiled.
 c. *Mary boiled.

(sy2) a. Mary ate the meat.
 b. *The meat ate.
 c. Mary ate.

(sy3) a. Mary cooked the meat.
 b. The meat cooked.
 c. Mary cooked.

(There are of course readings of the ungrammatical sentences on which they are grammatical; these readings are not relevant here.) How should we handle the distribution in (sy1) – (sy3)? For proponents of FGC, the answer is clear. (sy1) – (sy3), according to FGC, is evidence that there are three linguistic categories instead of two. That is, there must be a category A of causative words like *boil* which are not also object-deletion verbs, a category B of object-deletion verbs like *eat* which are not causatives, and a category C of causative object-deletion verbs like *cook*. Of course, further linguistic considerations will break each of A, B, and C into multitudes of even more fine-grained categories $A_1, \dots, A_m, B_1, \dots, B_n, C_1, \dots, C_k$, but I will not worry about these further subdivisions here. All that is important at present is to see how the FGC strategy works in a particular case.

The problem with FGC should now be clear. The linguistic evidence much more strongly supports the claim that there are two overlapping categories – the causatives and the object-deletion verbs – than it supports the claim that there are the three categories A, B, and C. The categories A, B, and C do nothing whatsoever to explain why A-type verbs and C-type verbs share some of their grammatical behavior, and B- and C-type verbs share other grammatical behavior. As far as A, B, and C are concerned, the similarities between *boil* and *cook* and *eat* and *cook* are utterly mysterious. The similarities of behavior just mentioned are easily explained by a theory that includes causatives and object-deletion verbs among its linguistic kinds: *boil* and *cook* share certain grammatical behaviors because they're both causatives, and *eat* and *cook* share other behaviors because they're both object-deletion verbs. So it's plausible that a linguistic theory will include at least the categories of causatives and object-deletion verbs. But could the theory also contain the categories A, B, and C? It is very hard to see why it would. In

general, anything that A, B, and/or C can explain are also explicable using the categories of causatives and object-deletion verbs. Another way to put this last point is that any laws of linguistics that could be stated using the categories A, B, and/or C are just as easily stated using the categories of causatives and object-deletion verbs. But if A, B, and C do not figure into any laws – linguistic, psychological, or otherwise – then that is good reason to hold that they are not real natural kinds. It is of course an empirical claim that A, B, and C are not needed in a successful linguistic theory. The total pattern of linguistic data could turn out so that overall it is vastly easier to use A, B, and C instead of causatives and object-deletion verbs. But as far as I know, there really isn't any evidence for this outcome, and there certainly isn't enough evidence to make it even slightly plausible that A, B, and C are genuine linguistic categories. We should conclude that there is no reason to hold that A, B, and C are natural linguistic kinds.⁶ With the loss of A, B, and C, the kinds FGC requires in order to account for (sy1) – (sy3) are simply not available. So FGC cannot be maintained as a viable linguistic theory, and so FGC fails as an attempt to preserve the systematicity of language.

The argument just given, which was based around (sy1) – (sy3), should not be surprising. After all, in the previous section, we saw that the distribution of an expression can be affected not only by what (morpho-) syntactic natural kinds the expression belongs to, but also by its membership in various semantic and phonological kinds. Although there are various famous points of interaction between these areas, it also remains true that syntax, semantics, and phonology are to some extent independent of one another. The work of practicing phonologists, for instance, is often isolated from semantics and syntax. Thus, it is to be expected that many words can have a given phonological, semantic, or syntactic property independently of many of the other properties they might have in the remaining two linguistic domains. When such properties are relevant to the grammatical behavior of the expression, we will have further instances of overlapping linguistic natural kinds of the sort seen in (sy1) – (sy3).

⁶ I stress that natural kinds, whether of linguistics, psychology, physics, or any other of the sciences, are determined by the best available scientific theories of the subject. Someone who wants to contend that A, B, and/or C are natural kinds needs to show why they are plausibly needed by some leading scientific theory.

4 What is systematicity?

The argument of the preceding section suggests that the problem with systematicity runs deeper than merely that statements like (S) are false. We saw that the fundamental problem with FGC is that it attempts to create a partition of all the words of the language. However, the empirical facts strongly suggest that the linguistic natural kinds of words do not form anything as neat as a partition; rather, they are a much more complex overlapping collection. A given word belongs to multiple overlapping linguistic kinds, and no one of these kinds provides a total account of the word's grammatical behavior. If this is right, then (S) is probably false, because it implies that a word belongs to exactly one linguistic natural kind. One might try to argue that there is only one linguistic kind per word that is relevant to systematicity. But such a position would be extremely hard to defend, since it would amount to something like the claim that the difference between *John ate* and **John boiled* is relevant to systematicity, but the difference between **the meat ate* and *the meat boiled* is not. In such a case, it would be hard to see what systematicity even amounts to, much less why it matters.

At this point, a likely suggestion is that although (S) is false, a better statement of systematicity can be had. According to this suggestion, a language is systematic iff for all words A, (S') holds:

- (S') A is a constituent of language only if for all B of *all* the same linguistic kinds as A, and all things C, C can compose with A (in a certain way) to form a sentence iff C can compose with B (in that way) to form a sentence.

The problem with (S') is that it is a tautology. The sum total set of linguistic kinds that a given word belongs to exhausts the word's grammatical properties, by the very definition of the linguistic kinds. So if A is a constituent of language, then it is trivially true that any B of all the same linguistic kinds as A will behave grammatically identically to A, and so *a fortiori* A composes with C to form a sentence iff B does too. But systematicity is not supposed to be a triviality. Systematicity is supposed to be an empirical property of human languages, one which is sometimes thought to be easily explicable by recursive, symbol-processing computational devices, but which is not so easily explained by connectionist devices that do not simply implement the former sort of machinery (e.g., Fodor 1987, p. 153, Fodor and Pylyshyn 1988, Fodor and McLaughlin 1990, Niklasson and van Gelder 1994, Smolensky 1995,). If

systematicity was a trivial property of language (or cognition), then there would be no need for any linguistic (or cognitive) theory to explain it. The worst that a theory or model could do would be to employ the wrong linguistic (or mental) kinds in its implementation of language (or cognition). But this sort of error would only be to get some of the empirical facts wrong, which is only a normal occurrence in the progression of any ongoing empirical science. So whatever property of systematicity cognitive scientists worry about, it is not (S'). (I should note in passing that this general problem of making systematicity trivial would apply to the FGC approach as well, if the linguistic kinds FGC requires actually existed.)

Here is where we have gotten. The really deep problem with a standard characterization of systematicity like (S) is that it makes the false (or at least unjustified) assumption that each constituent of language belongs to exactly one relevant category. Moreover, systematicity becomes trivial if we demand only that it apply only to constituents that belong to all of the same categories. At this point, the following thought might arise. Since what linguistic environments a given word can grammatically appear in is relative to the word's grammatical category, perhaps we can adjust the characterization of systematicity so that whether systematicity requires two words to be intersubstitutable depends not only on what linguistic kind(s) the words belong to, but also on what kind of linguistic environment we are dealing with. We might then entertain something like:

(S'') If A and B belong to category X, and C is an X-type linguistic environment, then C(A) is grammatical iff C(B) is too.

Thus, we might say that the environment *John ___ his gear under the bed* is a kind of transitive verb environment, and *put* and *stow* are transitive verbs, but *John ___ his gear* is a location-unspecified environment, and *stow* but not *put* is this type of transitive verb. Could this sort of strategy render language systematic? There are several problems with (S''). For starters, it suffers some by now well-known problems. First, (S'') is just a triviality, saying little more than language is systematic except when it is not. Second, it is not obvious that all of the requisite kinds X actually exist. Many linguistic environments require that an expression be contained in multiple categories (where the intersection of these categories is not a linguistic natural kind), which would make it hard to specify what type of linguistic environment it is. On a related point, using the categories of linguistic theory, (S'') is false. Returning to (10) (repeated below)

above, linguistic theory correctly predicts that both *show* and *describe* are of the same syntactic category, and so both should fit into the double object construction:

- (10) a. Alice showed Martha the book;
 b. * Alice described Martha the book.

The fact that *describe* does not fit into this construction is not due to the fact that it lacks membership in any relevant category, but that it also is a member of a certain phonological category, which prevents it from entering into certain constructions. Thus, being a member of a certain category can *prevent* an expression from entering into a construction that it would otherwise normally participate in. The grammatical behavior of a word is not determined merely by what categories it belongs to, but also by what categories it does not belong to.

At this point, it would be ideal to unveil a victorious characterization of systematicity. I would like to replace (S) (and (S') and (S'')) with something that is responsive to the linguistic facts and which has all the bite needed to maintain current discussions involving systematicity. Unfortunately, I do not have such a characterization. However, our discussion so far provides some clues as to what is needed from a useful notion of the systematicity of language. I will end this section with a catalogue of desiderata on a theory of systematicity.

- (d1) Systematicity should be sensitive to the organization and nature of linguistic natural kinds.

We have seen that linguistic kinds do not neatly partition the set of words in a speaker's lexicon, but are much more intricately interrelated. A successful theory of systematicity will work within the framework of these kinds, however complicated they turn out to be.

- (d2) Systematicity should be a contingent property of human languages.

Given the way that the systematicity of language fits into most contemporary discussions, systematicity should be an empirical property of language that a psycholinguistic model either captures or does not capture. All models capture trivialities.

- (d3) If only some linguistic kinds are used in a characterization of systematicity, there should be a reason why those particular kinds are used, and others are excluded.

As far as I can see, the most promising strategy for saying what systematicity is is to replace (S) with something that makes only some of the linguistic kinds that a word belongs to relevant to systematicity. But in such a case we will need an argument why some kinds are relevant and others aren't. It is utterly unclear to me how this might be done.

- (d4) It should be clear how the resulting version of systematicity is still useful and relevant for one's purposes.

A version of systematicity that meets desiderata (d1) – (d3) may look quite different from the unsuccessful but intuitive versions that we have reviewed. In such a case, it will be important to show how the successful version still accomplishes the work that one wants done from a theory of systematicity.

The desiderata in (d1) – (d4) form a challenge to philosophers of mind and language (and cognitive psychologists, computer scientists, etc.) who appeal to the systematicity of language. The challenge is of course to construct a theory of systematicity that both respects these desiderata and yet at the same time also expresses the combinatorial completeness that is essential to systematicity. It's partly an empirical matter whether or not this can be done, so I can't offer a definitive proof that there is no viable notion of systematicity. But when we see the kinds of troubles systematicity has, we should also ask whether the alleged systematicity of language is important or interesting enough to bother trying to theorize about. After all, whatever truth there is to the claim that language is systematic is due to other properties of the language. In particular, whatever systematicity language has is completely exhausted and explained by such things as the language's recursive combinatorial syntax (and perhaps semantics) over linguistic constituents and the mechanisms by which we acquire new expressions by organizing them into some of a limited number of linguistic natural kinds.⁷ It

⁷ I suspect that there is *some* truth in what Fodor and Pylyshyn say when they write that "Just as you don't find linguistic capacities that consist of the ability to understand sixty-seven unrelated sentences, so too you don't find cognitive capacities that consist of the ability to think seventy-four unrelated thoughts" (Fodor and Pylyshyn 1988, p. 40). What Fodor and Pylyshyn neglect to consider in this passage is that many of the systematic-looking patterns of natural language are often cases where sentences of apparently similar form have radically dissimilar underlying structures. For instance, we might appear to have systematic variants in sentences like *John promised Mary to clean the barn* and *John persuaded Mary to clean the barn*. However, a closer look suggests that the first sentence has little in common with the second. In fact, it is a better candidate for

strikes me as plausible that directly studying these features of language – how the grammar works and how the lexicon works – will in the long run be far more profitable than chasing after systematicity. Recursive syntax 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.

5 On the systematicity of thought

I turn now to a few remarks on the systematicity of thought. The term “thought” is used in many different ways in the literature. I will use it as a catchall term for mental representations, and the collective body of mental representations. Thoughts as I use the term are psychological states or processes that constitute what goes on inside of a person when she engages in an act of cogitation. (Although this is what I have in mind when I speak of thoughts, my arguments will hold – with appropriate adjustments made – against many other views of thought.) Just as with language, thought is frequently assumed to be systematic (e.g., Fodor and Pylyshyn 1988, Fodor and McLaughlin 1990, Evans 1982, Smolensky 1987). Moreover, the systematicity of thought is often considered to be a fairly unproblematic datum. Christopher Peacocke, for instance, writes approvingly that “the recombining of concepts to form new thoughts has been largely unquestioned in the published literature” (Peacocke 1992, p. 42). He also announces that “This recombining is about as general a phenomenon as one can hope to find in the realm of conceptual content” (Peacocke 1992, p. 41). But the most serious problems for the systematicity of language carry over straightforwardly to thought. Thus, the first problem for the systematicity of thought is to say what it is. If one thinks that the systematicity of thought is something along the lines of (S), then it will be important to say why this is the right form for thought, unlike (S’) or the form of systematicity for language, whatever that is. Secondly, one will need to specify the relevant mental categories, and to show that each mental primitive belongs to exactly one relevant category. In short, before the ball even gets rolling there are some formidable problems with developing a clear and justified characterization of the systematicity of thought. But without such an account we have nothing to discuss. I will therefore assume for the time being

being a systematic variant of *John gave Mary a ring* (Larson 1991). (I should note that there is only *some* truth in what Fodor and Pylyshyn say. If a being only had a small vocabulary of words available to it, its storehouse of basic sentences might indeed

that there is some acceptable characterization of the systematicity of thought which satisfies the mental correlates of (d1) – (d4) and which adequately expresses the combinatorial completeness of thought. In this section, I want to ask why one might think that thought is systematic. As far as I know, only two sorts of reasons have ever been offered to support the claim that it is. The first is that thought is similar to language in this respect, and the second is that it seems intuitively clear that thought is systematic. I will address these reasons in just a moment. Let me begin, however, by stressing that we should not be too quick to assume that a given empirical system (such as thought) is systematic.⁸

There are at least two reasons for doubting that a given empirical system is systematic. First of all, whatever precisely systematicity is, it appears to be rather rare in naturally occurring systems. For instance, the fact that one species of a genus can interact in some specified way with a feature of the environment does not imply that all species of that genus can do so too. I should mention, though, that not everyone agrees that systematicities are rare. Cummins claims that systematic systems appear to “abound in human psychology” (Cummins 1996, p. 604). Unfortunately, he gives no clear examples of this. His only putative example comes when he writes, “Consider, for example, the perception of objects in space: anyone who can see (imagine) a scene involving objects o_1 and o_2 can see (imagine) a scene in which their locations are switched” (Ibid.). But without specifying the psychologically relevant kinds of objects and locations, this claim is not obviously true. In this context, it seems natural to interpret “location” as location relative to other objects and background features. But although I can see (imagine) a small red box inside of a large glass sphere, I don’t think I can see (imagine) a large glass sphere inside a small red box. So without a clearer understanding of what systematicity is supposed to amount to in this case, it is not at all obvious that we have found a genuinely systematic system in human psychology. In fact, the only truly systematic systems of any complexity that I know of are the languages of logic and mathematics. But these systems didn’t arise naturally; they were *designed* to have systematic properties *precisely* in order to provide a more accurate medium for reasoning than natural language (and other systems of human psychology relevant to

be quite small and unsystematic-looking.)

⁸ At one point, Peacocke entertains the possibility that as a matter of stipulation, nothing would count as a thought (or an exercise of concepts) unless it exhibited the behavior characterized by systematicity (Peacocke 1992, p. 48). Such maneuvers do not eliminate the empirical problem, but only shift it over. In such a case, the relevant question would shift from “Is thought systematic?” to “Do human beings have thoughts?”

formal cognition). But the fact that these artificial systems are systematic is no evidence whatsoever for the systematicity of a naturally occurring system like thought.

The second reason to demand evidence that a given system is systematic is that systematicity is a strong mathematical property. For example, the view exemplified by (S) and other standard characterizations of systematicity appear to have the following structure.

- (gr) There exists a collection of sets of objects $C = \{G, C_1, \dots, C_n, \dots\}$, a designated element $G \in C$, and a set of functions $F = \{f_1, \dots, f_m, \dots\}$. The system is systematic iff for any (i) $f \in F$, (ii) k -tuple A_1, \dots, A_k , of elements from C , (iii) k -tuple a_1, \dots, a_k from $A_1 \times \dots \times A_k$, and any (iv) $x \in A_i$ (for any $1 \leq i \leq k$), it is true that $f(a_1 \dots a_k) \in G$ iff $f(a_1 \dots a_{i-1} x a_{i+1} \dots a_k) \in G$.

The abstract form of systematicity could be made even more general, but (gr) will suffice for now. Language is supposed to be systematic, because there is supposed to be some categories C_1, \dots, C_n and some syntactic operations F such that if the result of applying one of these functions to a k -tuple of expressions from the categories yields a grammatical sentence (i.e., an element of G), so will the result of applying the same function to any k -tuple that differs from the original at most in that it substitutes one element for another of the same category. But notice how easy it is for systems of various sorts to fail to be systematic. For instance, let G be the even numbers, and let C_1 be the prime numbers and C_2 be the composite numbers, and let f be the addition function. This system is not systematic: e.g., $f(3, 9) \in G$, but $f(3, 10) \notin G$. A bit more close to home, let C_1 be the set of words $\{Hesperus, Phosphorus\}$, and let $C_2 = \{Jones\}$ believes that ___ is the brightest star in the sky, and let $f(x,y) = 1$ iff the substitution of x into y yields a true sentence.

We've just seen that we really do need evidence that thought is systematic, so we should take the two sources of evidence seriously. Does the systematicity of language provide evidence for the systematicity of thought? Pretend for the moment that we can find characterizations of systematicity for language and thought which both satisfy desiderata (d1) – (d4). In such a case, the systematicity of language and thought will still be independent in the sense the neither entails the other. It's possible, for instance, for a being's language to have the sort of combinatorial completeness definitive of systematicity, even though the being's cognitive capacity lacks the corresponding property. In such a case, language as a means for expressing thought would then be like a tool that had more capabilities than it would ever be used for. On the other hand, a

being's language might fail to be systematic even though its capacity for thought is. In this case, the being would be capable of having thoughts to which it might not be able to give perfect expression. (Or perhaps it could give expression to such thoughts, but only by using roundabout circumlocutions that did not correspond in any obvious way with the structure of the thought itself.) However, even though the systematicity of language and thought are metaphysically independent from one another, the systematicity of thought is *evidentially* dependent on the systematicity of language. When it comes to giving reasons for thinking that thought is systematic, the purported systematicity of language looms very large. We see this in the large number of appeals to the (purported) systematicity of language made at the outset of discussions of the systematicity of cognitive processing in general (cf. the quotations at the beginning of this paper). This fact has been noted before. Cummins, for instance, characterizes our intuitions about the systematicity of thought as follows: "The intuitions in question are claims like the following:

(cu) Anyone who can think that *John loves Mary* can think that *Mary loves John*.

The systematicity of thought that this is supposed to illustrate is clearly derived from the systematicity of the language – in this case, from the fact that 'John loves Mary' is a permutation of 'Mary loves John'" (Cummins 1996, p. 596). The idea here is that the systematicity of thought is supported by a modus ponens argument: (i) language is systematic, and (ii) if language is systematic, then so is thought; hence, (iii) thought is systematic. Cummins goes on to argue against (ii), and we have seen that (i) is deeply problematic. So if, as Fodor says, "the only thing in the world that is as systematic as thought is language" (Fodor 1998, p. 26), the systematicity of thought is in serious trouble as well. The difficulties with the systematicity of language constitute a serious problem for the systematicity of thought, since the main empirical model of systematicity to which thought is supposed to be likened is natural language.

(Incidentally, we have surreptitiously raised another worry about premise (ii) of the argument just given. It could turn out, for example, that the appropriate formulations of systematicity for language and thought are radically different. Perhaps the systematicity of thought is adequately given by (S), but the systematicity of language is given by some vastly complex version of (S'). In such a case, it is hard to see how the systematicity of one of these

systems would provide any evidence whatsoever for the other. In fact, it is hard to see how the (S)-type systematicity for a given system would provide any evidence for that same system's having the other (S')-type form of systematicity. It is unclear what the two properties even have to do with one another.)

The second reason for asserting that thought is systematic is that it is intuitively obvious that it is. It's a little hard to see why our intuitions about systematicity are supposed to be evidence that thought is systematic, since so many of our other intuitions about cognitive structure have been shown to be way off base (consider, e.g., the large literature on the difference between our inferential practices and our reports of how we make various inferences). Nevertheless, I want to press the point a little, because I think it is all too often simply assumed that thought is systematic, and this sort of claim requires some kind of justification.

Language provides a strong reason to doubt the ability of introspective evidence to support claims of systematicity. After all, many people have found it intuitively obvious that language is systematic, and we have seen that our intuitions are not to be trusted here. Furthermore, we have seen that one of the reasons that simple versions of systematicity like (S) fail is that two words can fail to behave alike because of seemingly irrelevant syntactic or semantic or phonological reasons. If Peter helped Sue carry the piano upstairs, and you wanted to explain who Sue's partner was in this joint effort, you could not use (ex1a), even though it simply serves to create a kind of focus on Peter from the sentence in (ex1b):

- (ex1) a. *Peter is the guy who Sue and carried the piano upstairs together.
 b. Sue and Peter carried the piano upstairs together.
 c. Peter is the guy who Sue carried the piano upstairs with.

To make yourself clear, you might have to find some sort of circumlocution, like (ex1c). Thus, although you can get your idea across, the language does not allow you to do it in the simplest and most straightforward (and systematic??) way. A similar point may well hold for thought. Just as sentences are individuated by more than just their semantic content so too thoughts (qua mental representations) may be individuated by more than just their semantic content. Proponents of a language of thought are obviously committed to the possibility that the syntax of complex mental representations may have a structural properties that do not merely serve as a platform on which complex content (thought) is built out of concepts (e.g. Fodor 1975). If there

is a language of thought in any interesting sense, then it may well share some similarities with every other naturally occurring language-like system. Like every other non-artificial language, the syntax (or other properties) of the language of thought may simply prevent us from having certain thoughts, even when the content of that thought can be expressed by or at least closely approximated by other thoughts. If our language of thought is like English, for instance, then we may not be able to form the thought corresponding to (ex1a), even though we can form one with a similar content, in the form (ex1c). Of course, nobody knows whether syntactic properties play a significant role in the individuation of thoughts. Moreover, if there are syntactically relevant constraints on the individuation of thoughts, it is anybody's guess what those constraints are. Indeed, it is anybody's guess what mental representations will look like as we get clearer about the relevant phenomena and develop the tools to theorize about them. The point, though, is that we cannot simply assume that the properties relevant to the systematicity of thought will not be as varied and complicated as the properties that are relevant to the systematicity of language. To do so would be to adopt a very strong empirical assumption without empirical support, which is unacceptable.

6 Systematicity and the connectionism/classicism debate

The primary goal of this paper has been to critically consider the notion of systematicity as such, particularly as it pertains to language. However, a natural question arises concerning the consequences of our study of systematicity for the debates in which systematicity plays a part. Unfortunately, for any one of the debates a full answer to this question would be another paper in itself. So I will end by merely sketching how the present view of systematicity fits into one discussion, the debate between connectionists and classicists.

It will probably be important to get much clearer on what role systematicity plays in the connectionism/classicism debate. In addition to the questions about the form of systematicity – contrast (S) with (S'), for instance – there are also questions about what natural kinds of constituents (linguistic or mental) there are. We have seen that given the expressions of the language, if a theory makes available enough kinds of the right sort, then systematicity is trivially true. Reflection on this fact suggests that it is not fully clear what phenomenon connectionist devices are *unable* (in some robust sense of this term) to explain without

implementing a classical architecture. After all, when a connectionist system – or any kind of system, for that matter – learns language, it has to do several things, one of which is to figure out what the right linguistic (or mental) natural kinds of expressions there are, and how these kinds are organized. Our discussion of systematicity has shown, among other things, that two languages could share all the same words, and even a great many of the same sentences, and yet differ primarily because they organized the words into distinct natural kinds. In fact, there are no substantial logical restrictions whatsoever on what kinds a given set of words could be organized into. But this raises the following question: why shouldn't every connectionist device – and indeed every language-learning device whatsoever – be thought of as having learned a fully systematic language, but as having only induced an empirically incorrect organization of natural kinds of expressions? That is, for *any* set of grammaticality judgments that a language-learning system might make, there exists an organization of kinds of words under which the language is systematic. To see this point, assume for the moment that we have adopted some form of systematicity along the lines of (S), and that we have a connectionist device C which is being trained to learn English. Now suppose that C returns the following judgments:

- (cn1) 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. We can interpret the discrepancy between C's language and English solely in terms of a difference between what linguistic natural kinds serve to structure the language and to define systematicity for that language. Thus, we can consider C's language to be systematic, but the linguistic natural kinds of the language differ from those of English. In addition to learning a systematic language, it also appears that C provides a kind of implementational explanation of the systematicity of the language it has learned. C shows how to organize words into the right kinds (that is, the right kinds for that language, not for English), and thus accounts for the systematicity of that language. This interpretation of C's behavior and of the language it has learned suggests that C has only gotten the empirical details wrong. It learned a fully

systematic language, but didn't get the right linguistic kinds for English. But it is not a criticism of a research program like connectionism to say only that its models have so far not gotten all the empirical details right. Getting empirical details wrong is part of the normal course of the development of any empirical research program, and indeed is to be expected in one as new as connectionism. So connectionists do not appear to be challenged to explain systematicity, but to explain how English works. But if the challenge is just to learn English, then connectionists have only been challenged to do the very best cognitive science they can. And this is surely not a refutation of anything.

Have I just vindicated connectionism? Absolutely not. In fact, I have not altered the evidential status of connectionism at all. What I have done is shown that whatever problems there may be for connectionist models of language and cognition, systematicity is not one of them. There may be other ones, such as explaining how certain features or abilities could be innate, etc., but these are distinct from systematicity. At the end of section 4, I suggested that rather than trying to fix up all the problems with systematicity, we might do better in attending to the properties of natural language that give it whatever appearances of systematicity it has. These properties include its having constituents over which there is a recursive syntax, and the way in which new words enter the lexicon. Similarly, for at least three reasons, I suspect that when studying empirical models of linguistic and cognitive abilities, we might do better to ignore systematicity and study these other properties of language directly. First, if connectionist devices have problems modeling linguistic abilities, these problems are more likely to be exposed and understood by looking to the recursive structure of language and the way new words enter into the lexicon. After all, these are the mechanisms underlying systematicity. Second, any problems a theory has with systematicity will still be uncovered by focusing on these other properties. Finally, a deeper and more principled understanding of just what the empirical facts are will surely be useful for evaluating a given type of cognitive model. For example, many of the apparent failures of systematicity that we have looked at in this paper have been cases where individual words turn out to have slightly different features that have to do with the lexicon and how words are assigned to various linguistic categories, and not with the recursive structure of language. If so, then although the property of systematicity is frequently compared to productivity – which is just the compositional property of language – in fact, it sometimes has little to do with this recursive property. That is, many common instances of

systematicity appear to have a lot to do with the lexicon and little to do with recursive syntax. At the same time, other instances of systematicity have little to do with the lexicon and much to do with syntax; further instances are an admixture of several influences. The fact that systematicity and its failures is such a mixed bag would not be at all surprising if, as I have suggested, systematicity is not a real property of natural languages, in the way that a recursive syntax is.

7 Conclusion

The dominant theme of this paper has been that a theory of natural kinds of constituents is crucial to the notion of systematicity. Under a typical theory of systematicity, using standard grammatical categories, we saw that natural language is not systematic. An attempt to repair this problem by using more specialized grammatical categories exposed a deeper problem, namely that typical theories of systematicity make an empirically false assumption about the organization of linguistic natural kinds. The linguistic natural kinds that we have to work with do not form a partition of the words of the language, but form a complex overlapping structure. Given this system of linguistic natural kinds, it becomes unclear what it would be for natural language to be systematic. So before we can formulate an acceptable theory of systematicity, we need to answer both the question “What linguistic natural kinds are there?” and “How are these kinds relevant to systematicity?” The problems for the systematicity of language extend to the claims about the systematicity of thought as well, primarily because little or no substantial evidence for the systematicity of thought is offered besides thought’s putative similarity to language and our intuitive feelings that thought is systematic. Furthermore, given systematicity’s reliance on a theory of linguistic (or mental) natural kinds, it is unclear that connectionist devices really have a problem explaining systematicity. In fact, the explanation of systematicity appears to be nearly trivial, which suggests that the real problems lie elsewhere.

Despite all these difficulties with systematicity, it is still possible to end on a positive note. We have seen that there is a complex and challenging issue that directly concerns the nature of human language and the mind as well as the nature of systematicity, if there is such a thing. This issue is the question of what linguistic and mental natural kinds there are, and how they are organized. My guess is that a clear account of just how our expressions are organized

into kinds will also account for whatever intuitions about the systematicity of language or thought that we started with.

References

- Beghelli, Filippo and Stowell, Tim 1997, "Distributivity and Negation: The Syntax of *Each* and *Every*", in Anna Szabolcsi (ed.) *Ways of Scope Taking*, Dordrecht: Kluwer.
- Chomsky, Noam 2000, *New Horizons in the Study of Language and Mind*, Cambridge: CUP.
- Cummins, Robert 1996, "Systematicity", *The Journal of Philosophy*, 93, pp. 591—614.
- Evans, Gareth 1982, *The Varieties of Reference*, Oxford: OUP.
- Jerry A. Fodor 1975, *The Language of Thought*, Cambridge, MA: HUP.
- Fodor, Jerry A. 1998, *Concepts*, Oxford: OUP.
- Fodor, Jerry A. and Pylyshyn, Zenon 1988, "Connectionism and Cognitive Architecture: A Critical Analysis", *Cognition*, Vol. 28, pp. 3-71.
- Fodor, Jerry A. and Brian P. McLaughlin 1990, "Connectionism and the problem of systematicity: why Smolensky's solution doesn't work", *Cognition*, Vol. 35, pp. 183 – 204.
- Grimshaw, Jane 1990, *Argument Structure*, Cambridge: MIT Press.
- Jess Gropen, Steven Pinker, Michelle Hollander, and Richard Goldberg 1991, "Syntax and Semantics in the Acquisition of Locative Verbs" *Journal of Child Language*, Vol. 18, pp. 115-151.
- Hadley, Robert 1994, "Systematicity in connectionist language learning", *Mind and Language*, Vol. 9, pp. 247 – 272.
- Hadley, Robert 1997, "Cognition, systematicity, and nomic necessity", *Mind and Language*, Vol. 12, pp. 137 – 153.
- Hale, Kenneth and Keyser, Samuel Jay 1993, "On Argument Structure and the Lexical Expression of Syntactic Relations", in K. Hale and S. Keyser (eds.), *The View from Building 20*, Cambridge, MA: MIT Press, pp. 53-109.
- Harris, Zelig 1951, *Methods in Structural Linguistics*, Chicago: University of Chicago Press.
- Higginbotham, James 1987, "The autonomy of syntax and semantics", in J. Garfield (ed.) *Modularity in knowledge representation and natural-language understanding*, Cambridge, MA: MIT Press.
- James Higginbotham 1989, "Elucidations of Meaning", *Linguistics and Philosophy* 12, pp. 465 - 517.
- Johnson, Kent and Ernie Lepore forthcoming, Does Syntax Reveal Semantics? A Case Study of Complex Demonstratives, in J. Tomberlin (ed.), *Philosophical Perspectives*.
- Larson, Richard K 1991, "Promise and the theory of control", *Linguistic Inquiry*, Vol. 22, no. 1, pp. 103 – 139.
- Matthews, Robert 1997, "Can connectionists explain systematicity?", *Mind and Language*, Vol. 12, no. 2, pp. 154 – 77.
- Niklasson, Lars and Tim van Gelder 1994, "On being systematically connectionist", *Mind and Language*, Vol. 9, pp. 288 – 302.
- Parsons, Terence 1990, *Events in the Semantics of English*, Cambridge, MA: MIT Press.
- Peacocke, Christopher 1992, *A study of concepts*. Cambridge, MA: MIT Press.

- Pinker, Steven 1989, *Learnability and Cognition*, Cambridge, MA: MIT Press.
- Ross, J. R. 1967, *Constraints on Variables in Syntax*, unpublished doctoral dissertation, MIT, Cambridge, MA.
- Schiffer, Stephen forthcoming, "Meanings", in J. Campbell, M. O'Rourke and D. Shier (eds.), *Essays on Meaning and Truth*, Seven Bridges Press.
- Smolensky, Paul 1987, "The constituent structure of mental states: A reply to Fodor and Pylyshyn", *Southern Journal of Philosophy*, Vol. 26, pp. 137 – 160.
- Smolensky, Paul 1988, "On the proper treatment of connectionism", *Behavioral and Brain Sciences*, Vol. 11, pp. 1 – 23.
- Smolensky, Paul 1995, Constituent structure and explanation in an integrated Connectionist/Symbolic cognitive architecture, in C. Macdonald and G. Macdonald (eds.) *Connectionism: Debates on Psychological Explanation*, Oxford: Blackwell, pp. 223 – 290.
- Travis, Charles 1994, "On constraints of generality", *Proceedings of the Aristotelian Society*, Vol. 94, pp. 165 – 188.
- Travis, Lisa 2000, "Event structure in syntax", in C. Tenny and J. Pustejovsky (eds.) *Events as grammatical objects*, Stanford: CSLI Publications, pp. 145-85.