

What Systematicity Isn't: Reply to Davis

Robert Cummins

Jim Blackmon

David Byrd

Alexa Lee

Martin Roth

Journal of Philosophical Research. 30:405-408, 2005

The alleged systematicity of thought is widely believed to be strong evidence for the language of thought hypothesis (hereafter, LOT).¹ Cummins (1996) and Cummins et. al. (2001) have suggested that the systematicity of thought is not the obvious datum it is sometimes taken to be, and that, once this is appreciated, the systematicity argument as it is usually formulated appears to be question-begging. Wayne Davis (this journal) thinks we have this wrong, and that the systematicity of thought has an obvious and non-question-begging formulation. It is this:

(2) Anyone who can think Rab can think Rba.

Or this:

(3) Anyone who can think the thought expressed by a sentence of the form Rab can think the thought expressed by the corresponding sentence of the form Rba.

As Davis points out, (2) has to be taken as a schema. If we read it with 'Rab' as a constituent, the sentence is false or meaningless. But if we take it as a schema, what is 'Rab' schematic for, and how are we to understand the contrast between 'Rab' and 'Rba'? We think it is clear that the only way to understand this that is not question begging is to take it in the sense of (3), for it is unproblematic in this context to understand the contrast between 'Rab' and 'Rba' in (3) as involving switching the formal constituents 'a' and 'b.' But if we try to understand the contrast between 'Rab' and 'Rba' in (2) as involving switching thought constituents the question will have been begged. We conclude that (2) is best taken as shorthand for (3).

Davis objects to Cummins' complaint that, once we see that the systematicity of thought (ST) is formulated as (3), we can see that ST is derived from the systematicity of language understanding. Davis argues that the systematicity of thought can be used in the *explanation* of the systematicity of language understanding, even though we may *infer* the systematicity of thought from the systematicity of language understanding. We agree with Davis on this latter point. The inference in question might go something like this: Anyone who can understand a sentence of the form Rab can understand a sentence of the form Rba; anyone who understands a sentence must be able to think the thought it expresses; therefore, anyone who can think the thought expressed by a sentence of the form Rab can think the thought expressed by a sentence of the form Rba. Of course, ST is supposed to follow from the latter schematic assumption. Davis claims that all we need to assume is that instances of schema (3) be true. But it is clearly a matter for empirical investigation whether we have good reason to assume that they actually *are*. You would think that you don't find people who can see but believe they cannot (blindsight), or people who cannot see but believe they can (blindness denial) But you do. You would think that you don't find people who can write but not read. But you do. You would think you wouldn't find people who could recognize hats but not faces. But you do. And so on. In advance of

¹ LOT's central contemporary advocate is Fodor (1975 and many subsequent publications), but the idea goes back at least to Plato (Theaetetus 189e-190a Sophist 263e-264a), who appears to have believed that thinking is talking to oneself.

empirical demonstration, the jury is still out on the evidential status of schema (3), and, hence, on the very existence of the data that LOT is invoked to explain in the first place. It is worth emphasizing that (3) concerns *thought*, as opposed to, say language understanding. It is not, as advocates of the systematicity argument suggest, open to casual observation what thoughts one can and cannot have in the way it *is* open to casual observation that someone does or does not understand ‘Dogs hate cats.’

Furthermore, even if instances of schema (3) were demonstrated to be true, that would be nowhere near sufficient to ground ST. Consider:

(4) Anyone who can think the thought that dogs hate cats can think the thought that dogs hate skunks.

In the absence of sophisticated theory or careful empirical investigation, the evidence for (4) is on all fours with the evidence for (3). But clearly the truth of (4) has nothing to do with ST. To really support ST, you actually need to show that there is something special about thinking systematic variants. For example, you would need to demonstrate something like:

(5) Anyone who can think the thought expressed by sentence S can think the thoughts expressed by its systematic variants, and it is at least possible to think the thoughts expressed by S and its systematic variants yet not think the thoughts expressed by any other sentence of L.

We don’t know whether (5) is true or not, but it is certainly not the (contingent) truism that Davis takes (3) appears to be.² It certainly could not be established by casual observation of the folk, for example. So it seems that even if the data needed to support (3) *were* available, you still won’t get ST without the addition of sophisticated theory and/or evidence to support (5). Of course, until ST can be demonstrated, there is no need to explain it. And, hence, acceptance of LOT on the basis that it provides a “natural” explanation of this phenomenon is more than a little premature. But even if it weren’t, surely the history of science has by now taught us that the *naturalness* of an explanation in no way speaks to its *adequacy*.

Notice, too, that it will do no good to object that one could understand ‘Cats hate dogs,’ and hence ‘Dogs hate cats,’ even if you didn’t understand ‘Dogs hate skunks.’ While this is surely right, it is *not* what the doctor ordered. The LOT explanation of systematicity turns on the idea that understanders of a sentence S are likely to be understanders of systematic variants of S, because the same resources are required. But just what is it that is supposed to be missing in the case of someone who understands ‘Cats hate dogs’ and its systematic variants, but doesn’t understand ‘Dogs hate skunks’? The failure *could* be the result of lacking the thought that dogs hate skunks, but it is surely much more likely that your find-the-meaning-of-the-current-linguistic-input routine won’t find the thought that dogs hate skunks in response to hearing the sentence ‘Dogs hate skunks.’ In the latter case, although you might not be able to understand the sentence ‘Dogs hate skunks,’ you will be capable of thinking the thought that dogs hate skunks. And if you can think the thought that dogs hate skunks even though you cannot understand the sentence, then your failure to understand the sentence doesn’t bear on the systematicity of thought at all. For, to repeat, hands-off observation of the folk certainly will not show that

² Being a conceptual atomist, Fodor is committed to the possibility mention in (5). Being a concept/LOT nativist, he doesn’t think it actually happens (or anyway, not much). But if LOT is innate, this implies that we can all token all the sentences of LOT. One of those sentences is, presumably, |dogs hate skunks|. So, it simply follows from the innateness of LOT that anyone who can think that dogs hate cats can think that dogs hate skunks, or any other humanly possible thought.

someone who cannot understand ‘Dogs hate skunks’ cannot think the thought that dogs hate skunks, since that will be true of everyone who doesn’t know what ‘skunk’ means, but can still think about skunks.³

So: does the systematicity argument beg the question? We suspect that the systematicity argument looks most appealing to those who find little difference between being able to think about skunks, and being able, in some sense or other, to refer to skunks. If you are tempted to read ‘about’ in ‘thinking about skunks’ so that ‘thinking about skunks’ comes to ‘having a thought that *refers* to skunks’, you are already as near to accepting LOT as makes no difference, and the idea that there are systematic variants in thought that correspond to systematic variants in language will look like a no-brainer. There is no doubt that something like this is entrenched in “folk psychology”. If you ask yourself why you are inclined to believe that someone might be able to think that dogs hate cats and that cats hate dogs but not that dogs hate skunks, the answer that pops up is that someone might have the concepts DOG, CAT, and HATE but lack the concept SKUNK. But this *is* question begging, since the alleged datum is being inferred from the very theory it is supposed to support. Once we distinguish (3) from (5), and appreciate that it is (5), not (3), that the argument requires, the illusion that the systematicity argument proceeds from a contingent truism simply vanishes.

Works Cited

Cummins, R. (1996) “Systematicity,” *Journal of Philosophy*, 93: 591-614.

Cummins, R., J. Blackmon, D. Byrd, P. Poirier, M. Roth, and G. Schwarz (2001) “Systematicity and the Cognition of Structured Domains,” *Journal of Philosophy*, 98:167-185.

Davis, Wayne (this journal) “On Begging the Systematicity Question.”

Fodor, J. (1975), *The Language of Thought*, Thomas Y. Crowell.

³ We suspect that the systematicity argument looks most appealing to those who find little difference between being able to think about skunks, and being able, in some sense or other, to refer to skunks. If you are tempted to read ‘about’ in ‘thinking about skunks’ so that ‘thinking about skunks’ comes to ‘having a thought that *refers* to skunks’, you are already as near to accepting LOT as makes no difference, and the idea that there are systematic variants in thought that correspond to systematic variants in language will look like a no-brainer.