

EXTERNALIST THOUGHTS AND THE SCOPE OF LINGUISTICS

Kent Johnson

Abstract

A common assumption in metaphysics and the philosophy of language is that the general structure of language displays the general metaphysical structure of the things we talk about. But expressions can easily be imperfect representations of what they are about. After clarifying this general point, I make a case study of a recent attempt to semantically analyze the nature of knowledge-how. This attempt fails because there appears to be no plausible bridge from the linguistic structure of knowledge-how reports to knowledge-how itself. I then gesture at some other places where the connection between linguistics and metaphysics is commonly, but illegitimately, assumed.

1 An Alleged Connection between Semantics and Metaphysics

What is the relation between philosophy and linguistics? Many philosophers endorse, tacitly or explicitly, the view that linguistic theories can support philosophical theses in various substantial ways. In particular, there's a long tradition in philosophy of supposing that questions about the metaphysical nature of certain phenomena can be uncovered by investigating the semantics of the expressions we use to talk about them. This has been one of the central underlying components of the "linguistic turn" and of various philosophical programs of "semantic analysis". In short, the thought has been that for many philosophical topics, if you can discern the semantics, you'll uncover the metaphysics. The assumption that linguistics can supply evidence for metaphysical theses has been endorsed by many, including those who are well aware of the finer details of both philosophy and of contemporary linguistic theory. For instance, James Higginbotham, who has made numerous substantial contributions to both fields, has endorsed such an assumption in numerous places (e.g., Higginbotham 1989, 1992, 2001, 2004). Higginbotham suggests that linguistic theory can be a substantial aid in "the clarification of the nature of our thoughts, what we actually express when we understand one another" (Higginbotham 2004, 575). Initially, this might seem like a primarily psychological claim. However, Higginbotham then goes on to individuate (what he calls) "thoughts" in a

manner that makes them of central philosophical importance. In particular, he suggests that thoughts should be individuated, at least in part, by things in the world that thoughts and their parts refer to. For instance, he writes:

Assume, what is common enough although open to question, that this clarification calls first of all for the exposition of the truth conditions of sentences, as they occur as parts of total languages, and within the contexts of their potential utterance; and assume also that any correct account of what we are inclined to assert must, over a wide domain, make us pretty much right about the way things are. Then the truth conditions of much of what we believe must be such as to be actually met; and this implies that *what turns up in the metaphysics of semantic investigation cannot be passed off as a mere manner of speaking, but constitutes our best conception of the way the world is...* . To say this much is not to commit semantics to an all-out realism with respect to the elements invoked. It does, however, imply that the question of realism must be taken seriously, including familiar questions of possible reduction or relativization of the objects involved, as for instance whether individual events can be reduced to regions of space and time, or whether possible worlds can be modeled as maximally consistent sets of propositions (Higginbotham 2004, 575; italics added).

Here we see a decidedly “externalist” view of linguistics, in the sense that that the subject matter of linguistics is not limited to the states and processes in a speakers’ head. (Higginbotham has argued for such an interpretation of linguistics in numerous other places, e.g. his 1989.) As Higginbotham’s discussion makes clear, his “thoughts” are either identical or closely analogous to what other philosophers would call “propositions”, where the latter are interpreted in a similarly externalist way.

(It’s worth observing that externalism in linguistics contrasts starkly with the more common “internalist” view, whose roots lie in Chomsky’s work on syntax. According to internalism, semantic theory should be construed as a system of mental representations of the world. Thus, talk of an expression’s “reference” is really just shorthand for some kind of conceptual/intentional state or capacity of the speaker that is associated (in some linguistically relevant way) with the expression. Thus, internalism attaches much less importance to the question of realism that Higginbotham alludes to in the passage above.)

It would be great to see the sort of unification between philosophy and linguistics that Higginbotham (and numerous others; cf. below) propose. It is, however, one thing to suggest such a coordination between the disciplines, and another thing for such a view to command authority. In the next section (§2), I explain my skepticism about this view by presenting a picture of linguistics

and making some distinctions. These distinctions, I believe, are what really undermine Higginbotham’s two suggestions. However, in order to bring my point home, I then (§3) make a detailed case study of a recent attempt to establish a philosophical thesis by looking to current linguistic theory. We’ll see there that the distinctions drawn in §2 provide the crucial conceptual machinery for organizing an analysis of what goes wrong with this analysis. Next in §4, I gesture at a few more places where I think similar misapplications of linguistic data have appeared in philosophical theorizing. I conclude in §5.

2 Languages as representational systems

Regardless of how one interprets linguistics, it is hard to deny a claim like (Rep):

(Rep) Languages are representational systems, and so may produce imperfect representations.

(Rep) tells us that linguistics theorizes about representational systems. Whether our sentences “express thoughts” in the internalist sense of Chomsky (e.g., 1980, 230) or depict the world, they are representations of some sort.¹ As such, these representations may be imperfect. That representations are frequently imperfect is a ubiquitous phenomenon. For instance, authors of maps and diagrams often deliberately depict relevant objects (e.g. roads) as proportionally larger than objects less likely to be relevant (e.g. surrounding countryside) (cf. Matthews 1994 for discussion). Such representations also frequently omit much structure in what is being represented; e.g., a road that heads due north with a small bend in its middle may appear on a map as a straight line. In the linguistic case, consider that the semantic representation associated with the noun *water* is plausibly of a continuous, nonparticulate substance, even though water may be composed of discrete molecules. (In any case, it’s at least possible that language and the world would differ like this, which is all that really matters here.) Many further examples can be taken from the history of philosophy. For instance, a number of philosophers have echoed the sentiments of Thomas Reid, who observes that there are ‘phrases which have a distinct meaning’ other than what their overt form suggests, and that ‘there may be something

1 I take it that whatever else “expression” is, it is a form of representation.

in the structure of them that disagrees with the analogy of grammar or with the principles of philosophy... . Thus, we speak of pain as if pain was something distinct from the feeling of it. We speak of pain coming and going, and removing from one place to another' (Reid 1785, 167–8). Although we may say *John has a pain*, it doesn't follow that our language represents the situation so accurately, as involving two things, John and a pain. (One's metaphysics may lead one to disagree here, but my point is only that it's not apriori true that linguistic structure accurately mirrors the world.) Similarly, Chomsky has noted that the proper semantic representation of *The value of my watch is increasing* may be that of a certain kind of object moving on a unidimensional scale, but of course this has nothing to do with what it is for something to rise in value. (These sorts of representations are quite natural from a processing perspective, since they probably recruit our abilities for spatial cognition, which appear to be quite basic to human representation (e.g., Jackendoff 2002, Burgess et al. 1999).) Actually, for present purposes, we don't even need the claim that our sentences are representations, since if they're not, then they still don't provide perfect representations. All we really need to make the above points is that our sentences' representations can be imperfect. (Indeed, in the case of human psychology, it may be important that some of our representations are imperfect. The literature on learning and memory suggests that learning complex facts or skills often crucially involves learning to ignore irrelevant information; e.g. learning to count involves ignoring whether the objects being counted are apples or oranges.)

(Rep) acts as a degree of freedom lying between the structure an expression attributes to what it represents, and the actual structure of what it represents. In order to respect this degree of freedom, it will be helpful to distinguish three types of structure relevant to theories about language. First, *metaphysical* (or *empirical*) *structure* is the structure present in the external thing that an expression denotes. In the typical case, this sort of structure is uncovered by scientific or metaphysical investigations into some extra-linguistic phenomenon. E.g., our best chemical theories tell us that a given quantity of water is (primarily) composed of a discrete number of molecules of H₂O. The fact that water has this structure is independent of any linguistic properties of the English word *water*, which may well treat it semantically as a continuous, non-particulate substance. Thus, there may be little or no interesting relation between an object's metaphysical structure and the structure of the expression describing it.

Second, an expression's *syntactic structure* concerns how the lexical elements of a sentence are actually organized into constituents that form the target ex-

pression, regardless of how untrained users of the language tend to think about this structure. It's well known that much syntactic structure isn't accessible to conscious reflection. Some of this tacit structure is also relevant to semantic interpretation, which is the norm for every other cognitive ability that psychologists have investigated. It takes much experimental work, for instance, to see the kinds of structural features that are relevant to our interpretation of a visual scene (e.g. Hoffman 1998).

Third, *semantic structure* is the structure of the meaning (in the internalist's sense) of our expressions. An expression's semantic structure may not be identical to the structure of what it represents (e.g., the quote by Reid above). There is also no a priori guarantee that the semantic structure of an expression will mirror its syntactic structure. The relation between syntax and semantics remains an active area of research (cf. discussion and citations below). Often, though, syntactic and semantic structure go hand in hand, so I'll frequently use the phrase *linguistic structure* to lump them together.

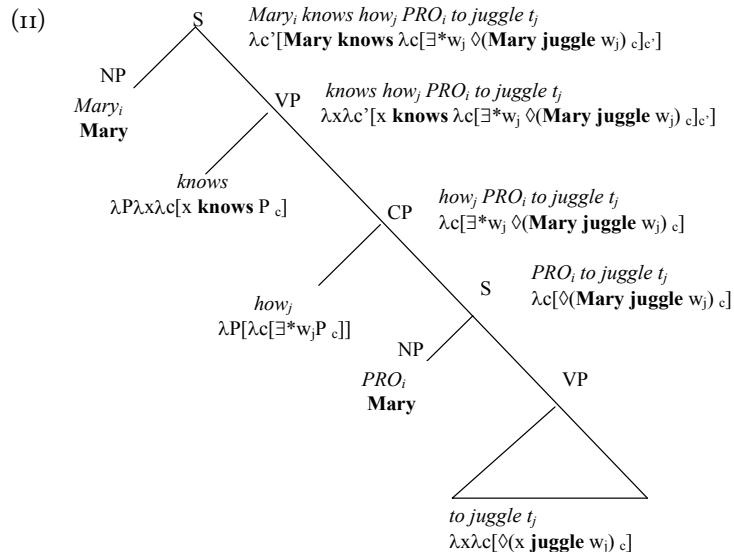
According to the strategy proposed by Higginbotham (and others), we can investigate the nature of the metaphysical structure of parts of the world by investigating relevant parts of the semantic structure of the language we use to talk about them. Using the terminology just introduced, this suggests that there should be some kind of significant connection between (certain aspects of) metaphysical and linguistic structure. However, a moment's thought shows that none of these kinds of structure logically entails anything about the other. Indeed, from our present perspective, it is hard to see how some of these forms of structure would have very much to do with one another at all. So inferences from one kind of structure to another require some defense. In particular, in order to have any hope of uncovering the underlying nature of some metaphysical phenomenon, we must forge some theoretical links between some aspects of linguistic structure and some aspects of metaphysical structure. The importance of arguing for such links should not be underestimated. In what follows, I want to illustrate the importance and difficulty of finding such links. I'll do this by examining a recent attempt to use linguistic structure to uncover the true nature of a philosophically interesting phenomenon, that of knowledge-how. The attempt fails precisely because the relations between linguistic and extralinguistic structure go simply ignored. To the extent that we cannot take such relations for granted, the general attempt to relate linguistics and philosophy remains incomplete.

3 Language as revealing our “externalist” thoughts

As Ryle (1946) and many others have observed, there appears to be a real difference between knowledge that something is the case (e.g., that snow is white) and knowing how to do something (e.g., how to juggle). The latter form of knowledge is distinctive, not least because it appears that one could know how to do something without being able to articulate, recognize, affirm, etc. any propositions that describe the content of that knowledge. People like Ryle have taken such data seriously, as marking a genuine theoretical distinction. Recently, however, the distinction between knowledge-how and knowledge-that has been attacked. Stanley and Williamson (2001) – hereafter SW – ‘contest the thesis that there is a fundamental distinction between knowledge-how and knowledge-that. Knowledge-how is simply a species of knowledge-that’ (SW 2001, 411). Their argument is based around the claim that our best semantic theories represent ascriptions of knowledge-how as ascriptions of knowledge of a proposition. E.g., they argue that the relevant interpretation of (9) has a semantic form similar to (10).

(9) Mary knows how to juggle.

(10) There is a way w such that Mary knows that w is a way that she can juggle.



Abstracting from many details, we can think of the derivation of (10) from the syntactic structure of (9) as being roughly as given in (11).

(11) shows how the meaning of (9) can be derived from its parts.² Let's assume that the semantics used in (11) is on the right track, interpreting ' \exists^* ' as the kind of quantifier appropriate to SW's theory. So in (11), \exists^* would be interpreted as having a wider scope than *know*. SW are silent about how to implement their theory, but for present purposes, I will just assume that these details theory can be satisfactorily worked out.³ Let us also assume that the trace t_j is a variable

2 (11) is a phrase marker that describes the relevant syntactic constituents that form (9). Following standard linguistic practice, (11) explicitly represents the unpronounced (and unwritten) subject of the lower clause (*PRO*). Also, since Mary must be the subject of the lower clause, *PRO* is represented as coindexed with *Mary*. Still following standard practice, *how* is assumed to originally appear in an adverbial position inside the lowest verb phrase. It then 'moves' up into a position near the top of the entire complement phrase. A 'trace' of *how* remains in the original position from which it moved. To indicate that the trace is of *how*, the two are coindexed. The general structure of (11) represents the sentence *Mary_i knows how_j PRO_i to ride a bike t_j* as being built out of a subject (*Mary_i*) and a complex verb phrase (*knows how_j PRO_i to ride a bike t_j*). This verb phrase is itself built out of a verb (*knows*) and a complement phrase (*how_j PRO_i to ride a bike t_j*), which is itself constructed of a moved *wh*-phrase (*how_j*), and a sentential clause (*PRO_i to ride a bike t_j*). The subject of this clause is *PRO_i*, and the verb phrase is an infinitive *to ride a bike t_j*. For discussion of the empirical motivation of the various devices used here, a standard introductory textbook in syntax should be consulted (e.g., Haegemann 1994). Alongside each constituent represented in (11) is the semantic value of the constituent as a standard semantic theory would represent it. I assume that the lowest VP receives as its interpretation a function from entities x to (incomplete) functions from possible circumstances c to truth values, such that the latter function yields the value true iff it's possible in c for x to ride a bike in way w_j (a free variable, to be bound later). Composition proceeds at each node by the application of an argument to a function. (In the CP node, I have additionally assumed some compositional mechanism for augmenting the proposition P . This can be done in many ways, but since such details are not relevant here, I omit them. For more details on the semantics, cf. a standard textbook, e.g. Chierchia and McConnell-Ginet 1990.)

3 The nature of the existential quantifier SW appeal to in (10) is curious. As SW argue, the quantifier must take wider scope than the verb that governs it, since Mary may not know how to juggle if she only knows that there's a way to juggle. However, the quantifier also seems constrained not to take any wider scope than that. In the normal reading of the sentences in (i), for instance, the alleged quantification over ways is *unable* to scope over other propositional attitude environments, negation, and quantified subject NPs:

- (i) a. John thinks that Mary knows how to juggle.
- b. Mary doesn't know how to juggle.
- c. Each girl knows how to juggle.

(E.g., (ib) does not simply mean that there's a way to juggle that Mary doesn't know; it means that it's not the case that there's a way that she knows.) As the sentence in (ii) show, this is not

bound by \exists^* in a syntactic position to indicate a way to perform the action denoted by the verb *juggle*. Together, these assumptions say that the meaning of the top S node in (11) is the proposition that there is some way w_j such that Mary knows the proposition (i.e., the function from possible worlds to truth values) that it is possible for her to juggle in way w_j . This precisely represents the meaning given in (10).

For now, let's assume that the true linguistic structure of (9) is exactly as (11) describes. Given these assumptions, we can examine the rest of SW's argument. SW argue that as (11) shows, *how to juggle* is a clausal direct object of *knows*, just like *that Sam laughed* in (12):

(12) Mary knows that Sam laughed.

Similarly, the semantics of these two clausal constituents are identical in all relevant respects. So since (12) ascribes propositional knowledge to Mary, (9) must too. Hence, knowledge-how is really a form of knowledge-that. In addition to the argument just sketched, SW also offer two further pieces of evidence for their thesis. First, they correctly observe some syntactic similarities between

the behavior we find with other quantifiers. *there*-constructions force a narrow scope reading, and the other sentences allow the relevant quantifier to take wider scope than SW's quantifier would allow:

- (ii) a. Everybody knows there is a man in the room with red hair.
 b. Each boy knows a man in the room with red hair (is singing).
 c. Mary doesn't know a man in the room with red hair.
 d. John thinks that Mary knows a man in the room with red hair.

(iia) lacks a reading that requires everybody to know the same man, and (iib-d) have readings in which *a man in the room with red hair* takes widest scope. Thus, on SW's proposal, *how* is semantically unlike other quantifiers. As I described in an earlier footnote, it's standard in linguistics to try to reduce the purely functional aspect of theories by looking for underlying similarities. Thus, SW's proposal seems rather ad hoc, treating the relevant kind of constructions as a *sui generis* type of mandatorily intermediate-scope quantification.

Finally, it's worth noting that other *wh*-adjuncts may not behave like *how*, either. I *think* that the sentences in (iii) allow *where* to take narrowest scope as well as widest scope. (The relevant test here should be performed with *some* taking narrowest scope – it is there to ensure that multiple *wheres* are possible, thus creating semantic differences in scopal ordering)

- (iii) a. Each boy knows where some cars were stolen.
 b. Mary doesn't know where some cars were stolen
 c. John thinks that Mary knows where some cars were stolen.

E.g., in (iiia) there might be a particular lot A that the speaker has in mind, and he knows that some of the boys know that a few Fords were stolen from lot A, and that the other boys know that several Toyotas were stolen from there too.

knowledge-how ascriptions and other sorts of indirect questions, such as those in (13):

- (13) a. Mary knows [_{CP} where [_S to go]].
 b. Mary knows [_{CP} when [_S to stop]].
 c. Mary knows [_{CP} what [_S to eat]].

(13a–c) all have syntactically clausal direct objects, just like (9). Let's also assume that they all have a similar semantics. But, SW argue, the sentences in (13) are clearly ascriptions of propositional knowledge, i.e., knowledge-that. So (9) must ascribe propositional knowledge, too. Second, some interpretations of knowledge-how sentences like (9) may ascribe propositional knowledge to the subject. E.g, among its many meanings, *Mary knows how to juggle* can mean either of the things in (14):

- (14) a. Mary knows how she ought to juggle
 b. Mary knows how one ought to juggle.

Both these interpretations, SW argue, ascribe propositional knowledge to Mary (SW 2001, 424ff.). So all interpretations of (9) must ascribe propositional knowledge to Mary. Thus, (13)–(14) expose substantial similarities between ascriptions of knowledge-how and clear ascriptions of knowledge-that.

SW conclude:

If these standard accounts of the syntax and semantics of embedded questions are correct, then ascriptions of knowledge-how simply ascribe knowledge-that to their subjects... Our view of ascriptions of knowledge-how is very straightforward. It is just that the standard linguistic account of the syntax and semantics of embedded questions is correct... Our view of ascriptions of knowledge-how is the analysis reached on full consideration of these constructions by theorists unencumbered by relevant philosophical prejudices (SW 2001, p. 431)

Unfortunately, this conclusion says nothing about knowledge-how. SW have *at most* defended a conclusion about the linguistic structure of ascriptions of knowledge-how (and we'll see shortly that even this conclusion is a little premature). But their thesis that 'knowledge-how is simply a species of knowledge-that' concerns the metaphysical structure of knowing how.⁴ SW hold that

4 It might seem more natural to refer to knowledge-how as an epistemic phenomenon. However, since what's relevant about knowledge-how (as opposed to linguistic ascriptions of

knowing how to do something is tantamount to standing in a certain relation to a proposition. So they're not primarily defending a claim about the structure of the language used to report the existence of such states. Rather their thesis concerns what knowledge-how *itself* is. But their evidence only concerns the structure of the language. Since their semantic theory represents knowledge-how as propositional, they maintain that knowledge-how must really be propositional. In other words, *at best* SW infer the metaphysical structure of the nature of knowledge-how from the linguistic structure of ascriptions of knowledge-how. But we've seen that an expression's linguistic structure needn't always mirror the metaphysical structure of its denotation. So we can't infer the structure of knowledge-how from the structure of knowledge-how ascriptions.⁵

SW's argument is incomplete as it stands. Could there nonetheless be some way to defend it? Could some general principles of the philosophy of science help here? Maybe, e.g., knowledge-how should be thought of as a species of knowledge-that simply because such a hypothesis presents a more unified, coherent, simpler story than the alternatives. That is, if our best linguistic theories assign knowledge-how reports a type of structure similar to sentences expressing knowledge-that, then, barring any further contravening considerations, perhaps we should conclude that the structure of the linguistic theory correctly mirrors the world. So maybe knowledge-how is a type of knowledge-that. This strategy *appears* to be how SW would choose to complete their argument. But it is compelling only if there aren't any significant reasons for distinguishing knowledge-how from knowledge-that. But there are.

In general, we deny that a type X is a species of another type Y whenever Xs and Ys are simply too different (physically or functionally) to naturally group Xs as Ys. Rocks are not a species of bird because they're too physically dissimilar. Similarly, for normal humans, singing is not a species of running because the mechanisms underlying these two activities are different. So we might ask, are knowledge-how and knowledge-that realized with sufficiently similar physical or functional mechanisms to make one a species of the other? As Ryle (1946) and many others since him have made clear, there is very little functional

knowledge-how) is that it is a non-linguistic phenomenon that occurs 'in the world' independently of language, I will continue to refer to it as a metaphysical phenomenon.

5 The empirical nature of knowledge-how – as opposed to the linguistic nature of ascriptions of knowledge-how – is the issue of traditional interest, as is very clear from the authors SW take pains to engage. (At the same time, though, SW are right to correct the researchers who mistakenly characterize and discuss knowledge-how by reference to the allegedly non-clausal syntactic structure of ascriptions of knowledge-how.)

similarity between knowledge-how and knowledge-that. In any case, an argument would need to be provided that knowledge-how is functionally speaking a subspecies of knowledge-that. (Of course, if there were such an argument, then the linguistic evidence in support of such a thesis would be largely otiose anyways!) A closer look into the more fine-grained functional and physiological specification of knowledge-how (in ordinary humans, at least) provides further evidence against treating it as a species of knowledge-that. In the neuroscientific literature, knowledge-that corresponds to 'declarative' knowledge or memory, while knowledge-how corresponds to 'procedural' knowledge or memory (e.g., Squire 1992a, b, Squire et al 1993, Ullman et al. 1997, Ullman 2001). These discussions often fall under the heading of 'memory', although the phenomena in question trivially include paradigmatic instances of knowledge-how and knowledge-that. Indeed, procedural memory is often referred to as knowledge-how (e.g., Cohen and Squire 1980). These two ways of storing, retrieving, and using knowledge appear to be instantiated in different parts of the brain. For instance, Ullman et al. discuss a division of the brain into:

two major kinds of memory systems... One is a *declarative memory* system underlying the learning and storage of information about facts and events. It is subserved by a medial temporal circuit connected largely with neocortical areas in the temporal and parietal lobes, with the medial temporal components consolidating memories that are eventually stored in neocortex... The other is a *procedural memory* system for the learning and processing of motor, perceptual, and cognitive skills. It is subserved by basal ganglia circuits connected largely with frontal cortex (Ullman et al. 1997, 267).

Similarly, Larry Squire has argued on neuroscientific grounds for a taxonomy of different kinds of memory which clearly distinguishes knowledge-how from knowledge-that.⁶ In short, the neurological evidence suggests that our possession and realization of certain skills (including knowing how to juggle) are physically instantiated in strikingly different brain regions than those areas that underlie our declarative memories of facts and events. Although memory is of course different from knowledge in some respects, the neuroscientific evidence suggests that knowledge-how is at least quite physically different from knowledge-that. Furthermore, these studies also show that the two forms of

6 In Squire's famous taxonomy (Squire 1992a, 205), memory is broken into two subsystems: declarative and nondeclarative memory. The former encompasses semantic memory (memory of facts) and episodic memory (memory of events). The latter encompasses skills (motor, perceptual, cognitive), priming (semantic and perceptual), dispositions, and nonassociative memory (habituation and sensitization).

memory are quite different in terms of their fine-grained functional properties relevant to cognitive processing. Thus, from the perspective of philosophy, psychology, and neuroscience, knowledge-how does not appear to be a species of knowledge-that.⁷

As a brief aside, it may be worth commenting briefly on the thoroughly ‘naturalistic’ stance I adopt in this paper. It should be clear that I’m only interested in a naturalistic game, and I’m happy to limit the scope of my arguments accordingly. A bit more narrowly, though, it’s hard to see how a non-naturalistic conception of knowledge-how could be justified by appealing to the rather fine-grained empirical details of theoretical syntax. E.g., presumably our “intuitive concept” of knowledge-how – assuming that there is such a thing – does not change with the scientific status of various theoretical elements of linguistics like functional projections, PRO, etc. in the same way that our intuitive concepts of space or causation do not change with advances in theoretical physics. Thus, appeals to the details of theoretical syntax force one’s position to be sufficiently naturalistic that they are vulnerable to the kinds of counterevidence offered above.)

What about the data discussed in (I3) – (I4)? Does it provide overwhelming linguistic evidence that knowledge-how is a species of knowledge-that? Is this data worth more than all the opposing evidence from philosophy, psychology, and neuroscience? No. As our previous discussion suggests, (I3) – (I4) only provide evidence regarding the *linguistic* structure of certain sentences, some of which have the form *NP knows how to VP*. Such data show only that various empirical phenomena can be *described* in human languages with verbs taking clausal complements. Some of these phenomena are instances of knowledge-how and others are not. In fact, it may even be that one and the same syntactic object can be used to describe different sorts of phenomena. If we accept SW’s discussion of the data, we can take (9) (i.e. *Mary knows how to juggle*) to be about multiple types of scenarios, some of which involve knowledge-how, and some of which don’t. That is, (9) sometimes describes an instance of knowledge-how, realized in Mary’s frontal and parietal cortices, basal ganglia, and dentate nucleus of her cerebellum. Other times (assuming SW’s discussion) (9) describes a phenomenon realized in Mary’s medial temporal lobe,

7 Some instances of knowledge-how, when they are being first learned, appear to be instantiated (at least in part) in the declarative memory systems of the brain, and then later are transferred to the procedural memory systems. But such facts if anything make things even worse for SW, since they suggest that the general question of whether knowledge-how is propositional or not cannot receive either the simple answer they endorse or the simple answer they reject.

hippocampus, elements of temporal and temporoparietal neocortex, anterior prefrontal cortex and some of the right cerebellum. The two interpretations of (9) also have strikingly different functional roles – both from a general philosophical perspective, and from the more fine-grained psychological perspective of cognitive processing. Neither phenomenon is a species of the other. The linguistic data only show that natural languages can group together disparate phenomena under a single type of linguistic expression. But this happens all the time.⁸ For instance, there are deep empirical differences between rocks, ideas, numbers, round squares, and trips to the mall. From the perspective of linguistic structure, however, they all appear to be very much the same insofar as natural language represents all of them with noun constructions. Similarly, *the book on the hill* can describe either the book on top of the hill or the book about the hill. These two uses of *on the hill* are very different, even though the relevant aspects of this phrase’s syntax are the same in both cases. Consider also how the sentence *John went through all the points* can be made true: John could walk, crawl, drive, or be catapulted through the points, or he could simply talk through them in his lecture. In a similar vein, (9) is true if one of several empirically different knowledge-states in Mary is realized, some of which involve one kind of circuit in the brain, others of which involve another.⁹

8 Of course, languages aren’t alone in grouping phenomena together in ways that don’t respect the empirical facts. For example, under the right conditions, various parts of the human visual system will give identical interpretations to the boundary between a mountain and a car that is in front of it as it does to certain patterns of multicolored light on a white surface. If it didn’t we wouldn’t have movies.

9 I take it to be a straightforward empirical question which true sentences of the form *X knows how to VP* are due to X’s having some bit of know-how as opposed to some bit of propositional knowledge. We settle this question by first clarifying the semantics of *X knows how to VP* to the point where we can collect a relevant range of cases in which sentences of this form are true (and perhaps another relevant range of cases in which it’s false). We then take our best going theory of what knowledge-how is, and when it is physically different from propositional knowledge, and we study how the various cases in question might be realized. Unless this aspect of language is unlike every other topic in cognitive science, we may find some interesting imperfect correlations and tendencies, but we will not find that particular interpretations are precisely determined in all normal humans by one sort of brain process and never another.

4 Some Broader Implications

So far, we saw that an analysis of the details undermines the purported connection between the linguistic structure of knowledge-how reports and knowledge-how itself. I want to briefly conclude by gesturing at another area of linguistically inspired philosophical research, namely the existence and nature of propositions. I believe that much of the work on propositions has foundered on problems similar to the ones we've seen above. Although I cannot here make the full case for this strong claim, I hope that the following remarks will indicate the general nature of my worry.

Many philosophers have attempted to draw inferences about the existence and nature of propositions and other abstract objects on the basis of the allegedly "logical" behavior of English (e.g., Schiffer 2003a, b¹⁰, King 1995, 1996, and many others). But such inferences appear to straightforwardly assume a strong connection between linguistic and metaphysical structure: if (our best theories of) our language sometimes treats the world as containing abstracta, then these abstracta really exist, and are quite literally, things that we mean when we speak (e.g., Schiffer 2003a). But we've seen ample reason to question such reasoning. A major impetus for putting propositions into one's ontology is that our language sometimes seems to treat clausal constructions as objectlike entities. But as mentioned above, it really shouldn't be too surprising that our linguistic system often works by linguistically "reifying" various phenomena, representing them with nominal constructions. Doing so would allow speakers to organize situations where others are being interpreted so that our remarkably well-developed human abilities for spatial cognition could be exploited. It may be that our language allows us to nominalize clauses (thus allegedly supporting the view that propositions exist) simply because such a procedure makes it easier for our minds to manipulate and reason with the information involved. It's clear that we often represent various situations spatially, even though there's nothing genuinely spatial involved. For instance, notice how common and easy it is to think of the empty set being literally *inside* the set whose only element is the empty set. But in actual fact, the axioms of set theory only tell us that these two abstract objects are related by a binary relation that observes certain logically expressible laws. (One could easily construct a model of set theory in which the empty set was identical to Mount Rushmore, and the set whose only element was the empty set was the Golden Gate Bridge.)

In a similar spirit, other researchers have suggested that many abstract objects can be identified with the metaphysical counterparts that are suggested by the mathematical tools used to represent human linguistic abilities. For instance, following Stalnaker (1984), Aczel (1980, 1987) and Chierchia and Turner (1988) both suggest that we treat properties and relations as functions from (n-tuples) of entities to propositions, where the latter are simply functions from possible worlds to truth-values. It would be hard to find a more straightforward assumption that semantic structure mirrors metaphysical structure than in this maneuver. Furthermore, general appeals to Ockham's Razor or to "the overall simplicity of one's total theory of the universe" are of little use here. We cannot infer that the semantics of natural language really is made up of functions from worlds to truth values simply because one particular choice of architecture for a semantic theory presents them as such. And we can't infer that the properties in the world are just as our language presents them, because our language may not accurately represent them. (In a similar spirit, one could easily design a language in which simple predications such as [NP VP] represented, in addition to the predication, that a golden mountain exists on X, an as-yet unobserved planet. But just because this language represents the world as being this way does nothing to imply that there is such a mountain on X, or even that X exists. Appeals to an overall simpler theory, here, as with propositions, are simply out of place, because too little is known about propositions and X to say – or to need to say – much about them.)

Once we demand an argument that some particular bit of linguistic structure actually mirrors a bit of metaphysical structure, human languages start to appear much less evidentially well-connected to the world than they sometimes seem. If I am right, then much of traditional philosophy of language must be reconsidered. We must realize that we cannot, without further argument, study the nature of the universe – e.g., the nature of modality and possibility and the (non-) existence of abstracta – by studying the nature of our language. (While it's a much more subtle issue, I even believe that we need to be very careful about inferring claims about psychological structure on the basis of (theories of) linguistic structure, and vice-versa.) Although I've focused here on language, I believe that a similar argument also shows that we cannot draw inferences about the nature of the universe solely on the basis of the structure of human psychology; e.g., the structure of our concepts and reasoning practices.

¹⁰ Cf. Schiffer's (2003a) endorsement of the "face-value theory" of propositions on p. 11.

5 Conclusion

Insofar as our expressions can represent the world, these representations can be imperfect. Where does this leave us? We have seen that a sophisticated attempt to uncover the nature of knowledge-how fails to achieve its goal, because of the difficulties with the inference from the linguistic structure of certain knowledge-how reports to the metaphysical structure of knowledge-how itself. Ultimately, I suspect that there is little metaphysics to be done by inspecting the structure of linguistic theories. My skepticism here is driven by the divergences between linguistic theories and the world, a few of which I've detailed in this paper. Of course, there may be some places where our linguistic theories accurately capture the structure of things in the world. But these seem to be cases where the worldly structure is quite simple and well understood independently of any linguistic analysis. The real challenge comes when we attempt to infer the existence or nature of more mysterious entities, which are not well understood independently of language. With regard to the study of the latter sort of phenomena, it appears that language may often fail to be a useful exploratory device.

References

- Aczel, P. 1980. Frege Structures and the Notions of Propositions, Truth, and Set. In *The Kleene Symposium*, edited by J. Kiesler, J. Barwise, and K. Kunen, 31–59. Amsterdam: North-Holland
- Aczel, P. 1989. Algebraic Semantics for Intensional Logics I. In *Properties, Types and Meaning*, edited by G. Chierchia, B. Partee, and R. Turner, 17–45. Dordrecht: Reidel.
- Burgess, N., K. J. Jeffrey, and J. O'Keefe. 1999. *The Hippocampal and Parietal Foundations of Spatial Cognition*. Oxford: Oxford University Press.
- Chierchia, G., and S. McConnell-Ginet. 1990. *Meaning and Grammar*. Cambridge, MA: MIT Press.
- Chierchia, G., and R. Turner. 1988. Semantics and Property Theory. *Linguistics and Philosophy* 11: 261–302.
- Chomsky, N. 1980. *Rules and Representations*. New York: Columbia University Press.
- Cohen, N. J., and L. Squire. 1980. Preserved Learning and Retention of a Pattern-Analyzing Skill in Amnesia: Dissociation of Knowing How and Knowing That. *Science* 210: 207–210.
- Haegeman, L. 1994. *Introduction to Government and Binding Theory* (2nd ed.). Oxford: Blackwell.
- Higginbotham, J. 1989. Knowledge of Reference. In *Reflections on Chomsky*, edited by

- Alexander George, 153–174. Oxford: Blackwell.
- Higginbotham, J. 1992. Truth and Understanding. *Philosophical Studies* 65: 3–16.
- Higginbotham, J. 2001. On Referential Semantics and Cognitive Science. In *The Foundations of Cognitive Science*, edited by J. Branquinho, 145–156. Oxford: Clarendon Press.
- Higginbotham, J. 2004. On Linguistics in Philosophy, and Philosophy in Linguistics. *Linguistics and Philosophy* 25: 573–584.
- Hoffman, D. 1998. *Visual Intelligence: How We Create What We See*. New York: W. W. Norton and Co.
- Jackendoff, R. 2002. *Foundations of Language*. Oxford: OUP.
- King, J. 1995. Structured Propositions and Complex Predicates. *Nous* 29: 516–35.
- King, J. 1996. Structured Propositions and Sentence Structure. *Journal of Philosophical Logic* 25: 496–521.
- Matthews, R. 1994. The Measure of Mind. *Mind* 103: 131–46.
- Reid, T. 1785. *Essays on the Intellectual Powers of Man*, A. D. Woolzey (ed.) (1941). London: Macmillan and Co.
- Ryle, G. 1946. Knowing How and Knowing That. *Proceedings of the Aristotelian Society* XLVI: 1–16.
- Schiffer, S. 2003a. *The Things We Mean*. Oxford: OUP.
- Schiffer, S. 2003b. Knowledge of Meaning. In *Epistemology of Language*, edited by Alex Barber, 303–24. Oxford: OUP.
- Squire, L. 1992a. Memory and the Hippocampus: A Synthesis from Findings with Rats, Monkeys and Humans. *Psychological Review* 99: 195–231.
- Squire, L. 1992b. Declarative and Nondeclarative Memory: Multiple Brain Systems Supporting Learning and Memory. *Journal of Cognitive Neuroscience* 4: 232–243.
- Squire, L., B. Knowlton, and G. Musen. 1993. The Structure and Organization of Memory. *Annual Review of Psychology* 44: 453–95.
- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stanley, J., and T. Williamson. 2001. Knowing How. *The Journal of Philosophy* 98: 411–444.
- Ullman, M., S. Corkin, M. Coppola, G. Hickok, J. Growdon, W. Koroshetz, and S. Pinker. 1997. A Neural Dissociation within Language: Evidence that the Mental Dictionary is Part of Declarative Memory, and that Grammatical Rules are Processed by the Procedural System. *Journal of Cognitive Neuroscience* 9: 266–76.
- Ullman, M. 2001. A Neurocognitive Perspective on Language: The Declarative/Procedural Model. *Nature Reviews: Neuroscience* 2: 717–726.