Linguistic solutions to philosophical problems
Barbara Abbott
Michigan State University
abbottb@msu.edu

In this paper I want to focus on a particular attempt to apply linguistic evidence to the solution of a philosophical problem. The problem is whether or not there is a distinction between knowledge how and knowledge that. Ryle (1949) argued that there is. Jason Stanley and Timothy Williamson (hereinafter SW) have argued, based on linguistic analysis, that knowledge-how is not different from knowledge-that, but is instead just a subcategory of propositional knowledge (SW 2001). At least two issues can be distinguished here. One is, what exactly is the linguistic evidence telling us about knowledge how vs. knowledge that? Another is, should we believe what linguistic evidence tells us? The first section of the paper will briefly review SW's paper; the following two sections will take up the two issues in order.

1. SW on knowing how.

The construction of interest is that illustrated in (1), which must be distinguished from sentences like (2) and (3). (All three sentences, as well as (4) below, are from SW 2001.)
(1)   Hannah knows how to ride a bicycle.
(2)   Hannah knows how Bill rides a bicycle.
(3)   Hannah knows that penguins waddle.
(1) is distinguished from (2) and (3) by the infinitival (untensed) form of the complement verb (*to ride*), and the lack of an overt subject for this verb. SW point out that, under currently popular theories of syntax, sentence (1) would receive the analysis in (4).
(4)   Hannah knows [how PRO to ride a bicycle *t*]
They remark: '"PRO" here is a phonologically null pronoun that occurs, according to standard syntactic theory, in the subject position of untensed clauses' (419). (The *t* in (4) is called a 'trace'; it marks the spot where *how* is assumed to have originated.) It is easy to see that on this analysis, the verb *know* is taking a clausal or sentential complement, and hence a complement which might be expected to express a proposition. As far as arguing that knowledge how is propositional in nature, one could almost say 'QED' at this point, except for the fact that almost nobody would be convinced by one brief paragraph. Also, there are some other complications.
        One complication is the interpretation of PRO. As SW note (423), the fact that both examples in (5) seem acceptable
(5)   a.   Mary knows how to behave herself at a horse race.
          b.   Mary knows how to behave oneself at a horse race.
(although (5a) sounds somewhat better than (5b)), suggests that PRO may receive either an interpretation on which it is coreferential with the subject of *know*, or one on which it

stands for an unspecified individual, somewhat similar to the interpretation of *one*. For the most part for our purposes, only the first type of interpretation will be relevant.

Another complication is the exact interpretation of the complement clause. Complements like *how to ride a bicycle* in (1) (and *how Bill rides a bicycle* in (2)) are called embedded questions. SW first evoke the well known analysis of Karttunen 1977, on which embedded questions actually denote, not a single proposition, but instead a set of propositions – those that would constitute true answers to the question. Thus *how PRO to ride a bicycle*, as it occurs in (1), would denote the set of propositions expressed by true sentences of the form *x is a way for Hannah to ride a bicycle*. However subsequently SW acknowledge that (1) does not actually attribute to Hannah knowledge of **all** the ways she can ride a bicycle, but instead only at least one. At this point (leaving out some details of SW's presentation) (1) is being claimed to be true 'if and only if, for some contextually relevant[1] way *w*…Hannah knows that *w* is a way for her to ride a bicycle' (426). We now have equivalence between (1) and a sentence with *know* and a tensed *that*-complement.

A brief digression: thus far the SW analysis is quite similar to that of Brown (1970). However Brown finds two variants of the *know how to* construction. Brown analyses *x knows how to V* in one sense as 'x knows, of some way *w*, that it is **a** way one can V', while the other type he analyses as 'x knows, of some way *w*, that it is **the** way to V'. The difference can be illustrated by the examples in (6) (from Brown, pp. 219 and 241, respectively):

(6)   a.   John knows how to build frame houses.
        b.   John sure knows how to build frame houses.

(6a) would most likely be taken the first way – as merely attributing to John the ability to build frame houses, but (6b) attributes to John knowledge of a very good way, or the best way, to build them. Two other of Brown's examples of the second, 'superlative', sense are given in (7) (his examples 4 and 5 respectively).

(7)   a.   John knows how to move about in a canoe.
        b.   The accused does not know how to address the magistrate.

Brown's analysis seems to capture the difference between the two understandings. On the other hand the phenomenon illustrated here appears to be comparable to other instances of hyperbole in describing skills and abilities, of which the sentences below are further examples.

(8)   a.   Louise sure can dance.
        b.   You painted the hell out of that picture!
        c.   I couldn't add my way out of a paper bag.

Ultimately it may not be necessary or even advisable to recognize an additional sense specifically of the *know how to* construction to account for this type of thing. Rather, it might be better accounted for with a general theory of hyperbolic usage.[2]

---

[1] It's not clear what is the significance of the phrase 'contextually relevant', which was introduced by SW without explanation. At first glance it seems unnecessarily restrictive; it seems that someone can tell me out of the blue that, e.g., Kim knows how to make model airplanes out of Jell-O, without implicating any particular way of doing so, contextually relevant or otherwise.

[2] It should be noted that Brown did not ultimately reject any sort of knowledge how vs. knowledge that distinction. Instead, his position seems to be that the labeling of the distinction is inaccurate and misleading, given that the distinction itself cross-cuts our grammatical resources for describing knowledge. Brown suggested the term 'practical knowledge' as better to capture what Ryle had in mind with

Returning to the main thread, SW note a problem for their analysis (apparently one not noticed by Brown). Suppose that Hannah does not know how to ride a bicycle. Then one could hardly teach her how to ride, in the relevant sense, by simply pointing out to her another person riding a bicycle in the usual way. The final adjustment in the analysis is to assert that (1) attributes to Hannah knowledge of the proposition in question under what SW call 'a *practical* mode of presentation' (429; italics in original). Not much more is said about what this practical mode of presentation involves; only that 'there is a conventional connection between the use of constructions that embed instances of the schema "how to *F*", and practical modes of presentations of ways' (429), and that the practical mode of presentation 'undoubtedly entails the possession of certain complex dispositions' (429). These complex dispositions would ordinarily entail the ability to ride a bike. However two provisos must be noted. First of all, SW agree with Ginet (1975) in holding that loss of ability to do something because of injury or other abnormal circumstance does not mean loss of knowledge how to do it; 'a master pianist who loses both of her arms in a tragic car accident still knows how to play the piano' (SW, 416). And secondly, if the action in question is complex we might attribute knowledge how in the absence of ability. SW give the example of a ski instructor who knows how a complex stunt should be performed, but cannot do it herself. And Koethe (2002) points out that we might say that it is true that Paul Revere knows how to alert the populace to the imminent arrival of the British (namely, as Koethe says, by getting on a horse and riding around to every Middlesex village and farm), even if Paul doesn't know how to ride a horse.

Yet as the actions become more basic, and riding a bicycle would be included here, it seems that knowledge how must imply ability in normal circumstances, e.g. barring a physical disability of some kind. This is acknowledged by SW themselves, towards the end of their paper, when they are discussing David Lewis's response to Frank Jackson's famous 'knowledge argument' (Jackson 1982). Lewis (1988) (following Nemirow 1980) argues that what Mary has acquired following her release from her black and white prison is a set of abilities – specifically, the ability to recognize the color red, to imagine what red looks like (both to yourself and to others), and to remember the experience of seeing red. Lewis says:

> knowing what an experience is like just *is* the possession of these abilities to remember, imagine, and recognize. It isn't the possession of any kind of information, ordinary or peculiar. It isn't knowing that certain possibilities aren't actualized. It isn't knowing–that. It's knowing-how. (Lewis 1988, 593.)

SW's reply is, of course, that knowledge how **is** knowledge that. Furthermore even if Lewis were to have left off this comment and spoken just in terms of abilities, it would not help, since the possession of this type of intentional ability really coincides with knowledge how. They say both that 'it is plausible that the ability to imagine an experience of red entails knowing how to imagine an experience of red', and that it 'seems absurd' to hold that Mary knows how to imagine an experience of red while she is

---

'knowledge how', whether this is described with or without the words 'how to'. Ryle had noted, in describing the special kind of knowledge he was concerned with, that it is demonstrated by doing rather than explaining, it is 'the exercise of intelligence in practice' (Ryle 1949, 40), and Brown concurs that this is the essence of the distinction. Cf. Brown 1975, 247f.

in her black and white cell: 'If she knows how to imagine an experience of red, why is she unable to imagine such an experience?' (SW 2001, 443). Thus it seems that we should take the practical mode of presentation introduced by SW to imply ability when the action in question is a simple one, and the actor would not be prevented from carrying out the action by some other circumstance. (Another possibility, proposed by Rosefeldt (2004), is that *know how to* is ambiguous depending on whether or not ability is implied.)

Modes of presentation have been used elsewhere in philosophy of language as a way of making the Russellian analysis of propositions consistent with some problematic aspects of propositional attitude contexts. As such they are sometimes deemed to be semantic, but other times only pragmatic. In the current case, although the connection between the 'how to' construction and the practical mode of presentation is said to be conventional, SW assert that that does not preclude it 'from being of only pragmatic significance', and declare themselves 'neutral on the issue of whether modes of presentation affect semantic content' (429, n. 31). It is worth noting that if the practical mode of presentation was held not to affect semantic content, in the situation described above, where Hannah does not know how to ride a bicycle and we merely show her a person who is riding one, (1) would be said to be true. But as just noted, knowledge how, when the action is a basic one and there are no impeding circumstances, should entail ability.

2. What is the linguistic evidence really telling us?

The analysis of sentences like (1) which SW propose can be questioned at several points. To begin with, although the assumption that infinitival complements as in (1) have a PRO subject may be considered to be standard in current linguistic theory, this is not the only possibility, or even the only possibility within recent linguistics. Montague (1973), for example, analyzed similar infinitival VPs – specifically, the complements of *try* and *wish* as in *Mary tried/wished to find a unicorn* – simply as VPs. This type of analysis has also been adopted in, e.g., GPSG and its descendent HPSG; see Gazdar et al., 1985, ch. 6; Pollard & Sag, 1994, ch. 3. Applying this analysis in the present case, *know* would be subcategorized both for a full tensed sentential complement and for an infinitival VP complement. The structure of (1) would be that in (9)

(9)   Hannah knows [$_{VP}$ how to ride a bicycle]

The fact that the abilities in question are abilities of the knower could be attributed to the real world fact that it is one's own knowledge, and not that of another, that enables one to do things. The agreeing type of reflexive pronoun, as in (5a), repeated here as (10)

(10)   Mary knows how to behave herself at a horse race.

would automatically fall under Principle A of Chomsky's binding theory, which says, roughly, that reflexive pronouns must agree with an antecedent in the smallest containing sentence which has a subject. The less felicitous nonagreeing reflexive pronoun, *oneself* in (5b), repeated here as (11)

(11)   Mary knows how to behave oneself at a horse race.

would have to be dealt with as a special case (just as in the 'standard' analysis), perhaps by regarding *oneself* as a special type of non-referring reflexive.

Dowty, Wall & Peters (1981, 235f) point out that Montague's approach has the advantage of automatically ruling out complements with an expressed subject for predicates like *try to* and *wish to*. Similar facts hold for *know how to*, as in (12).

(12)　　　　* Hannah knows how Bill to ride a bicycle.

A Montague-type analysis also accords better with our intuition that the know-how construction relates entities with abilities, exactly what is being denied by SW. (Of course ultimately the grammar must recognize that *ride* is a predicate which takes two arguments, and that (1) in the relevant sense attributes to Hannah an ability in which she instantiates the first of those argument slots.)

From their own perspective of trying to show that knowledge how is simply another type of knowledge that, the weakest part of the SW analysis is their invocation of the practical mode of presentation. As they acknowledge (433f), it may well be the case that this concept cannot be explicated without crucial reference to knowing how. Koethe (2002) points out that knowing that w is a way to F under a practical mode of presentation seems to amount to knowing how to instantiate w oneself,[3] in which case SW have a problem either of infinite regress, or of a different type of knowing how which they have not analyzed.

In any case ultimately there is a strong inclination on the part of those who remain unconvinced by SW to say, concerning the practical mode of presentation, that that **is** the crucial difference between knowledge how and knowledge that – the requirement that the knower acquire skills, abilities. SW have only a half-hearted response to this objection. They say: '…one might worry that significant disanalogies still remain between knowing-how and other kinds of knowing-that' (435), but go on at this point to discuss just one possible disanalogy – susceptibility to Gettier cases. The allusion is to Ed Gettier (1963), who challenged the 'justified true belief' analysis of knowledge with examples where the 'justification' for the belief in question is independent of the facts, and only accidentally accurate. We have a strong intuition that such a belief does not constitute knowledge.

SW proceed to construct a Gettier case for knowledge how: Henry is an unscrupulous but also incompetent flight instructor imposter who trains Bob on a scrambled, but accidentally accurate, flight simulator and with advice that he believes is faulty but which happens to be correct. Bob passes the course with flying colors, but SW assert that 'there is a good sense in which he does not *know* how to fly' (435). I have to confess that my intuitions are different from those of SW on this example. It seems to me that if passing the course with flying colors results in the ability to fly a plane (which it certainly ought to), then Bob does know how to fly a plane whether or not his acquisition of that knowledge was via deception.

In any case I do not think that one's intuitions about this Gettier case are crucial. Even if SW were correct in their claim that knowledge how is as susceptible as knowledge that to such cases, it would not detract from the fact that there remains a significant disanalogy between the two, which is the fact that knowledge how implies the acquisition of skill (the 'practical' mode of presentation) and knowledge that does not.

---

[3] Rosefeldt (2004) makes a similar observation, and notes in addition that ultimately, this clause of the analysis of knowing how (the one stipulating the practical mode of presentation) is the one that matters, and the one attributing propositional knowledge to the knower-how (i.e. the clause saying that w is a way to F) is not only unnecessary but inaccurate (375).

There are some other aspects of SW's paper which have been challenged. SW address the possibility of invoking an ambiguity in the verb *know*, corresponding to its use in sentences like (1) as opposed to its use in sentences like (3), repeated here as (13), which have overt tensed *that* clauses.

(13)   Hannah knows that penguins waddle.

They assert that conjunction reduction, as in (14)

(14)   Hannah knows that penguins waddle, and Bill, how to imitate them.

does not result in anomaly (as would be expected were *know* ambiguous). They also say: 'In the languages with which we are familiar, the uses of "know" in [1] and [13] are translated by the same word' (437). However both claims have been disputed.

Rumfitt (2003) says, concerning the example below:

(15)   John knows both how to twitch his ears and that his mother is sickened by facial tricks. [= Rumfitt 2003, ex. 23]

'I am not alone in finding [15] zeugmatic' (165). Rumfitt also points out that although French does use the same verb (*savoir*) for the two kinds of knowledge sentences, it does not use embedded questions for the know how type of examples. Thus (1) would be translated as something like (16a) in French, and not as (16b).[4]

(16) a.   Hannah   sait   monter   à   vélo.
                     knows   to ride   to   bicycle

      b. # Hannah   sait   comment   monter   à   vélo.
                     knows   how   to ride   to   bicycle

(# is the symbol conventionally used within linguistics for anomaly.) According to Rumfitt, neither Greek nor Latin uses an embedded question for this construction either.

Rumfitt notes further that Russian **does** use two different verbs for the two constructions: *umet'* for knowledge how, and the unrelated verb *znat'* for knowledge that. Thus sentences (1) and (13) would be translated as (17a) and (17b) respectively.[5]

(17) a.   Hannah   umeet   katat'sia   na   velosipede
             Hannah+nom   knows   to roll self   on   bicycle+loc

      b.   Hannah   znaet   chto   pingviny   khodiat   vrazvalku
                      knows   that   penguins+nom   walk+3pl   waddlingly

As Rumfitt notes, *umet'* occurs only with infinitives while *znat'* does not occur with infinitives; the two verbs could not be switched in the examples in (17). Similar facts hold for the Native American language Montana Salish, the verb *yo'* being used to

---

[4] There is some difference of opinion on how (1) is best translated into French. In (16a) I have used a variant of the translation provided by Rumfitt (specifically, his ex. 5, p. 161). However a colleague in the French Department at MSU (Anne Violin-Wigent) offers (i) below, and the online AltaVista Babel Fish translation service suggests (ii).

(i)   Hannah   sait   faire   du   vélo.
                      knows   to make   of the   bicycle

(ii)   Hannah   sait   monter   une   bicyclette
                                                  .
                      knows   to ride   a   bicycle

Despite these differences, all agree in not having *comment* ('how') in the translation.

[5] I am grateful to my colleague Jason Merrill for these translations.

express knowledge how, and *míy* for knowledge that.[6] And finally Rosefeldt (2004) points out that German, with which SW implicate familiarity, also uses two different verbs – *wissen* (for knowledge that), and *können* (for knowledge how). (Rosefeldt asserts that ch. 2 of Ryle's *The Concept of Mind* ('Knowing How and Knowing That') is properly translated 'Wissen und Können' (377).)

We should also note the existence of the compound *know-how*, in English, meaning '[t]he knowledge and skill required to do something correctly' (Soukhanov, 1992). As Rumfitt points out, this term, like *savoir faire* in French, is used specifically for the type of knowledge we are talking about. This is reflected in the naturalness of (18a), which contrasts with (18b) and (18c).

(18) a. Knowing how to fix a leak is part of any competent plumber's know-how.
 b. # Knowing how Trotsky was killed is part of any competent historian's know-how.
 c. # Knowing why Trotsky was killed is part of any competent historian's know-why.

(These examples are Rumfitt's (26)-(28), with the anomaly sign (#) added for the last two.)

In sum, it is not so clear that the linguistic evidence is telling us that there is no significant difference between knowledge how and knowledge that, and that the former is just a subtype of the latter. Instead, at least part of the evidence is pointing in the opposite direction. But we should consider how seriously to take linguistic evidence of this type.

3. Should we believe linguistic evidence?

J.L. Austin is probably the philosopher best known for advocating the investigation of natural language in the search for insights into philosophical issues.[7] In 'A plea for excuses' he suggested a justification of that approach with the following comment:

> [O]ur common stock of words embodies all the distinctions men have found worth drawing, and the connexions they have found worth marking, in the lifetimes of many generations: these surely are likely to be more numerous, more sound, since they have stood up to the long test of the survival of the fittest, and more subtle, at least in all ordinary and reasonably practical matters, than any that you or I are likely to think up in our arm-chairs of an afternoon – the most favoured alternative method. (Austin 1956-7, 182.)

We can acknowledge the truth of what Austin says here, but still find some limitations in this application of language data. (Austin himself noted a couple, one of which we will come to below.)

There are at least three reasons why one might question the use of linguistic evidence to provide answers to the question of whether knowledge how is different from knowledge that, and one or more of these reasons may also apply in other cases of the application of linguistic evidence to problems in philosophy. These three reasons are

---

[6] The prime in *yo'* represents a glottal stop. I am grateful to Sarah Thomason for telling me about the Montana Salish facts, and providing me with many examples of the two verbs in action.
[7] I am grateful to Rich Hall for reminding me of the relevance of Austin's work to the topic of this paper, and pointing out the quotes from Austin below.

first, unclarities in the evidence itself, because of questions about linguistic analysis, shakiness of intuitions, or some other factor; second, the fact that a language may not incorporate distinctions which might be of philosophical importance; and finally, the fact that the linguistic evidence may embody a false view of the area in which insight is sought. Let us look at each of these in turn.

3.1. Questions about the linguistic evidence. The first pitfall in using linguistic evidence to solve philosophical problems has been amply illustrated in the foregoing. It involves both questions about particular linguistic analyses, as well as questions about the nature and significance of more or less raw pieces of linguistic data.

In the present case, SW's analysis of knowledge how sentences like (1) includes a subject for the embedded complement, which is a crucial part of their argument that this complement is a clause, and thus that knowledge how is propositional. However superficially 'knowledge how' sentences do not have fully clausal complements, leaving it open to propose a different analysis of the linguistic data, one without an embedded clause subject. As often occurs in this kind of analysis, there seems to be a trade-off between two ways of doing things, without any absolutely decisive evidence in favor of one over the other.

We also saw variation in the claims about the data. Is the verb *know* in English ambiguous? Ambiguity claims are notably difficult to either establish or rebut, depending as they do on intuitions about example sentences like (14) and (15) above – intuitions which may be no clearer than intuitions about the words themselves. We have seen that at least German, Russian, and Montana Salish use different verbs where English would use *know*, but the uses of *wissen*, *znat'*, and *míy*, or *können*, *umet'*, and *yo'*, most likely do not correspond exactly to each other, or to the different uses of *know* which are marked by different complement structures in English. Even within what seems like a single use of *know* in English – namely, the *know how to* construction – we have seen suggestions that there are actually two senses depending on whether or not ability is implied (Rosefeldt), or whether mere ability or a relatively high level of skill is implied (Brown).

I would not want to say that linguistic evidence is never clear, or that there are no analyses which are well agreed on, but only that there may be difficulties of both types, and that the case with *know* seems to exemplify them.

3.2. Relevant linguistic distinctions might not exist. As Barwise & Perry (1983) have pointed out, natural language is efficient. That means, among other things, that single forms can perform multiple services of expression. Barwise & Perry were referring specifically to the indexicality of language, e.g. the fact that a sentence like *I'm right and you're wrong* (their example, p. 5), can be used by different people on different occasions to make different claims. But natural languages are efficient in other ways too – they are loaded with generality, vagueness, and polysemy. For example Green (1996, 11) points out that an order to write your name backwards can be construed in a number of different ways (e.g. to write the last name first, or to write the letters in reverse order, or to write the letters backwards, or to stand backwards while writing your name). And others (e.g. Grice 1975, 1989; Sperber & Wilson 1986; Recanati 1989; Bach 1994) have for some time stressed the extent to which we elaborate and fill in information in an

utterance using (among things) our knowledge of how things work in the world. To take one often cited type of example, a word like *red* is interpreted differently when we are talking about apples (the skin is red), grapefruits (the edible part is red), hair (it's really orange), and so forth. The point here is that even if the verb *know* in English is unambiguous, that does not mean that it might not be used to refer to two (or more) very different kinds of knowledge.

A similar point goes for complement constructions. Syntactically there are relatively few complement types in English, each of which can signify quite different kinds of things. Thus *that* clauses may be used to refer to, e.g., facts, propositions, events, or utterances, as illustrated in the examples below.

(19) a. <u>That mice love cheese</u> makes them easy to catch.
      b. Mary disputed the idea <u>that mice love cheese</u>.
      c. We saw <u>that the car had crashed</u>.
      d. Lee was so astounded <u>that she choked on her cereal</u>.
      e. Max mumbled <u>that the damned car had crashed</u>.

Simply put, we should not expect a natural language to grammaticalize all the different distinctions we might want to draw. That would not be efficient. So again, even if SW are correct in their argument about the syntactic structure of the complement of *know*, that does not mean that all such sentences are talking about the same kinds of knowledge.

There are other reasons why a difference in kinds of knowledge might not be marked linguistically. For one thing, there might be similarities among these different kinds of knowledge in how they function to guide our behavior. Or people might not be aware of these different kinds of knowledge, or they might be aware of them, but not feel the need to discriminate among them in their everyday talk. Linguists and linguistic philosophers do tend to seize on perceived linguistic distinctions as evidence for important distinctions in nature, and language data may provide a kind of one way test for that – i.e., if the linguistic distinction exists, it may be evidence for a distinction in nature (although it also may not be, as we will see in the following subsection). But language definitely cannot be taken as a two way test for significant distinctions; a missing linguistic distinction need not be taken as evidence for the absence of a corresponding distinction in the world.

3.3. Language may misrepresent the world. Supposing we do find a linguistic distinction, it still might not be the case that that distinction would point to a valid or useful real world distinction. If we are looking to natural language to provide evidence for how the world is, we need to remember (as Austin reminds us) that natural languages evolve via the expressive needs of their speakers, who are ordinary human beings. Thus whatever patterns and distinctions may be drawn will be no better or worse than the insight of those speakers. Practitioners in the hard sciences typically would not look to natural language for advances in their respective fields; botanists (*qua* botanists) do not really care about the distinction ordinary speakers draw between fruits and vegetables, and astronomers would be ill-advised to pay much heed to the fact that we describe the sun as rising in the morning, and going down in the evening, as evidence bearing on the structure and operation of the solar system. As Austin himself said (this was one of the limitations of the linguistic method alluded to above):

9

> If a distinction works well for practical purposes in ordinary life (no mean
> feat, for even ordinary life is full of hard cases), then there is sure to be
> something in it, it will not mark nothing: yet this is likely enough to be not the
> best way of arranging things if our interests are more extensive or intellectual
> than the ordinary. And again, that experience has been derived only from the
> sources available to ordinary men throughout most of civilized history: it has
> not been fed from the resources of the microscope and its successors. (Austin
> 1956/7, 185.)

The present case involves the mind, and more specifically philosophical speculations about ways of knowing. In this case what we can expect from natural language is an incorporation of folk psychology. As Churchland (1988) has argued so convincingly, there is no reason to think that future developments in the relevant sciences will not supercede the received folk psychological wisdom on what kinds of knowledge there are, making the contribution of linguistic analysis irrelevant. The sciences of the mind and brain might find that there are indeed significant differences among the ways different kinds of information are stored and processed, but these different ways may not correspond exactly to any difference in ways we talk about knowledge.

A relevant example comes from the study of long term memory. For some time psychologists have postulated a major difference between explicit (or declarative) memory and implicit (or nondeclarative, or procedural) memory. Probably the most vivid piece of evidence in favor of a distinction along these lines is the famous case of the patient H.M., part of whose temporal lobes were removed in 1953 in order to control seizures. After the surgery it was found that H.M. could not retain new factual information; nevertheless he was able to acquire new skills, such as how to solve the 'Tower of Hanoi' puzzle, or how to trace patterns in a mirror (although he had no memory of the experiences of learning these skills). The distinction in kinds of memory made vivid by H.M. seems to correspond roughly to the knowledge that vs. knowledge how distinction, although the correspondence is not exact. Indeed, we can certainly **describe** him as having acquired new knowledge that – e.g., the knowledge that he can't remember things (which is apparently the case – see Milner 1970, 37). But the failure of his condition to correspond exactly to our ordinary ways of attributing knowledge does not detract from the significance of his case for the kinds of knowledge we have.

Furthermore cognitive psychologists commonly divide the two major types of long term memory into a number of subcategories. Within the category of explicit memory, memory for facts (also called 'semantic' memory) is distinguished from memory for events ('episodic' memory); and within the category of implicit memory acquired skills and habits may be distinguished from nonassociative learning and from classical conditioning. There is also evidence, as might be expected, for the involvement of different brain structures in different types of memory. Thus the significant types of long term memory appear to go well beyond our ordinary ways of talking. I should note that this brief sketch does not begin to do justice to the complexity of this area of research, in which psychologists are currently testing a number of competing models of memory. Nevertheless there seems to be general agreement that memory is not uniform.[8]

---

[8] Cf. e.g. Johnson 2005. I am grateful to Rich Hall and Carol Slater for reminding me of H.M. and the relevance of psychological studies of memory to the issue of knowledge how vs. knowledge that, and to Rose Zacks for confirming that relevance.

In the course of responding to Jackson's knowledge argument (mentioned briefly above in section 2), Churchland makes a comment which is very much apropos of the topic of this paper. Citing applications of a knowledge how vs. knowledge that distinction to Jackson's argument, he says:

> While the approach is well motivated, this binary distinction in types of knowledge barely begins to suggest the range and variety of different sites and types of internal representation to be found in a normal brain. There is no reason why we must be bound by the crude divisions of our prescientific idioms when we attempt to give a precise and positive explication of the equivocation displayed in Jackson's argument. (Churchland 1989, 572.)

Without taking a stand on whether or not there is equivocation in Jackson's argument, I want to concur on a couple of points. One is Churchland's implicit assumption that the issue of types of human knowledge is ultimately one to be resolved empirically, through the findings of the cognitive sciences. Thus I want to include the issue of types of knowledge among those issues Austin indicated might not be best served by an examination of natural language, evolving as it has without being 'fed from the resources of the microscope and its successors'. The other point is that when the cognitive sciences have finished their work, it seems certain that the types of knowledge discovered will go far beyond the 'crude divisions of our prescientific idioms'.

4. Concluding remarks.

In this paper I have tried to cast doubt both on SW's analysis of English sentences describing knowledge how, and on the conclusions which they draw from their analysis. In so doing I have tried to point out some of the ways in which linguistic data and/or the analysis thereof might be misleading or unenlightening about facts about the world and hence about the philosophical issues surrounding such facts. Of course this is not to say that linguistic solutions for philosophical problems can never be found; indeed, Austin's work stands witness to the fact that sometimes they can be. My point has only been that in the present case, and others which are similar in the relevant respects, natural language might not be the best place to look for answers.[9]

References

Austin, J.L. 1956-7. A plea for excuses. Proceedings of the Aristotelian Society. Reprinted in J.O. Urmson & G.J. Warnock, eds., 1979, J.L. Austin: Philosophical Papers, 3rd edition. Oxford: Clarendon Press, 175-204.

Bach, Kent. 1994. Conversational impliciture. Mind & Language 9, 124-162.

Barwise, Jon and John Perry. 1983. Situations and attitudes. Cambridge, MA: MIT Press.

Brown, D.G. 1970. Knowing how and knowing that, what. In Oscar P. Wood & George Pitcher, eds., Ryle: A collection of critical essays. Garden City, NY: Anchor Books, 213-248.

Gazdar, Gerald, Ewan Klein, Geoffrey Pullum & Ivan Sag. 1985. Generalized phrase structure grammar. Cambridge, MA: Harvard University Press.

Churchland, Paul M. 1988. Matter and Consciousness. Revised edition. Cambridge, MA: MIT Press.

Churchland, Paul M. 1989. Knowing qualia: A reply to Jackson. In A neurocomputational perspective. Cambridge, MA: MIT Press, 67-76. Reprinted in Ned Block, Owen Flanagan, and Güven Güzeldere, eds., The nature of consciousness: Philosophical debates. Cambridge, MA: MIT Press, 571-577.

Dowty, David R., Robert E. Wall & Stanley Peters. 1981. Introduction to Montague semantics. Dordrecht: D. Reidel.

Gettier, Edmund L. 1963. Is justified true belief knowledge? Analysis 23, 121-123.

Ginet, Carl. 1975. Knowledge, perception, and memory. Dordrecht: D. Reidel.

Green, Georgia M. 1996. Pragmatics and natural language understanding. Second edition. Mahwah, NJ: Lawrence Erlbaum.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan, eds., Syntax and semantics, vol. 3: Speech acts. New York: Academic Press, 41-58.

Grice, H. Paul. 1989. Studies in the way of words. Cambridge, MA: Harvard University Press.

Jackson, Frank. 1982. Epiphenomenal qualia. Philosophical Quarterly 32, 127-136.

Johnson, Mark K. 2005. Developmental cognitive neuroscience. 2nd edition. Malden, MA: Blackwell.

Karttunen, Lauri. 1977. Syntax and semantics of questions. Linguistics and Philosophy 1, 3-44.

Koethe, John. 2002. Stanley and Williamson on knowing how. Journal of Philosophy 99, 325-328.

Lewis, David 1988. What experience teaches. In J. Copley-Coltheart, ed., Proceedings of the Russellian Society. Reprinted in Ned Block, Owen Flanagan, and Güven Güzeldere, eds., The nature of consciousness: Philosophical debates. Cambridge, MA: MIT Press, 579-595.

Milner, Brenda. 1970. Memory and the temporal regions of the brain. In Karl H. Pribram and Donald E. Broadbent, eds., Biology of memory. New York: Academic Press, 29-50.

Montague, Richard. 1973. The proper treatment of quantification in ordinary English. In Jaakko Hintikka, Julius Moravcsik & Patrick Suppes, eds., Approaches to natural language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics. Dordrecht: Reidel, 221-242. Reprinted in Richmond Thomason, ed., 1974, Formal philosophy: Selected papers of Richard Montague. New Haven: Yale University Press, 247-270.

Nemirow, Lawrence. 1980. Review of Mortal Questions, by Thomas Nagel, Philosophical Review 89, 475-6.

Pollard, Carl & Ivan Sag. 1994. Head-driven phrase structure grammar. Chicago: University of Chicago Press.

Recanati, François. 1989. The pragmatics of what is said. Mind and Language 4, 295-329.

Rosefeldt, Tobias. 2004. Is knowing-how simply a case of knowing-that? Philosophical Investigations 27, 370-379.

Rumfitt, Ian. 2003. Savoir faire. Journal of Philosophy 100, 158-166.

Russell, Bertrand. 1912. The problems of philosophy. Oxford: Oxford University Press.

Ryle, Gilbert. 1949. The concept of mind. London: Hutchinson.

Soukhanov, Anne, executive ed. 1992. The American heritage dictionary of the English language. Third edition. Boston: Houghton Mifflin.

Sperber, Dan & Deirdre Wilson. 1986. Relevance: Communication and cognition. Cambridge, MA: Harvard University Press.

Stanley, Jason & Timothy Williamson. 2001. Knowing how. Journal of Philosophy 98, 411-444.