

The Baltic International Yearbook of
Cognition, Logic and Communication

October 2010 Volume 5: *Meaning, Understanding and Knowledge*
pages 1-17 DOI: 10.4148/biycl.c.v5i0.281

MITCHELL GREEN
University of Virginia

LANGUAGE UNDERSTANDING AND KNOWLEDGE OF MEANING¹

ABSTRACT: In recent years the view that understanding a language requires knowing what its words and expressions mean has come under attack. One line of attack attempts to show that while knowledge can be undermined by Gettier-style counterexamples, language understanding cannot be. I consider this line of attack, particularly in the work of Pettit (2002) and Longworth (2008), and show it to be unpersuasive. I stress, however, that maintaining a link between language understanding and knowledge does not itself vindicate a cognitivist view of the former.

1. AN EPISTEMIC CONDITION ON LANGUAGE UNDERSTANDING (ECLU)

Language has phonological, grammatical, morphological, semantic and pragmatic components. Understanding a language is no less complex, and in what follows I'll restrict inquiry to its semantic dimension in the form of the question: What is involved in understanding the meaning of a word, phrase or sentence? Is such understanding underwritten by an exercise of propositional knowledge, by an exercise of practical ability, or instead is such understanding better conceived on the model of perception? One clarification our question immediately called for involves distinguishing between having the ability to understand a bit of

language such as a word, phrase or sentence, and realizing that ability in an act of understanding. Accordingly, in the "ability" construal of understanding, one understands a sentence in one's home language even if one has never encountered it before, either in thought or communication. By contrast, the "act" construal of understanding requires a dateable cognitive event of grasping that sentence. In the "ability" construal of understanding, then, I've understood the sentence 'My hovercraft is full of eels,' since about age five, whereas on the "act" construal of that notion I only understood that sentence upon first encountering it as a teen-ager enjoying British comedy.²

Abilities to understand play a large part in explaining acts of understanding. Just as my riding a bicycle on a given occasion is due in large part to my ability to do so and my choice to realize that ability on that occasion, so too my act-understanding a word, phrase, or sentence is due in large part to my ability to understand it and something activating that ability. Generally, however, what activates that ability is not my choice, but rather my being confronted with a word, phrase, or sentence: act-understanding such a thing is typically something that befalls me rather than something that I do. This is one of the observations encouraging an account of language understanding on the model of perception.

It is nevertheless difficult to resist the suggestion that act-understanding an uttered sentence requires as a necessary condition knowing what its constituent words mean. Knowing what its words mean is not a sufficient condition for understanding a sentence because one must also know how its constituent words are composed therein. Also, understanding a sentence is not in general sufficient for understanding what a speaker is doing in uttering it. 'Bob is on time,' means that Bob is on time and no more, but very often in uttering that sentence a speaker will be suggesting that there is something remarkable about Bob's punctuality. I shall assume in what follows that one can understand this sentence without knowing all of what its speaker, in any given utterance, was using it to convey.

In order to assess the above-mentioned connection between language understanding and knowledge, let's make more precise what has been called the Epistemic Condition on Language Understanding (Pettit (2002)). Where *S* is a speaker and α a word, phrase or sentence

that means M, then we have

ECLU: S act-understands α only if S knows that α means M.

I'll prescind here from complexities that would arise from cases in which α is indexical, ambiguous, or in some non-indexical way context-sensitive. Accommodating these phenomena would involve relativizing, within the definition of ECLU, to interpretations of α , contexts of utterance of α , and so on, but without a corresponding increase in clarity for our purposes.³ Also, in what follows I shall use 'understanding' to refer to act-understanding.

I state ECLU only as a necessary condition of language understanding because knowing that means M is not a sufficient condition for understanding α . After all, α might be a mile-long sentence which I'm told on good authority means that $0 \neq 1$. I thereby know what this sentence means, but it's doubtful that I understand it. Instead, understanding this sentence would require grasping its meaning in light of its compositional structure. Again, I might be told by a reliable source that a sentence S of a language with which I am quite unfamiliar, Hopi, say, means that some alligators are bigger than most footballs. I now know what S means, and indeed know that it means that some alligators are bigger than most footballs. Yet it's doubtful that I understand S. Or consider the sentence

1. No brain trauma is too minor to ignore.

It's easy to fall under the illusion of understanding (1), which at first blush appears to mean that one should pay attention to all brain traumas no matter how minor.⁴ However, closer inspection reveals that what (1) really says is that one should ignore all such traumas, no matter how minor. But someone might simply be unable to get that reading, and instead have to be told that (1) means that one should ignore all brain traumas. In this case, too, she knows what (1) means without understanding (1). The same goes for the well worn 'Buffalo buffalo Buffalo buffalo buffalo buffalo Buffalo buffalo'.

So I'll restrict myself to ECLU. In spite of being formulated as a necessary condition, the thesis appears to do some explanatory work. For instance, on the assumption that ECLU is true, we can start to make

sense of the fact that speakers can grasp both atomic and non-atomic expressions, as well as of the fact that they can interpret utterances of those expressions by others. This is not to say that the knowledge imputed by ECLU to speakers is conscious; perhaps it's tacit in the way that most people's knowledge of grammar is. Nevertheless, ECLU offers a link between semantics and psychology of a sort that enables us to illuminate speakers' competence.⁵

In spite of its credentials, in recent years ECLU has come in for criticism. One such criticism is in the form of a Gettier-style case in which a speaker seems to acquire understanding of a word's meaning in spite of not knowing what it means. In the next section I'll discuss the case and consider what those who have proffered it have inferred from it.

2. A GETTIER-STYLE COUNTEREXAMPLE TO ECLU

Pettit 2002 invites us to imagine a traveler in Germany whose German is moderately good, but who encounters an unfamiliar word, 'Krankenschwester.' He sees a kindly-looking, elderly gentleman sitting on a bench and asks him what this word means. The old man answers with an air of authority, "It means *nurse*." This is, of course, what the word means, and satisfied with the answer, the traveler thanks the old man and goes on his way. However, unbeknownst to the traveler, the old man is quite senile and doesn't know a word of English. In fact he says, 'It means *nurse*,' to any tourist that he encounters, and so it's only by sheer luck that he happened to say something true and relevant in this instance.

Pettit contends that in spite of this bizarre state of affairs, the traveler now understands what 'Krankenschwester' means. He writes,

As a result of this exchange, you are now able to use this previously unfamiliar word correctly and correctly interpret it as it is used by other speakers of German. If a German speaker assertively utters the sentence 'Die Krankenschwester ist nett,' for example, you will correctly take the speaker to be asserting that the nurse is nice. Or if you want to say in German that the nurse is coming, you will correctly express this thought with the sentence, 'Die

Krankenschwester kommt.’ In short, in a familiar sort of way, you have come to understand the word ‘Krankenschwester.’ (Ibid, pp. 519-20)

Pettit’s description and diagnosis of the case contain many strands, so let’s disentangle some of them. First of all, assuming for the moment that I’m the traveler in question, we are evidently being told that after hearing the old man’s answer and taking it at face value, I form the disposition to use ‘Krankenschwester’ in a certain way even before I start using it, or responding to its use by others. Further, Pettit seems to be claiming that by virtue of this disposition, I understand the word. That is, *immediately* after hearing the old man’s “answer” to my question about the unfamiliar word and taking that answer at face value, I understand the word.

This inference should be distinguished from another that we might at first blush discern in what Pettit says. For he might also be construed as suggesting that after just a few subsequent conversations with native German speakers that include a use of ‘Krankenschwester’, I, finding no resistance to my interpretation of this word as meaning ‘nurse’, understand that word. Those exchanges are, indeed, most likely sufficient to constitute my understanding of the word, just as they would most likely be sufficient to justify my belief about the meaning of the word if it wasn’t justified already. However, Pettit makes clear that he means to claim that I understand the word even before any such conversations take place: his claim is that immediately upon forming the disposition to use ‘Krankenschwester’ to mean ‘nurse’, I understand it. It is my *disposition* to use the word in a certain way that allegedly endows me with an understanding of it. Accordingly, Pettit would have to grant that the same conclusion would follow had I used a random interpretation generator to come up with a meaning for the word. So long as this process results in my forming a disposition to use the word properly, his view will be that I understand it even if I don’t know its meaning.

I stress this difference between having a disposition to use a word in a certain way, and using it that way based on that disposition, because casual conversation tends to collapse them. Even Pettit elides this distinction. At one point in his 2002 he is arguing for a special form of the thesis that one can understand a word without knowing its meaning, and for this purpose imagines that you are struck by lightning

in a public square in Munich. Before the lightning strike you were a monolingual English speaker. Now that the lightning has struck, your brain has been altered in such a way that you are disposed to use a great number of German words, and many rules of German grammar, properly. Pettit goes on:

You are not seriously hurt, but you are a bit dazed. A crowd of German speakers gathers around to check on your condition, and to your amazement you have the impression of being aware of what they are saying when they speak German, even though you previously didn’t understand a word. I’m not imagining that it sounds to you as though they are speaking English, as if you were watching a dubbed movie. You are aware that they are speaking German, but their speech seems intelligible to you in just the way English is intelligible to you. (2002, p. 35)

Pettit insists that in this case you understand German after the lightning strike. But the very experience of interpreting the words of the Germans surrounding you in a way that makes them come out as roughly rational is confirmation of your (perhaps otherwise unjustified) beliefs. You don’t interpret them as saying things like, “My hovercraft is full of eels,” or “We have lumps of it around back.” Rather, you interpret them as saying things like, “The poor Yank bastard looks singed but otherwise unhurt,” and “Don’t touch him! He might still carry a charge.” These latter two sentences are more or less rational things to say in light of the circumstances. But then, this experience provides evidence of the propriety of your interpretation of their words and grammatical categories. As such, it can’t be used to drive a wedge between understanding of language and knowledge of language.

In assessing Pettit’s case, then, we need to attend carefully to the distinction between having a disposition to use a word in a certain way and using it that way based on or in light of that disposition. With that warning in mind, we may now see that if Pettit’s characterization of the case is right, then one can understand a word without knowing what it means: The traveler does not know what ‘Krankenschwester’ means, Pettit contends, because he does not possess propositional knowledge that it means ‘Krankenschwester’. He does not possess such propositional knowledge because although he believes it means ‘nurse,’ and

that belief is true, his belief lacks justification—or at least justification of the sort required for knowledge. Nonetheless, Pettit contends, the traveler understands the word ‘Krankenschwester.’ On this basis, Pettit concludes that ECLU is not true.

3. ASSESSING THE COUNTEREXAMPLE

Responding to Pettit’s example, Weatherson and Sennett (ms) contend that in this case the traveler, whom they call Nogot, does in fact know what the word means. In support of this claim they ask us to imagine three similar cases:

The mixed room: Thirteen people are in the ballroom. Six of them are native German speakers. Six of them are monolingual English speakers. And the thirteenth is Mr. Nogot. How many people in the room know what “Krankenschwester” means in German? We think the intuitively plausible answer here is *seven*, not *six*.

I’m afraid I don’t see the force of this intuition at all. Of course, the example tries to play upon our tendency to distinguish Nogot from the monolingual English speakers. By hypothesis, Nogot doubtless knows more German than they do. What is more, he has beliefs about ‘Krankenschwester’ that the monolingual English speakers most likely lack, with the resultant tendency to use that word in a certain way. But what does that show? Surely it won’t imply that Nogot does in fact know the meaning of that word. He’s still in an inferior position to that of the German speakers with respect to that word. Let’s instead consider Weatherson’s and Sennett’s next case:

The homogenous room: As in the mixed room, except the six monolingual English speakers leave. Is it now true that everyone in the room knows what “Krankenschwester” means in German? Again, it seems very plausible to say *yes* here.

I confess I don’t see that this twist adds anything new. As before, Nogot is in an epistemically inferior position to the monolingual German speakers with respect to ‘Krankenschwester’. The fact that some

people have left the room strikes me as entirely irrelevant to our judgment of this situation. So let’s consider a third case offered by Weatherson and Sennett:

The bet: Herr Sieger and Herr Verlierer are placing bets on all sorts of things while they wile [sic] the day away. (Fortunately, they conduct their betting business in English.) As they see Mr. Nogot approach, Seiger says to Verlierer, “I bet that he knows what ‘Krankenschwester’ means in German.” and Verlierer accepts the bet. After some simple research, they discover all the facts about Nogot as described above. Who do you think would win the bet? We think that it’s very intuitive that Sieger wins, or at least should if the games are being fairly played.

Let’s assume that betting is carried out as fairly as the parties in question can ensure. The intuition that Nogot knows what the word means rests on the assumption that a true belief in this betting context is sufficient for knowledge. But surely that is not an intuition that most of us, or at least most reflective people, share. Obviously, the bettors in the case might agree, implicitly or explicitly, that ‘knowing’ shall here be a matter of true belief, just as evolutionary biologists are happy to talk about ‘deception’ as applying to tree frogs bearing bright warning coloration without in fact being poisonous. On that understanding, however, the bettors’ opinions about whether Nogot knows the meaning of Krankenschwester won’t settle any questions about what he knows. I assume, that is, that Weatherson and Sennett are not simply changing the subject. Rather, let’s agree that ‘know’ is not being used in a technical sense. On this understanding, the case adds nothing new to the authors’ cases considered already.

Strikingly, though, Weatherson’s and Sennett’s is an intuition that Pettit also harbors, albeit in a different way. For his reasoning in the quotation with which this paper began, moves from the correctness of the speaker’s use and interpretation of the word in question, to his understanding thereof. That is, as we have seen, Pettit points out that the speaker is disposed to use and interpret ‘Krankenschwester’ correctly as a result of his experience with the senile man. From this he infers that the speaker understands this word.

This should immediately give us pause. I might make a correct guess concerning the number of jelly beans in a jar. The correctness of the guess doesn't tend to show that I know the number of jelly beans in that jar even if I believe that my guess is correct. Getting lucky, even with the aid of the courage of my (unfounded) convictions, isn't sufficient for knowledge. Nor does the correctness of my guess make it the case that I understand how many jelly beans there are in the jar. That is in spite of the fact that I am, by virtue of my newly-formed belief, disposed to answer correctly questions such as, "How many beans are there in the jar?" and, "Are there more beans in the jar than there are quarters in your pocket?" Being disposed so to behave doesn't give me understanding in this case even if my dispositions to behavior are the same as if I had formed my belief on sufficient evidence.

This last point is easy to overlook and, once acknowledged, easy to forget, so let me re-state it. My lucky guess might endow me with a true belief that most of us will agree doesn't constitute knowledge. In so doing it will also, in many cases, endow me with a disposition or set of dispositions, both verbal and non-verbal, to behave as if I know the answer to the question how many jelly beans the jar contains. But that does not imply that I understand how many jelly beans there are in the jar. So a properly configured disposition doesn't give me understanding, at least in non-verbal cases.

So too, in the case that Pettit imagines, it's not controversial that the tourist got lucky in acquiring the correct definition of the word 'Krankenschwester.' Let us also agree, *pace* Weatherson and Sennett, that he does not, immediately after the exchange with the senile man, know the meaning of that word. Let us also agree that he can—that is, is disposed to—correctly use the word. I suggest that this does not show that he understands it.

Why not? One would have thought that being disposed to use a word correctly is at the very least a large part of understanding it. What more is there, one might ask, to mastery of a word than being disposed to use it properly in a broad and open-ended variety of contexts?

Let's consider the issue from a larger perspective. Words are a type of artifact, so I suggest we consider the question whether being disposed to use an artifact correctly is sufficient for understanding how to use it. It is not difficult to see that the answer is no. Suppose I

win a gadget in a raffle, but have no idea how to use it or for that matter what it's for. None of the raffle organizers can tell me anything about the gadget—perhaps explaining why there was a raffle in the first place. At any rate, I bring the thing home, open it up, and, flummoxed, decide as a sheer guess that I'll attach it to my road bike's power meter. *Voilà!* In fact, as I head out for a ride I find that this appears to be precisely what my prize is for, and it works marvelously there to add information to what was already being captured by the meter about my power output. Now I know not just my wattage but also my altitude, hill category, VO2 output, and efficiency! Here, then, are two questions:

Question 1: Did I know, after receiving the gadget but before even venturing a guess as to what it's for, what its function is?

Question 2: Did I understand, after receiving the gadget but before even venturing a guess as to what it's for, what its function is?

Surely the answers to both of these questions is No. So now imagine some time has passed and I've now made a completely random guess as to the gadget's function and on that basis have attached it to my bike's power meter.

Question 3: Did I know (after making my guess but before using it) what its function is?

Question 4: Did I understand (after making my guess but before using it) what its function is?

Here too, I submit that the answer to both questions is clearly in the negative. No doubt, we may grant that after a few tries with the gadget attached to my power meter, I know, as well as understand, what to do with it. That, however, implies nothing of interest about my epistemic state, or about my state of understanding, right after attaching the gadget. After all, following a few tries I get enough feedback to see that the gadget is working in a way that provides useful information for me as a rider, and since I realize that few tech gadgets have multiple uses, it's now a good bet that I've made a very lucky guess. Had I attached the

thing to my electric toothbrush, bread machine, or catalytic converter, nothing (or at least nothing good) would have happened.

In this case, then, I used the gadget correctly but neither knew how to use it nor understood its use simply on account of forming a belief about its use: I just got lucky. So too, in Pettit's case I used the word 'Krankenschwester' correctly after my encounter with the senile man. But that doesn't show that I understood the word before my disposition so to use it was confirmed by the acquiescence of others. For, no doubt, after a few exchanges with people using that word that confirm my belief about how it is to be used, I do come to understand it. If I had been misusing or misinterpreting uses of the word, someone would have corrected me; the fact that my uses and interpretive acts with the word didn't attract comment or generate puzzlement is itself enough to imbue my correct belief about the word's use with the justification required for knowledge. It is also enough for me to now be said to understand the word. Yet none of this shows that I understood the word immediately after hearing and accepting the definition given by the senile man.

Even if I am right that in the gadget example we don't find a cleavage between knowledge and understanding, this won't show that Pettit's diagnosis of his example is incorrect. However, it does show that there is no good inference from the fact that someone is disposed to use an artifact correctly, to the conclusion that he understands how to use it. We can also understand how this fact can be difficult to spot. Using something correctly is intuitively conceived as extended over time with many chances for self-monitoring and self-correction. Being *disposed* to use a thing correctly is, as a result, naturally thought of as sensitive to any need for modification or repair. Accordingly, being disposed to use something correctly goes hand in hand with competence. But the two are not logically tied, as the example of the gadget makes clear.

4. UNDERSTANDING AND KNOWING HOW

In the article in which his counterexample to ECLU appears, Pettit observes that in the original Krankenschwester case, the traveler acquires the same disposition to use the word that he would have acquired had he acquired that word's definition from a competent speaker. Accord-

ingly, Pettit argues, if the traveler comes to understand the word in the latter case, he comes to understand it when he hears the definition from the senile man.

The issue here is whether identity of dispositions to use a word implies identity in level of understanding of that word. We can readily see that it does not do so with a minor variation on the example of the bike gadget. Suppose that instead of my deciding what to do with it by virtue of a lucky guess, I attached the gadget to my power meter because of a small slip of paper on the gadget that said, "Attach this to your power meter." Now *this* case can bifurcate into two:

PowerMeter1: The slip of paper was put there in the factory at which the gadget was made, by technicians responsible for building it, and who correctly prescribe its use.

PowerMeter2: The slip of paper was put there by our same old senile man who also puts a slip with the same words on every gadget he can find.

Evidently I acquire the same disposition whether *PowerMeter1* or *PowerMeter2* obtains. But in the second case it is also clear that I don't understand what to do with the gadget. I form, on the basis of the slip of paper, a belief about what to do, and that belief is correct, but I only acquire understanding when I try it on and learn that it works well attached to the power meter. By contrast, it seems reasonable to suppose that I do understand what to do with the gadget in *PowerMeter1*. If this is correct, then it is not true that identity of dispositions implies identity in levels of understanding.

One might reply that understanding words is a different kind of affair from understanding artifacts. Although words are themselves artifacts, perhaps something special sets them apart from the larger genus. Perhaps what is special is this: understanding language is more like seeing than inferring. Just as, in important respects, we are passive recipients of the deliverances of our senses, so too, above a certain level of competence, grasping the meaning of a sentence—or phrase—token is not something we can help doing once we see or hear the sentence or phrase. This is why we can sometimes truly say that we couldn't help overhearing another person's conversation. It may also account for the phenomenon of illusions of understanding of the sort we encountered

before. Language understanding, that is, is perception—like whereas understanding the proper use of other artifacts does not appear in general to be.⁶

As a phenomenological point, the perception-like character of language understanding is indisputable.⁷ Observe, however, that the perception-like aspect of language understanding applies paradigmatically to results of compositional processes rather than to the semantic atoms with which those processes begin. If I know the meanings of the words out of which a phrase or sentence is composed and can discern its syntax, then my grasp of that phrase or sentence is very likely automatic and, accordingly, perception-like. The same may be true of complex words whose components we already grasp: if I already know the meanings of words like ‘driver’, ‘walker’ and ‘rider’, I can probably discern the meaning of ‘reader’ in a perception-like way even if I’ve never encountered that word before.

We may remain neutral on what precisely are the semantic atoms on which compositional processes operate: they may be words, or instead, more plausibly, morphological elements out of which those words are composed. Yet whatever are the relevant semantic atoms for compositional purposes, it should by now be plausible that our understanding of those items is no different in kind from understanding the function of other types of artifact such as tools. Scissors are for cutting, telescopes are for seeing distant objects with, and ‘Krankenschwester’ is for referring to nurses in speech acts. In all three of these cases, too, understanding the artifact’s use requires and is assured by knowing how properly to use it, and this in turn depends on more than just being disposed to use it properly.

We are now prepared to respond directly to Pettit’s challenge. He had claimed that if the traveler comes to understand ‘Krankenschwester’ on the basis of being told its meaning by a reliable source, then he also understands it when he is told its meaning by the senile man. We now see, however, that grasp of word meaning is more akin to grasp of matters of fact such as those pertaining to the behavior of artifacts generally, than it is to grasp of sentences or perceptible landscapes. As a result, two dispositions might be identical with regard to the sorts of (non-epistemically described) behavior they are prone to produce in an agent; yet one might be case of knowing how while the other is not.

Further, since we have established that a mere disposition to use an artifact in a certain way is not sufficient for knowledge of how to use it, it is not true that identity of dispositions to use a word in a certain way is sufficient for identity in level of understanding of its meaning. We can, for instance, imagine that a person has no idea of the meaning of ‘Krankenschwester’, but for amusement uses a random definition generator that attaches words to definitions, and, completely by coincidence tells you what is in fact the correct definition of this word. For some reason he believes what this interpretation generator tells him, and so forms a correct belief about this word’s definition. That does not endow him with knowledge how to use the word, nor does it endow him with understanding of that word. At the same time, a brief conversational exchange in which the word occurs may be all that is needed to endow him with know-how, and thereby understanding, of the word.

5. LONGWORTH’S FAÇADES

Longworth (2008) offers the following schematic example as a way of further challenging ECLU. In his discussion, a sentence façade is a string of words that appears to be grammatical but is in fact not. The schema then runs as follows:

Suppose that, in the midst of being presented with a barrage of utterances of sentence façades—again in a psycholinguistics lab—one is presented with an utterance that one parses veridically. Plausibly, a case of that sort can be constructed that corresponds with standard barn façade cases, so that one’s unreliability—or lack of safety—in the environment of the lab precludes one’s knowing what the utterance means on the basis of the veridical parse. If one would nonetheless count as understanding the utterance, then the case would stand as a counterexample to the propositional knowledge view. (2008, p. 64)

Given the examples of parsing that Longworth elsewhere provides, he apparently takes this process to be primarily a matter of constructing a semantic interpretation on the basis of a sentence’s grammar and

meaningful words. Accordingly, let's assume that the lab subject knows the meanings of and understands the component words of the sentence in question. Then parsing is, in the standard way, a matter of constructing an interpretation on the basis of those words' meanings and the sentence's syntax. In a veridical parse one interprets the sentence correctly by means of this process. In spite of this, there are a lot of other sentences hanging around in the lab that do not yield a grammatical interpretation.

If that is indeed the proper understanding of Longworth's case, then once again this example, albeit schematic, does not do the work it is intended to do. The reason is that one can see immediately and directly if one's interpretation of a sentence is correct. The 'Buffalo' sentence is an example of this: After some squinting, and some construction of analogous sentences in order to elicit a proper reading of this one, one simply sees that this sentence is perfectly grammatical. The presence of a lot of sentence fakes nearby does nothing to undermine one's knowledge of what it means, just as one can know the answer to a mathematical problem even when there are a lot of trick problems around.

This is not to say that one can never be mistaken in interpreting a sentence as meaning one thing when in fact it means something else. The brain trauma sentence, (1) and similar "verbal illusions" show that such things are possible. Rather, my contention is that *if* one understands a sentence (rather than just knowing its meaning), then one is in a position to know that one does—just as if I discern the answer to an arithmetical problem, I can also determine that I have done so.

6. CONCLUSION: KNOWING AND UNDERSTANDING LANGUAGE

Challenges to ECLU seem to be inspired by the intuition that an account of language understanding tied to knowledge over-intellectualizes it. The thought seems to be that what is crucial to our semantic competence is a capacity to use words in a certain way, and this capacity need not be underwritten by belief and justification. Yet knowledge and understanding cannot be separated in the way that this line of thought demands. The reason is that dispositions of the sort needed to underwrite understanding of language must also be competencies, which in

turn require a justificatory structure analogous to that of knowledge.

This conclusion does not, however, show that language understanding is best construed on the model of knowing that rather than that of knowing how. For as I argue elsewhere (Green (ms)) if it is plausible to reduce knowing how to knowing that, it is no less plausible to reduce knowing that to knowing how. In brief, an agent's knowing that P can be reduced to her knowing how correctly to answer the question whether P. Such know-how is not evinced in a correct guess or even in a plausible conjecture. What's more, an agent who lacks the use of her tongue and vocal chords because of, say, Locked-in Syndrome, is still able to answer questions if only in the privacy of her own thoughts.

A reduction of knowing that to knowing how will do nothing to show the latter is more fundamental than the former, or that knowing that is in some sense specious and that all that "really" exists is knowing how. (Either conclusion would betray a muddle about the nature of reductive claims.) Rather, if there is a moral to draw it is that a speaker's grasp of a language is at once cognitive and practical. ECLU is but one manifestation of this multi-dimensional nature of our grasp of language.

Notes

¹Earlier versions of this paper were presented at the Meaning, Understanding and Knowledge Conference, Riga, Latvia, August, 2009; and at the University of Missouri in St. Louis, September, 2009. I am grateful to audiences on those occasions for their perceptive comments, as well as to Corin Fox and Barry Smith for comments on an earlier draft of this paper. Research was supported in part by a Sesquicentennial Fellowship from the University of Virginia, and by a Visiting Fellowship from the Virginia Foundation for the Humanities. Both sources of support are gratefully acknowledged here.

²Hunter 1998 discusses this distinction in more detail.

³Pettit (2002) formulates a similar thesis in terms not just of necessary but also sufficient conditions. As I'll explain presently, the sufficient condition is a non-starter.

⁴The example is inspired by a similar one from Longworth (2008).

⁵Pettit (2002) explains these links in more detail.

⁶McDowell (1998, p. 99) eloquently defends the perception-like character of language understanding.

⁷See, however, Smith 2009 for a challenge to an over-ambitious interpretation of this phenomenological observation.

References

- Franklin, R. L. 1981. 'Knowledge, Belief and Understanding'. *The Philosophical Quarterly* 31, no. 124: 193–208.
- . 1983. 'On Understanding'. *Philosophy and Phenomenological Research* 43, no. 3: 307–328.
- Green, M. ms. 'Knowledge that reduces to knowledge how.'
- Gross, S. 2005. 'Linguistic Understanding and Belief'. *Mind* 114: 61–66.
- Heck, R. 1995. 'The Sense of Communication'. *Mind* 104: 79–106.
- Hunter, D. 1998. 'Understanding and Belief'. *Philosophy and Phenomenological Research* 58: 559–80.
- Larson, R. & Segal, G. 1995. *Knowledge of Meaning*. MIT.
- Longworth, G. 2008. 'Linguistic Understanding and Knowledge'. *Noûs* 42: 50–79.
- McDowell, J. 1998. *Meaning, Knowledge and Reality*. Harvard.
- Pettit, P. 2002. 'Why Knowledge is Unnecessary for Understanding Language'. *Mind* 111: 519–50.
- . 2005. 'Belief and Understanding: A Rejoinder to Gross'. *Mind* 114: 67–74.
- Smith, B. 1979. 'Speech Sounds and the Direct Meeting of Minds'. In M. Nudds & C. O'Callaghan (eds.) 'Sounds and Perception: New Philosophical Essays', 183–210. Oxford.
- Wason, P. C. & Reich, S. S. 1979. 'Belief and Understanding: A Rejoinder to Gross'. *Quarterly Journal of Experimental Psychology* 31: 591–7.
- Weatherston, B. & Sennett, A. ms. 'Knowing and Understanding: Reply to Pettit'. Available at: <http://www.brian.weatherston.net/papers.html>.