

Michael Devitt, *Ignorance of Language*.
Oxford: Oxford University Press, 2006. xiii + 320 pp.

Michael Devitt's book makes a starting assumption that generative linguistics has been a very successful scientific program, and Devitt's general aim is to do philosophy of generative linguistics in the sense that one might do philosophy of physics or philosophy of biology—indeed, as one would do the philosophy of any branch of science. Accordingly, Devitt purports to be interested in the proper interpretation of theory and practice in generative linguistics and not to be prescriptive about how generative linguists should conduct inquiry.

Quite properly, Devitt does not think we should automatically assume that linguists correctly interpret their own work. Just as philosophers of physics needn't defer to Einstein on the proper analysis of relativity theory, there is no reason why philosophers of linguistics should automatically defer to Chomsky. Accordingly there is nothing wrong with the fact that Devitt takes issue with Chomsky on several core claims about the nature of linguistic theory.

For example, Chomsky and many other linguists take generative linguistics to be engaged in the study of a cognitive faculty that underwrites our linguistic competence. They often maintain that linguistic competence involves our knowing certain rules or principles that determine whether linguistic structures are well formed or “legible” and how those structures are interpreted.

Devitt disagrees. On his view, individuals don't know or follow linguistic rules—at least not in the sense that the rules constitute knowledge—nor in the sense that there are rules that govern linguistic processing; rather, the sense in which we have linguistic rules is in the sense of there being rules that describe the products of an *ability* or *skill* that we have. Moreover, Devitt “tentatively” proposes that the skill is not part of a task-specific cognitive module, but is rather determined by constraints imposed by our need to map linguistic forms into the language of thought. Less tentatively, Devitt proposes that the syntactic forms studied by linguists are not properties of mental representations, but are abstract relational properties of external linguistic objects (utterances and inscriptions).

As a corollary Devitt has much to say about topics ranging from the nature of linguistic intuitions (they are not the voice of competence but are rather judgments grounded in higher-level processing) to questions about the conventional nature of language.

The problem with Devitt's work is threefold. First, a key argument he offers for his positive proposal begins with a premise that he labels “uncontroversial” but which in point of fact appears to be rejected by most generative

Thanks are due to Michael Devitt for helpful and lively comments on several previous iterations of this review. This is not to say that he is happy with the result. Thanks are also due to the editors of the *Philosophical Review* for comments on matters of both substance and style.

linguists, and the ensuing argumentative structure is often elusive. Second, while Devitt purports to be offering a proposal that is faithful to linguistic practice, the range of linguistic phenomena and explanations he surveys is limited. This limited survey helps lead him badly off target. Third, while Devitt cautions us against basing our interpretation of linguistic theory on the written proclamations of linguists, he often seems to base his own account of linguistic practice on written statements by linguists (and philosophers) that, in my opinion, are often taken out of context. Let's consider these concerns in order.

1. On the Soundness of Devitt's Arguments

When Devitt defends his positive thesis on the relation between language and thought, he offers an argument that I found difficult to track. He begins with what he declares to be a "theory-neutral" assumption: "The most theory-neutral view of competence in a spoken language comes with position (M), the minimal position on the psychological reality of language. It is the view that this competence is the ability to produce and understand sentences with the sounds and meanings of that language. . . . This view is so neutral that even an eliminativist about thoughts could adopt it" (128).

On the initial (and natural) reading, this sounds like something that linguists would reject outright since the whole point of the competence/performance distinction is that linguistic theory is not a theory of comprehension and production. But apparently this isn't what Devitt intends. For Devitt, the "ability to produce and understand sentences with the sounds and meanings of that language" in his sense should not be taken as a theory of comprehension and production, but rather as a theory of a *state* that is *utilized* by the processes involved in comprehension and production. Devitt considers this proposal to be theory neutral in that even a Chomskyan could accept it: for the Chomskyan the ability would include the knowledge of grammar—possibly knowledge of a set of propositional rules.

In my view this is not a very happy use of the term 'ability'. By itself knowledge of linguistic rules leaves one far short of the ability to produce or understand anything. One also needs things like ears and tongues and mouths and vocal chords, or at least eyes and hands and all the motor and perceptual processes driving them, not to mention a parser (which is usually, if not always, considered outside of the theory of competence), dedicated memory, and so forth. On Devitt's view a state of knowing linguistic rules counts as "an ability to produce and understand sentences" because the relevant comprehension/production processes would *utilize* the knowledge state. The problem is that those processes utilize many other states as well, from the relevant memory states and the computational state of the parser all the way down to physical and biological states of the perceptual/articulatory system as well. Do we want to say that those states are also part of our linguistic competence? If not, then the definition of competence

needs to be sharpened to exclude them. If the definition of competence includes such states, then there is nothing theory neutral about Devitt's definition of competence.

Another concern here is that people could know or cognize linguistic rules but be unable to speak or comprehend due to physical damage. Devitt (section 6.3) thinks this is a red herring; of course we have abilities that cannot be causally engaged—consider someone who has the ability to ride a bicycle but cannot do so because of a broken leg. The problem for Devitt is that linguists have hypothesized that competence could be *completely* causally inert—indeed some linguists (for example, Chomsky) have entertained the possibility of the grammar being a kind of mutation that was intact in human ancestors thousands of years before it was utilized in speech perception and comprehension.

But the problem goes deeper than the question of whether knowledge of linguistic rules could be causally inert. Devitt's formulation of competence makes it clear that the state is *individuated* by or partly constituted by its role in linguistic production and comprehension. Or as Devitt puts it (23) “the theory of competence begins with a theory of its outputs. A competence is a competence to produce outputs with certain natures; those natures are partly constitutive of the competence.” Even if we broaden this to include outputs that *would be* produced if the grammar were connected to a well-functioning perceptual/articulatory system, this is not a theory-neutral view of competence. Many linguists think that inputs and outputs (actual or possible) play no role in how the grammar and its underlying states are individuated.

Devitt's characterization of competence is not benign. A big part of Devitt's pitch is the idea that linguists are not studying a species of knowledge that, but rather knowledge how. If you define competence so that it now means something like having an ability to produce and understand linguistic expressions, the knowledge-how thesis is arguably already baked into the cake. Maybe the knowledge-how thesis is correct, but it is no defense of the thesis to pack the conclusion into a key premise and declare the premise “theory-neutral” (or, for that matter, the “most theory-neutral” option).

The problem with Devitt's argument is not limited to his premise, which, as I noted, I consider to be question begging. The subsequent argumentative structure is often difficult to track, and insofar as it is trackable, it appears to rest on a number of invalid steps.

For example, Devitt segues into a discussion of the doctrine (this one merely declared “*relatively* uncontroversial”) that language expresses thought (LET). LET and M (the view that this competence is the ability to produce and understand sentences with the sounds and meanings of that language) now come together in the following way:

Our relatively uncontroversial assumption [LET] leads to a more theory-laden view: that competence is the ability to use a sound of the language

to express a thought with the meaning that the sound has in the language in the context of utterance; and the ability to assign to a sound a thought with the meaning that the sound has in the language in the context of utterance. In brief, the competence is the ability that matches token sounds and thoughts for meaning. (128–29)

I gather that the strategy is this: Take “uncontroversial assumption” LET, add the “most theory-neutral” view M, and we are allegedly led to this:

LET+M: Competence is the ability to use a sound of the language to express a thought with the meaning that the sound has in the language in the context of utterance.

The question is, what does Devitt mean when he says that his “uncontroversial” assumption “leads to” LET+M? It certainly isn’t a valid inference, given the intensional aspects of terms like ‘ability’. I might intend to (be able to) say *X*, and *X* might express *Y*, but it doesn’t follow that I intend to (am able to) use *X* to express *Y*.

Is there another sense of “leads to” that is weaker than logical or metaphysical entailment but stronger than “leads Devitt to suppose”? Perhaps there is an abductive argument in the offing, but we are not given a hint of what it might be.

So far in Devitt’s argument we have a premise that is labeled “uncontroversial” (but which in my view is very controversial), followed by a second claim that allegedly “follows” from this premise and LET (but doesn’t follow logically or in any other obvious way). The next step is no less troubling. “LET [Language Expresses Thought] has led to the view that linguistic competence is the ability to match sounds and thoughts for meaning. If this is right then it is immediately apparent that any theory of linguistic competence, and of the processes of language comprehension and production, should be heavily influenced by our view of the nature of thoughts” (129).

Of course far from being “immediately apparent,” the step is obviously fallacious. In the first place we may have little or no handle on the nature of thoughts, in which case it would seem imprudent to shackle our theory of linguistic comprehension to whatever it is we think we know about the nature of thoughts. But second, even if we had robust knowledge about the nature of thoughts, the principle is in general bad advice. I might be interested in how birds navigate by celestial bodies, and correlatively by their ability to detect and represent these bodies, but it does not follow that my theory of bird navigation should be heavily influenced by our view of the nature of celestial bodies—for example, that they are massive distant bodies in which nuclear fusion is taking place.

Devitt isn’t finished however. From the contention that our view of comprehension and production should be “heavily influenced” by our view of the

nature of thoughts, he draws yet another conclusion: “So, the psychological reality of language should be investigated from a perspective on thought” (129).

There is a difference between having a theory “influenced by” x and our theory being “investigated from a perspective on x .” Watson and Crick’s theory of the human genome was influenced by factors in quantum chemistry, but we wouldn’t say that they investigated the human genome from the perspective of quantum chemistry.

This is just a case study in a kind of disease that infects the book. Premises are declared “uncontroversial” or “theory neutral” (and competing assumptions are declared “implausible”) and then are linked together with argumentative steps that supposedly “come with” or “follow from” or are “led to” or are “immediately apparent” from the allegedly benign premises. But the premises are not benign and the alleged argument is but a simulacrum of a real argument—it borrows the vocabulary of logical reasoning to walk us through a series of claims that Devitt considers platitudinous but do not follow and may well be false. And here we get to the issue of Devitt’s limited survey of linguistic phenomena and explanatory mechanisms.

2. On Linguistic Phenomena and Explanatory Mechanisms

As I noted earlier, Devitt doesn’t object to what linguists do, he just objects to the standard interpretation of linguistics on which it is construed as a theory of the psychological reality of language. On Devitt’s view, the theory of grammar is not a theory of psychological reality, it is rather a theory of *linguistic* reality. On his view, studying the psychological reality of language makes perfect sense, but that would require studying it from the perspective of thought (thus the argument we saw in the previous section). There are thus two different enterprises, one that is concerned with the psychological reality of language and one (linguistics) that is not. The question is why would Devitt think that the kinds of structures and mechanisms posited by linguists would not be psychological. Here, *prima facie*, the chips seem to be stacked against him.

For example, one of the demands that the language faculty apparently puts on question formation in natural language is that it satisfy what Ross (1967) called *island constraints*. Here is one illustration:

1. *Who(m) did John hear the story that Bill hit ____?
2. Who(m) did John hear that Fred said that Bill hit ____?

Example (1) is pretty robustly bad, but (2) seems like a well-formed question. Notice that it is perfectly clear what a sentence like (1) would mean if it was well-formed—we can express it in “Quinese” like this: Who is such that John heard the story that Bill hit him/her? There doesn’t seem to be a problem about (1) not being apt for thought, or being in conflict with our conceptual competence. It is just an odd fact about natural languages that structures like this are

ill-formed. The explanations for this fact—sometimes called “subjacency”—have varied over the last forty years, but the basic phenomenon is that one can’t move a WH element (question word starting with “wh”) over both an S (sentential clause) and an NP (noun phrase) node unless there are intermediate landing sites, as in (2).

Facts like this are typical of what syntacticians have studied for the last forty years. They study these quirky linguistic phenomena and attempt to deduce why the language faculty gives rise to them. It does not appear that assumptions about the communicative function of language or the aptness for expressing thought offer any hope of explanation here; hence syntacticians look elsewhere for the answers.

Linguists are also obviously interested in differences between languages. For example, we might be interested in why, in languages like English, the reflexive pronoun ‘himself’ can be associated with ‘John’ but not ‘Bill’, while in languages like Swedish the reflexive can be associated with ‘Bill’ in the following sentence:

3. Bill said that John likes himself.

Baker (2001) provides a useful summary of current work investigating these questions and describes how the parametric model of the grammar can explain a vast array of linguistic variation of this type by positing a handful of binary parameters. This parametric explanation does not intersect at any point with the role of language in communication or thought, nor is it even clear how it might do so.

The problem with Devitt’s interpretation of what linguists are up to (an interpretation on which linguistic properties are abstract relational properties of the externalia and are determined in part by convention) is that it seems hard to imagine that the acoustical signals at whatever level of abstraction (and whatever role convention might play) carry all of the information about the properties of interest to linguists. To begin with, it is difficult to even detect word boundaries from an acoustical signal, but the representations posited by linguists consist in part in a rich array of inaudible structures and forms that interact in robust ways. A good example involves the analysis of sentences like ‘John wants a unicorn’ as initially proposed by Quine, introduced within linguistic theory by McCawley (1974) and developed by den Dikken, Larson, and Ludlow (1996). The basic idea is that we can account for a number of facts, ranging from quantifier scope ambiguity to adverb attachment ambiguity, if we posit an implicit clause, and correlative an implicit (phonologically unrealized) subject PRO and a phonologically unrealized predicate HAVE:

4. John wants [PRO (HAVE) a unicorn].

The reading in which John has a particular unicorn in mind corresponds to the case where the quantified expression ‘a unicorn’ takes wide scope over the entire

sentence, as in (5), leaving behind a coindexed “trace” *e*:

5. [a unicorn]_{*i*} John wants [PRO (HAVE) *e_i*].

The result is an entire clause that has no phonological content. Nevertheless all of the inaudible syntactic content still interacts in robust ways with the rest of the grammar. For example, the trace must be properly governed, the PRO must satisfy the PRO theorem, and so forth. This isn’t a particularly rare phenomenon in linguistics—phonologically unrealized structure is rather the norm in generative linguistics.

While in subsequent papers Devitt (2006, 2008) has attempted to address the problem of PRO (not successfully, in my view), it is completely obscure how all of the structure in an example like (5) (and the relations between its components and those of the voiced structure) could be a feature of the acoustical signals that we produce or could be established by convention (consider the difficulty in our following simple audible conventions like when to say ‘whom’ instead of ‘who’). In example (5) the lower clause corresponds to absolutely nothing in the sound wave that is produced—not even a pause. Whether ultimately right or wrong, it at least makes sense for linguists to think of these structures as being psychological objects—data structures, for example.

3. On Devitt’s Reading of Others

Earlier I mentioned that Devitt is not particularly charitable in his interpretation of what linguists have to say about the nature of their enterprise. A good illustration of this problem can be found in Devitt’s discussion of linguistic intuitions (chap. 7). The problem begins with his definition of linguistic intuitions: “We should start by clarifying what we mean by ‘linguistic intuitions.’ We mean fairly immediate unreflexive judgments about the syntactic and semantic properties of linguistic expressions, metalinguistic judgments about acceptability, grammaticality, ambiguity, coreference/binding, and the like” (95).

This is an odd place to begin. Linguists typically do not claim to have judgments of grammaticality and certainly not of binding facts. Rather, they claim that we have judgments of acceptability and (in some cases) possible interpretations of linguistic forms. These judgments provide evidence for linguistic phenomena (like binding), and the theory of grammar in turn explains these phenomena and ultimately our judgments. In other words, we don’t have judgments about theoretical linguistic phenomena. Those theoretical phenomena are discovered by sophisticated higher-level theorizing—this is what linguists are for.

Devitt compounds the error by supposing that linguists take these judgments of acceptability to involve some inner “voice of competence.” Devitt is also quick to give it a label that, for anti-Cartesians like me, is maximally negative: “I need a word for such special access to facts. I shall call it ‘Cartesian’” (96). The problem is this: who would suppose that judgments of acceptability involve

special Cartesian access to the voice of competence? Devitt argues that many linguists do hold this position, and in support of his claim, he offers a string of quotes (some from nonlinguists), none of which speak to the point in question (96). Consider:

It seems reasonably clear, both in principle and in many specific cases, how unconscious knowledge issues in conscious knowledge . . . it follows by computations similar to straight deduction. (Chomsky 1986, 270)

We cognize the system of mentally represented rules from which [linguistic] facts follow. (Chomsky 1980, 9)

We can use intuitions to confirm grammars because grammars are internally represented and actually contribute to the etiology of the speaker/hearer's intuitive judgments. (Fodor 1981, 200–201)

[A speaker's judgments about the grammatical properties of sentences are the result of] a tacit deduction from tacitly known principles. (Graves et al. 1973, 325)

Our ability to make linguistic judgments clearly follows from our knowing the languages that we know. (Larson and Segal 1995, 10).

Note that not one of these quotes talks about special access or Cartesianism or “the voice of competence.” Consider the quote from Larson and Segal, for example. Far from talking about an inner voice of competence, they even eschew the term ‘intuition’ for the less loaded ‘linguistic judgments’. Moreover, is there any thoughtful person who could possibly disagree with this quote? Surely the linguistic judgments that I make follow from my knowing the language that I know. If, for example, I knew Japanese, I would have very different judgments about my language.

The second quote from Chomsky doesn't even speak to linguistic judgments, but rather to linguistic facts. There is a difference. Here is one way of understanding the difference: linguistic data (intuitions or judgments) provide evidence for phenomena (like binding facts or “island effects”) that are explained by the theory of grammar. The picture is something like this:

Theory ==> Linguistic Facts <== Data (for example, judgments)
 (explains/predicts) (are evidence for)

When Chomsky says that we cognize a system of mentally represented linguistic rules from which linguistic facts follow, he is talking about the relation between the theory of grammar and linguistic phenomena. Data (and judgments) are not even under discussion in that quote.

Of course there is a path from the grammar to the judgments I have, just as there is a path from what Chomsky writes to what Devitt reads and judges Chomsky to be saying, but the fact that Chomsky's writing “contributes to the

etiology” (Fodor’s phrase) of what Devitt judges Chomsky to be saying does not mean that Devitt has Cartesian access to Chomsky’s intended meaning; likewise for judgments of acceptability. Those judgments would be very different if the grammar was different, but that does not mean we have special Cartesian access to the grammar via our judgments, and certainly not that our judgments are always correct.

4. Conclusion

There have been few book-length treatments of the philosophy of generative linguistics, and Devitt deserves credit for attempting to contribute to this important topic. There are weaknesses in his effort, however, ranging from the soundness of the arguments offered, to the fidelity of the readings of others. Apart from these weaknesses, there is the equally serious problem that the work does not successfully reflect the richness of the phenomena and explanations linguists traffic in. The errors, individually and in concert, seem to play key roles in Devitt’s unflinching and at times desperate attempt to lead us to the conclusion that linguistics is not a chapter in cognitive psychology. I don’t object to the conclusion per se; it is certainly a serious candidate position in the move space. Here my objection is with the limited range of linguistic practice that Devitt examines as well as with the many questionable steps and assumptions Devitt enlists to try and persuade us of his conclusion. A much stronger effort is required to secure the conclusion he is aiming for.

References

- Baker, M. 2001. *The Atoms of Language*. New York: Basic Books.
- Chomsky, N. 1980. “Rules and Representations.” *Behavioral and Brain Sciences* 3: 1–14.
- . 1986. *Knowledge of Language*. New York: Praeger.
- den Dikken, Marcel, Richard K. Larson, and Peter Ludlow. 1996. “Intensional Transitive Verbs.” *Rivista di Linguistica* 8: 312–48. Abridged version reprinted in *Readings in the Philosophy of Language*, ed. P. Ludlow, 1041–54. Cambridge, MA: MIT Press, 1997.
- Devitt, M. 2006. “Defending Ignorance of Language: Responses to the Dubrovnik Papers.” *Croatian Journal of Philosophy* 6: 571–606.
- . 2008. “A Response to Collins’s Note on Conventions and Unvoiced Syntax.” *Croatian Journal of Philosophy* 8: 249–55.
- Fodor, J. 1981. “Introduction: Some Notes on What Linguists Are Talking About.” In *Readings in the Philosophy of Psychology*, ed. N. Block, vol. 2, 197–207. Cambridge, MA: Harvard University Press.
- Graves, C., J. J. Katz, Y. Nishiyama, S. Soames, R. Stecker, and P. Tovey. 1973. “Tacit Knowledge.” *Journal of Philosophy* 70: 318–30.
- Larson, R., and G. Segal. 1995. *Knowledge of Meaning*. Cambridge, MA: MIT Press.
- McCawley, J. 1974. “On Identifying the Remains of Deceased Clauses.” In *Language Research* 9: 73–85. Reprinted in *Adverbs, Vowels, and Other Objects of Wonder*, 84–95. Chicago: University of Chicago Press, 1979.

Ross, J. 1967. "Constraints on Variables in Syntax." PhD diss., Massachusetts Institute of Technology.

Peter Ludlow

Northwestern University

Philosophical Review, Vol. 118, No. 3, 2009

DOI 10.1215/00318108-2009-009

Owen Flanagan, *The Really Hard Problem: Meaning in a Material World*.
Cambridge, MA: MIT Press, 2007. 288 pp.

If, as the lazy and oversimplified philosophical tradition has it, you can't derive *ought* from *is*, what is a good naturalist to do about ethics? Does morality just have to hang there in the air, underivable and insupportable, at best the creature of some sort of "existential" choice that can only be justified internally, within one's chosen moral compass? Or can there be a sort of scientific investigation of what is the best way to live one's life *all things considered*? Can meaning, in the grand sense of the meaning of life, not just the semantics of one language or another, be found—and confirmed—in the natural world, and if so, how? That is what Owen Flanagan thinks is the "Really Hard Problem," and he proposes to solve it, or at least to sketch out the best path to its solution, in an exercise of *eudaimonics*, "the attempt to say something interesting and systematic about what makes for human flourishing and that gives life meaning—that is, if anything does" (xii).

This is well-trodden territory, of course, but mainly explored by amateur, not professional, philosophers: people who have thought hard—but not "rigorously"—about the issues, often from the perspective of some religious tradition. Perhaps for that very reason most professional philosophers have shunned the topic as too ambitious, too naively formulated, to be tractable. We should be grateful to Flanagan for bucking that trend, for he conducts his inquiry with erudition, calm open-mindedness, cautious optimism, and ingenuity. Flanagan's choice of '*eudaimonia*' as his term for the most important form of human flourishing—allowing him to acknowledge that some people are quite *happy* in a familiar sense while not yet conceding that their lives are all they could or should be—signals his willingness to draw heavily on the philosophical tradition going back to Aristotle, but he is eclectic, drawing heavily on the thinking of the Buddha (primarily as interpreted by the Dalai Lama) and also on such contemporary moral thinkers as Amartya Sen, John Rawls, and Martha Nussbaum, and