

Could Competent Speakers Really Be Ignorant of Their Language?¹

ROBERT J. MATTHEWS
*Department of Philosophy,
Rutgers University*

This paper defends the commonsense conception of linguistic competence according to which linguistic competence involves propositional knowledge of language. More specifically, the paper defends three propositions challenged by Devitt in his Ignorance of Language. First, Chomskian linguists were right to embrace this commonsense conception of linguistic competence. Second, the grammars that these linguists propose make a substantive claim about the computational processes that are presumed to constitute a speaker's linguistic competence. Third, Chomskian linguistics is indeed a subfield of psychology, in the business of characterizing the linguistic competence of speakers.

Introduction

Commonsense, at least in the mouths of English speakers, has it that competent speakers *know* the language that they speak: They do not simply *know how* to speak the language that they speak; they actually *know* that language in a way that goes beyond practical know-how. Many linguists and philosophers of language sympathetic to the broadly Chomskyan program of generative linguistics have embraced, at least until very recently, these deliveries of (English) commonsense. They have held that (i) knowledge of language is distinct from knowing how to speak a language, (ii) this knowledge constitutes a speaker's linguistic competence, (iii) this knowledge is propositional knowledge of a grammar for the language, and (iv) speakers use this knowledge in the course of language production and understanding.² And because it is the task of linguistics to characterize this knowledge that constitutes a speaker's linguistic competence, linguistics is therefore a subfield of

¹ A version of this paper was presented at the Philosophy of Linguistics Conference, Dubrovnik, September, 2005.

² See, e.g., Chomsky [1980], [1986].

psychology, at least a subfield of commonsense propositional attitude psychology.

These Chomskyan elaborations of (English) commonsense have always been the source of considerable controversy. Especially controversial have been the claims that competent speakers have propositional knowledge of a grammar for their language and that such knowledge is constitutive of their linguistic competence. Critics have challenged both claims. Epistemologists, for example, have long complained that the knowledge of grammar that Chomskyans take to be constitutive of linguistic competence fails to satisfy a number of the standard criteria for being genuine knowledge (it isn't available to consciousness, it lacks the requisite epistemic warrant; etc.).³ A number of critics have argued that knowledge of grammar, or more specifically, knowledge of semantic theory, is not required for linguistic competence; indeed some have argued linguistic competence involves no knowledge whatever beyond practical know-how. As Devitt [2006] puts it, a competent speaker might be *completely ignorant* of his language. Thus, for example, proponents of translationalist semantics such as Schiffer [1987] and Fodor [1990] have argued that semantic competence is not a matter of knowing anything at all about one's language, certainly not a matter of knowing a semantic theory for that language. It is simply a matter of mastering a translation procedure that maps heard utterances in a public language into sentences in one's language of thought, sentences that express beliefs about what a speaker said in uttering what he did. Schiffer describes the mastery of such a procedure in terms of the acquisition of a non-intentionally described partial recursive function.⁴ Proponents of practical know-how accounts of linguistic competence have also been critical of intentional (and hence epistemic) conceptions of linguistic competence. Hornsby [2005], for example, is especially critical of the notion that "ordinary users of language have knowledge, deployed in understanding, of propositions that semantic theorists articulate" (109). She thinks that it is a "devastating" objection to point out that "in order for someone to know the theorist's propositions, they would need intellectual resources which are simply not needed for using their language" (109). Linguistic competence, she insists, is a matter of practical know-how.

Even Chomsky [2000] himself, after decades of seeming to defend an epistemic conception of linguistic competence, has become sharply critical of intentional (and hence epistemic) conceptions of linguistic competence, at least as such conceptions might figure in a scientific theory of linguistic competence. He has argued variously over the last few years that (i) the so-called 'representational' states to which linguistic theories advert are not genuinely representational in the sense of being about something represented (Chomsky [2000], 159), (ii) linguistic

³ See, e.g., Barber [2001], Pettit [2002].

⁴ For discussion, see Matthews [2003].

knowledge is not constitutive of linguistic competence, or at least that concepts such as 'knowledge of language' or 'knowledge of grammar' have no role to play in scientific linguistic inquiry (Chomsky & Stemmer [1999], 397), and (iii) intentional attribution has no explanatory role to play in cognitive science (Chomsky [2000], 23). Speaking of Marr's theory of vision, which has often been presented by philosophers as evidence of the intentional character of computational cognitive science, Chomsky argues, "The theory itself has no place for the [intentional] concepts that enter into the informal presentation, intended for general motivation" ([2000], 161).⁵

Clearly the commonsense epistemic (and hence intentional) conception of linguistic competence, at least as that conception has been elaborated within Chomskyan linguistics, has fallen on hard times. I am sympathetic to Chomsky's idea that intentional attribution has no role to play in a naturalistic scientific theory of linguistic competence, other than the informal presentational and motivational role that he readily acknowledges. I am also sympathetic to the idea that the knowledge constitutive of linguistic competence cannot be knowledge of a grammar or semantic theory, for many of the reasons adduced by critics of what I have been calling the Chomskyan elaborations of the commonsense epistemic conception of linguistic competence. But I nonetheless think that the commonsense epistemic conception is correct: propositional knowledge of language *is* constitutive of linguistic competence, with the consequence that it is simply false that a competent speaker could, as Devitt and others have claimed, be *completely ignorant* of his language, in the sense of having no such knowledge. What a competent speaker minimally has to know, is the pairing of sounds and meanings that his or her language effects. A competent speaker has to know this pairing because otherwise we would find ourselves unable to justify certain basic truths about language use, e.g., that in the usual situation a speaker *knows* what he or she has said and *knows* what another with whom he or she is conversing has said, truths that seem fundamental to any account of our rational reliance on language as a reliable means of communication. In making a case for the commonsense epistemic conception of linguistic competence, I shall focus on Devitt's argument to the contrary, an argument that he takes also to establish that linguistics is *not* a subfield of psychology.

Devitt's Ignorance of Language

In his *Ignorance of Language* Devitt argues for three basic theses:

- i. A competent speaker could be *completely ignorant* of his language (in the sense of not having any propositional knowledge of that language); indeed, he could fail to have any propositional attitude whatever towards his language—linguistic competence is rather a matter of practical know-how.

⁵ For a discussion of this matter, see Egan [1995].

- i. Linguistics is not a subfield of psychology concerned to characterize the linguistic competence of speakers—linguistics is rather concerned with the products of linguistic competence.
- iii. The grammars attributed by linguists to speakers, to the extent that they correctly characterize the ‘structure rules of language’, impose only the following minimal adequacy condition (M) on theories of those speakers’ linguistic competence, and only to that extent can grammars claim to be psychologically real:

(M) A competence in a language, and the processing rules that govern its exercise, respect the structure rules of language: the processing rules of language comprehension [take] sentences of the language as inputs; the processing rules of language production yield sentences as outputs.

Against Devitt I want to argue that (1) it is empirically unlikely that any theory of a speaker’s linguistic competence, i.e., a grammar, will satisfy condition (M); nevertheless, (2) grammars do in fact make a substantive empirical psychological claim about the computational processes that presumably constitute a speaker’s linguistic competence; moreover, (3) competent speakers of a language do have propositional knowledge of their language, at least they have knowledge-that of their language, and that knowledge is constitutive of their linguistic competence. Let me begin with (1) and (2), which concern the empirical import of linguistic theory, specifically of the grammars proposed by linguistics, for computational psycholinguistic theorizing. I will then turn to (3), which directly concerns the epistemological question of whether, as Devitt claims, a competent speaker might be *completely ignorant* of language, in the sense of having no knowledge of language other than a practical knowledge of how to speak this language.

Devitt, we noted, argues that that the grammars attributed by linguists to speakers enjoy only a minimal psychological reality: to the extent that they correctly characterize what he calls the ‘structure rules of language’, grammars impose the above-mentioned ‘minimal’ adequacy condition (M) on theories of that speaker’s linguistic competence, and to that extent they can claim to be psychologically real. Put another way, the most that can be claimed of a grammar by way of psychological reality is that the processing rules that govern exercises of a speaker’s linguistic competence *respect* the grammar, in the sense that such exercises successfully generate or interpret linguistic products that conform to the structure rules of language which Devitt takes the grammar to characterize. In effect, grammars are compact (finite) characterizations of the linguistic objects that a linguistically competent speaker is competent in producing and understanding, and as such grammars tell us only what sort of linguistic object the competence is a competence for. To borrow Devitt’s analogy, grammars tell us as much about linguistic competence of speakers as a theory of horseshoes tells us about the blacksmithing competence of the blacksmiths that produce them.

I agree with much that Devitt has to say about the various positions that he canvases regarding the psychological reality of grammars. All

must be rejected—some on empirical grounds, some on grounds of computational intractability, and still others on grounds that they rest on conception confusion (e.g., thinking that grammatical rules are processing rules). The simple fact is that most philosophers and many linguists are mistaken about the sense in which grammars are psychological hypotheses.⁶ I also agree with Devitt that linguistics, as actually practiced, is concerned primarily with the products of linguistic competence and only indirectly with the competence itself, though I would hasten to add that if one is interested in linguistic competence, and some linguists clearly are, I see no other way of studying it at this point in time. But I do *not* endorse Devitt’s condition (M) as the minimal positional on the psychological reality of grammars that even someone who rejects each of the five stronger positions on psychological reality should accept. It seems to me quite unlikely that any grammar of the sort that might be proposed by linguists using current linguistic methodology would satisfy (M), and not simply because the grammars currently being proposed are not plausibly construed as rule sets (they are better construed as constraint sets, which specify conditions on the pairing of sounds and meanings effected by a language). The requirement of (M) that competence in a language, and the processing rules that govern its exercise, respect the structure rules of the language is clearly too strong, and it is probably too weak as well.

The requirement is too strong inasmuch as if by ‘the structure rules of the language’ one means the sort of grammar that a linguist might propose, viz., a compact rule/principle set that captures in a perspicuous manner the sort of structural relations that linguists find theoretically significant, then it seems very unlikely that such a grammar will satisfy (M). The problem here is not simply that linguists are prepared to count as grammatical garden path sentences such as *the canoe floated down the river sank* or multiply center-embedded sentences such as *sailors sailors sailors fight fight fight*, despite most speakers’ inability to understand these sentences without considerable priming or coaching. There is also the obvious fact that English speakers are often able to understand perfectly well sentences and expressions of English speakers that would receive *no* interpretation on any grammar of English that linguists would be likely to hypothesize, e.g., Davidson’s (1985) ‘a nice derangement of epitaphs’. Moreover, the languages defined by the grammars hypothesized by linguists are often, if not typically, computationally intractable,⁷ such that whatever the competence of speakers is a competence for, it is not a competence for understanding languages defined by these grammars, since certain sentences of the languages so defined cannot be parsed efficiently by *any* computational procedure, even in some cases over the fragment of the language that speakers

⁶ See Matthews [1991].

⁷ For a discussion of this problem as it afflicts EST-era grammars, see Barton et al. [1987].

might actually use. Or if the languages defined by these grammars are computationally tractable, they are not biologically plausible, in the sense that the computational processes required to parse them do not seem to be biologically plausible given the computational demands that they impose.

The requirement of (M) is probably too weak as well inasmuch as computational psycholinguists working on language processing can (and do) distinguish between processing models (and acquisition models) that implement the grammars of one or another kind of linguistic theory (Minimalism, P&P, EST, LFG, GPSG, etc.), where the distinctions that they draw do not turn on the fragment of the language (the set of sentences) that each model can process, but rather on the nature of the processes hypothesized. But if different processing models can be so distinguished, then it is plausible to assume that empirical data could be brought to bear on the question of which kind of model provides the best characterization of a speaker's linguistic competence, thus showing that (M) is probably too weak.

So it seems simply wrong to think that any grammar of the sort that linguists might propose could turn out to be psychologically real in the sense given by Devitt's (M), viz., that the linguistic competence respects that grammar. But this, so far as I can see, entails nothing whatever as regards whether grammars are psychologically real, i.e., whether they are psychological hypotheses, whether linguistics is a subfield of psychology? Grammars might be psychological hypotheses about linguistic competence, despite the fact that they are invariably (and perhaps, given linguistic methodology, necessarily) false. Indeed, they might be quite fruitful hypotheses, despite being false. Whether grammars are psychological hypotheses is to be decided, not by determining whether grammars bear some privileged implementation relation to language processing models, e.g., satisfying Devitt's condition (M), but by investigating their *actual* role in psycholinguistic theorizing. And when one investigates their role in such theorizing, one discovers that grammars are standardly taken as *specifying intensionally the pairing that the language effects and that speakers compute in the course of language processing*, albeit only under (significant) idealization and approximation. The procedure that specifies this pairing is not, of course, either of the functions that the speaker computes in the course of language production and language understanding. Those functions are mappings from sounds to meanings or vice-versa, whereas what the grammar specifies is a function that has as its range not either sounds or meanings but sound-meaning pairs, i.e., specifically the pairing that the given language effects and a speaker of that language computes. Thus, for example, within the Chomskyan minimalist program, a grammar for a language specifies a function from lexical items into sound-meaning pairs by means of a recursive procedure called 'Merge'. Such a grammar is a specification of the speaker's competence, and not a model of performance, precisely because the function specified intensionally by

this grammar is not the function actually computed by a speaker of that language in the course of language processing, but the one function does specify the other (again, under idealization and approximation). As Chomsky has emphasized time and again, the grammar specifies *what* speakers do in the course of language use, viz., compute the specified pairings, but not at all *how* they do it. The grammar is not a processing model, not even a component of such a model. Computational theorists regard the grammar in just the way that one might expect, viz., as specifying the pairing of sounds and meanings that their processing model must largely respect, modulo considerations of computational tractability, at least over the fragment of sentences that competent speakers actually encounter.

Once one begins to think of grammars in this way, and this arguably is just the way that Chomsky and computational psycholinguists do, it seems quite clear that grammars *are* empirical hypotheses about psychological processes, albeit hypotheses articulated at a level of significant abstraction, inasmuch as there is no claim as to how speakers manage to compute the pairing specified by the grammar.

An obvious objection to my proposed construal of the claim that grammars are psychological hypotheses, and hence that linguistics is a subfield of psychology, would be to point out, as Devitt [2006] does, that in constructing linguistic theories, linguists are largely unresponsive to the sorts of psycholinguistic data to which one might expect them to be sensitive, if grammars were indeed psychological hypotheses about language processes. But this objection is mistaken. Most of the psychological evidence to which it is often claimed linguists should be sensitive is simply not of the right sort: at best it bears on the computational procedures by which speakers compute the function that they do, rather than on the function computed. In effect, this evidence is at the wrong level of description. Of course, there might be evidence to which linguists should be, but aren't, responsive, perhaps about the computational complexity of the functions that the grammars specify. But even here matters are not that straightforward: it is extremely difficult to know just how to bring such complexity considerations to bear on linguistic theories. Linguists, however, have become quite responsive to learnability considerations, *which* are articulated at the level of the function computed, a responsiveness that to my mind goes a long way towards establishing their claim to be doing psychology. If they were, as Devitt claims, just providing a compact characterization of the products of linguistic competence, learnability considerations would be irrelevant.

The important question is not whether grammars are psychological hypotheses, they clearly are, but whether they are methodologically fruitful. Certainly the computational models of language processing developed thus far have drawn heavily on existing linguistic theory, so much so that it is commonplace in the field to describe these models as 'implementations' of one or another linguistic theory. Only time will tell whether this research strategy leads to empirically plausible theories of

language processing. Even if grammar-based computational modeling turns out not to be very fruitful in developing a computational theory of language processing, I suspect that the large amount of computational mucking around that they have inspired will turn out to have been very fruitful in giving us some idea about just what, computationally speaking, linguistic competence amounts. Other empirical hypotheses in psychology should do so well!

But Could Competent Speakers Really Be Ignorant of Their Language?

If grammars are specifications of the pairing that languages effect and that speakers of those languages compute in the course of language processing, then grammars are clearly psychological hypotheses, and linguistics, which is in the business of constructing grammars, is a subfield of psychology. And linguistics is such, it should be noted, irrespective of whether linguistics trades in intentional attribution. (It is this fact, incidentally, that explains how Chomsky is able to hold onto his claim that linguistics is a subfield of psychology while abandoning an intentional conception of linguistic competence.) But what about Devitt's central claim that competent speakers could be completely ignorant of their language, in the sense of lacking any knowledge of their language except for practical knowledge how to speak the language? Devitt argues that while speakers clearly have considerable linguistic know-how, most are largely ignorant of their language; they don't have knowledge-that of their language, specifically they don't have knowledge-that of a grammar for their language. They lack such knowledge, Devitt claims, because they lack what he takes to be essential for knowledge or indeed for any other propositional attitude, namely, a mental representation of that knowledge. Specifically, they lack a semantically contentful mental representation of the language, i.e., a mentally represented grammar for the language, which could play the right sort of causal computational role in language processing. Devitt thinks that we can be quite certain that competent speakers don't have such representations of their language, or at least don't need them, from the fact that computational theories of language processing don't attribute to speakers representations of grammars of the sort proposed by linguists.

In thus arguing from the lack of such representations to the conclusion that competent speakers could be completely ignorant of language, Devitt presumes the so-called Representational Theory of Mind,⁸ according to which someone possesses an attitude A towards some proposition p (e.g., knows that p) just in case that individual has a mental representation that has the propositional content p and which plays the appropriate causal role in the possessor's psychological economy. The problem here, as Devitt sees it, has nothing to do with the usual sorts

⁸ See, e.g., Fodor [1987].

of objections raised against the commonsense epistemic conception of linguistic competence. It has rather to do with the lack of the requisite mental representations, which Devitt thinks can be established on both conceptual and empirical grounds. Devitt's argument is simply this: if there are no representations, there are no propositional attitudes. And if there are no propositional attitudes, then there is a fortiori no propositional knowledge, i.e., no knowing-that, and no cognizing-that either, for that matter. But there are no representations; hence, there is no knowledge-that.

Georges Rey [2005] shares Devitt's conviction that knowledge requires (semantically contentful) representation of what's known, but unlike Devitt he is not willing simply to abandon Chomskyan claims to the effect that speakers know a grammar for their language. At least he is not willing to abandon claims to the effect that linguistic competence is a matter of knowing something about one's language, because to do so would be tantamount to abandoning the supposition that language competence is to be explained in intentional terms. Thus, Rey feels compelled to challenge Chomsky's claim ([2000], 159) that the linguistic 'representations' that speakers have and use lack semantic and intentional properties.

Not being a fan of the Representational Theory of Mind (what Rey calls the 'Computational Representational Theory of Thought', CRTT for short), or indeed of any representational account of the attitudes, I am not at all inclined to abandon the commonsense intuition that competent speakers really do know their language for the sorts of reasons that Devitt offers. If speakers should turn out to lack contentful mental representations that are plausibly interpretable as the contents of what commonsense takes speakers to know about their languages, then so much the worse for the Representational Theory of Mind. The case for knowledge being constitutive of linguistic competence is ever so much stronger than is the case for the Representational Theory of Mind. Two considerations seem especially relevant. First, it seems that we cannot explain our rational reliance on language as a reliable means of communication if we can't attribute to competent speakers certain knowledge about their language. Language is a reliable means of communication precisely because we know what we and others are saying when we utter the things that we do. But we can't be said to *know* what we or others have said in uttering the sentences that we do unless we know *inter alia* the pairing of sounds and meanings that we effect in speaking the language that we do. In other words, competent speakers know what they and others have said in uttering the sentences that they do in virtue of knowing this pairing, and knowing this pairing is precisely what makes competent speakers competent. The second consideration has simply to do with how commonsensically we understand competence: competent individuals just are individuals who are knowledgeable about whatever the competence is a competence for, and it is this knowledge that is constitutive of the competence in question. So to the

extent that we are concerned with a commonsense conception of linguistic competence, which is surely the sort of conception that the epistemic conception is, then knowledge *is* constitutive of linguistic competence. The knowledge in question counts as *propositional*, or at least an instance of knowledge-that, simply because when we endeavor to say with any precision just what this knowledge constitutive of the competence is, we find ourselves forced to the knowledge-that locution, not because we are therefore picking out some special kind of knowledge, viz., propositional as opposed to practical, but rather because that particular locution is the only one available to us. Thus, for example, when we want to characterize what the child who knows how to add knows, we find ourselves forced to say that this child knows, e.g., that when the numbers to be added in a column sum to greater than 9, then the tens digit must be carried into the next column of numbers. And similarly when we undertake to characterize in commonsense terms the competence of a competent speaker of English, we find ourselves forced to say that this speaker knows, e.g., that reflexive pronouns must be bound in their governing category, though here again it is not clear that we are attributing to the speaker some special kind of knowledge, e.g., 'propositional' as opposed to 'practical' knowledge. Such competence characterizations, we need to remind ourselves, are, as Chomsky says, 'informal'; they are just our commonsense way of characterizing competences.

The conclusion that there can be knowledge without representation, at least without explicit representations of the sort over which computations are often defined, should not be surprising: Much cognitive scientific theorizing presumes this to be the case, and some cognitive scientists (most notably Newell [1981]) have explicitly argued that knowledge attributions (what Newell called 'knowledge-level' descriptions) entail nothing as regards computational architecture. But if there can be genuine knowledge of grammar without any mental representation of what's known, then there is the following question: How then are we to explain our conviction that states of knowledge (e.g., of a language or a grammar), and propositional attitudes more generally, have semantic and intentional properties? Some may feel themselves forced into a representationalist construal of the attitudes, even to the point of embracing a view according to which the knowledge states that constitute speakers' linguistic competence are representations of essentially nonexistent entities (see Rey [2005]). I myself don't see any insurmountable difficulty here that should force us into a representationalist construal of the attitudes. The empirical support for the computational models that cast doubt on representationalist construals of the attitudes is arguably much stronger than whatever (minimal) theoretical support these representationalist construals may enjoy, especially given the fundamental unclarity of the crucial notion of representation that these construals employ. And in any event, there are plausible non-representationalist construals of the attitudes that can easily accommodate the commonsense epistemic conception of linguistic competence even in the

absence of representations of the sort that representationalists presume such knowledge to require.⁹ But this is a matter for another time. The point I want to emphasize here is the Devitt's claim that competent speakers might be completely ignorant of their language is on weak ground indeed if its only justification is the lack of computationally explicit representations that could express the propositional content of such knowledge. It's going to take a more direct argument both for the claim and for this claim that linguistic competence is simply a matter of practical know-how.

References

- Barber, A. [2001], "Idiolectal Error," *Mind & Language*, 16, 263–283.
- Barton, G., R. Berwick, and E. Ristad (eds.) [1987], *Computational Complexity and Natural Language* (Cambridge, MA: MIT Press).
- Chomsky, N. [1980], *Rules and Representations* (New York, Columbia University Press).
- _____[1986], *Knowledge of Language* (New York: Praeger).
- _____[2000], *New Horizons in the Study of Language and Mind* (Cambridge: Cambridge University Press).
- _____ & Stemmer, B. [1999], "An On-line Interview with Noam Chomsky: On the Nature of Pragmatics and Related Issues," *Brain and Language*, 68, 393–401.
- Davidson, D. [1986], "A Nice Derangement of Epitaphs," in R. Grandy and R. Warner (eds.), *Philosophical Grounds of Rationality* (Oxford: Clarendon Press), 157–174.
- Devitt, M. [2006], *Ignorance of Language* (London: Oxford University Press).
- Egan, F. [1995], "Computation and Content," *Philosophical Review*, 104, 181–203.
- Fodor, J. [1987], *Psychosemantics* (Cambridge, MA: MIT Press).
- _____[1990], "Review of Stephen Schiffer's *Remnants of Meaning*," in *A Theory of Content and Other Essays* (Cambridge, MA: MIT Press).
- Hornsby, J. [2005], "Semantic Knowledge and Practical Knowledge," *Supplement to the Proceedings of the Aristotelian Society*, 79, 107–130.
- Matthews, R. [1991], "Psychological Reality of Grammars," in A. Kasher (ed.), *The Chomskyan Turn* (London: Basil Blackwell), 182–199.
- _____[forthcoming], *The Measure of Mind* (London: Oxford University Press).
- Newell, A. [1981], "The Knowledge Level," *AI Magazine* (Summer Issue), 1–20.
- Pettit, D. [2002], "Why Knowledge is Unnecessary for Understanding Language," *Mind*, 111, 519–550.
- Rey, G. [2006], "Conventions, Intuitions and Linguistic Inexistents: A Reply to Devitt". *This volume*.
- Schiffer, S. [1987], *Remnants of Meaning* (Cambridge, MA: MIT Press).

⁹ See Matthews (forthcoming).