

Why We Still Need Knowledge of Language¹

BARRY C SMITH
*School of Philosophy, Birkbeck College,
University of London*

In his latest book, Michael Devitt rejects Chomsky's mentalist conception of linguistics. The case against Chomsky is based on two principal claims. First, that we can separate the study of linguistic competence from the study of its outputs: only the latter belongs to linguistic inquiry. Second, Chomsky's account of a speaker's competence as consisting in the mental representation of rules of a grammar for his language is mistaken. I shall argue, first, that Devitt fails to make a case for separating the study of outputs from the study of competence, and second, that Devitt mis-characterises Chomsky's account of competence, and so his objections miss their target. Chomsky's own views come close to a denial that speakers have knowledge of their language. But a satisfactory account of what speakers are able to do will need to ascribe them linguistic knowledge that they use to speak and understand. I shall explore a conception of speaker's knowledge of language that confirms Chomsky's mentalist view of linguistics but which is immune to Devitt's criticisms.

We tend to take the speech of the Chinese as inarticulate gurgling.
Someone who understands Chinese will recognize language in what he hears.

(Wittgenstein, *Culture and Value*, p. 1)

What enables us to recognize language in the sounds we hear? Crucially, what matters is our knowledge of language: knowledge of which sequences of phonemes are words, which word strings are grammatical,

¹ My thinking on these topics was influenced by invaluable discussions with Noam Chomsky, John Collins, Ophelia Deroy, Dunja Jutrović, Guy Longworth, Peter Ludlow, Robert Matthews, Stephen Neale, Paul Pietroski, Georges Rey and of course, Michael Devitt, whose consummate good nature and good humour have kept things lively and positive through the most robust of exchanges. He has put his case vigorously and in detail in a challenging book and should be thoroughly commended for doing so. My thanks to all, and my special thanks to Michael Devitt for providing several opportunities in Dubrovnik and in print for this exchange.

and what these words and sentences mean to us. It is this knowledge of language that Noam Chomsky takes to be the subject matter of linguistics, and this makes linguistics a branch of psychology:

Linguistics is simply that part of psychology that is concerned with one specific class of steady states, the cognitive structures employed in speaking and understanding (Chomsky [1975], 160)

Michael Devitt disagrees. According to Devitt, Chomsky and his followers have misconceived their task and are mistaken about the whole enterprise of linguistics. In a striking passage from his [2006] book, *Ignorance of Language*, Devitt declares:

I shall argue, a person could be competent in a language without representing it or knowing anything about it: she could be totally *ignorant* of it. (*Ignorance of Language*, 5)

He offers two main reasons for this bold claim. First, we do not need to study competence to engage in linguistic inquiry: the real subject matter of linguistics is the output of linguistic competence. Second, those who, like Chomsky, mistakenly take linguistics to be concerned with the psychological reality of a speaker have an implausible conception of linguistic competence as consisting in propositional knowledge of the rules of a grammar. In Devitt's alternative vision, the linguist theorises about the outputs of the speaker's competence by producing a grammar for the sentences speakers produce. Where Chomsky is mistaken, according to Devitt, is that he goes on to attribute to the speaker an unconscious knowledge of the rules of the language the linguist comes up with. This further move is unwarranted, and unnecessary, according to Devitt, since speculation about the nature of linguistic competence is a worthy subject but one that plays no part in linguistics proper. The linguist can study aspects of linguistic reality without referring to any facet of speakers' competence. Linguistics is simply the study of the symbols or expressions of a language.

Much of Devitt's book is taken up with the critique of the competence view he ascribes to Chomsky, but space is given early on to his attempt to separate the study of linguistic outputs from the study of the competence that produces them. I shall argue that the attempt to divorce the theory of linguistic outputs from the theory of linguistic competence fails. And although Devitt is right to reject the specific claim that speakers have *propositional* knowledge of the rules of their language, the rejection of this claim misses its target when aimed at Chomsky. A proper conception of the speaker's knowledge of language is immune to Devitt's criticisms, and an account of the extent and nature of this knowledge is still needed if we are to explain the linguistic properties of natural languages.

1. *Devitt's Rejection of Linguistic Mentalism*

The generalisations about language that linguists are interested in concern the syntactic structure of language: the structure of complex

expressions. Natural languages are essentially structured, and an account of their syntactic structure plays a vital part in any theory of language. However, the precise nature and location of syntactic properties is far from clear. We know that they cannot be aligned with the acoustic properties of speech or the orthographic properties of inscriptions. The syntax of a language does not reside on the surface of speech or writing. The real issue is where these essential linguistic properties *are* to be found. What sort of facts are linguists describing when they specify the syntactic properties of a language? Our answer to this question should have consequences for how we conceive of language and its place in nature.

The trouble is that people often begin with an a priori conception of what a language is and struggle to accommodate syntactic facts, or even the range of evidence linguists draw upon in forming syntactic hypotheses. We have had linguists and philosophers of language who were behaviourists, some who have been Platonists, and, of course, there are mentalists, like Chomsky. Behaviourism fared very badly by restricting the type of evidence that could count in favour of a linguistic hypothesis. Platonism left unclear what relation there is between languages conceived as independently existing abstract objects and the evidence we can elicit from speakers.² By contrast, Chomsky's view is that the real subject of linguistic inquiry is the syntactic structure of language – the essential properties and relations that make language *language*. And in bringing syntactic facts into focus he concludes that much of our pre-theoretical notion of language is simply irrelevant to the scientific study of language. The reality of language consists in the underlying facts that give form and character to the linguistic items we perceive ourselves and others to be producing. Chomsky's mentalism consists in his taking these underlying facts to be in the mind/brains of speaker/hearers.

Devitt keenly contests this view. He is strongly attached to the folk conception of a language, while at the same time ready to accept the findings of generative grammar with its postulations of syntactic structure. At first, this appears an unworkable option since the descriptions generative grammarians give of syntactic structure depart radically from the recognisable descriptions of grammar given by the folk. However, Devitt thinks of the linguist as theorising about the observable bits of language we ordinarily encounter in order to arrive at hypotheses about the syntactic properties possessed by those ordinary items of language. Two questions arise: (a) What *is* his conception of language as the observable evidence for linguistics?; and, (b) What aspects of reality are linguists describing in their detailed hypotheses about the syntactic structure of expressions?

² The Platonist also faces the problem of defining the actual language relation between speakers and their languages. See Smith [2006a] for a discussion of this problem.

The answer to (a) is the folk notion of language. The folk notion sees languages as out there, existing independently of us; a language is something all around us that we gradually become immersed in and avail ourselves of in order to express and communicate our thoughts. The philosophical interest, and the philosophical difficulty, is to see whether this notion of language makes sense, and if not, how it should be revised.

The cunning answer to (b) is the claim that linguistics is about linguistic reality, and that the hypothesised syntactic properties of ordinary language are part of that linguistic reality. This is fine as far as it goes. But it is precisely the nature of that reality which is at issue. We cannot just assume linguistic reality as a given any more than we can assume mathematical or moral reality as a given in a world of matter and causes. The task is to understand the true nature of linguistic forms and to understand how they can be accommodated in the natural world. We need a proper conception of linguistic reality in order to understand the ultimate nature of language. Is it, for example, a psychological matter of what is represented in the mind of speakers? Is it a matter of communal practices? Is it a set of conventions in a community? Is it a set of abstract objects existing independently of all speakers? All these are candidates and it would be surprising if those who study language and its linguistics properties did not have a view about the nature of language and linguistic properties. Somehow Devitt both wants to postpone discussion of this issue and attack Chomsky's mentalist stance.

Devitt's case against the mentalist conception of linguistics begins by denying that the study of linguistics is the study of a speaker's competence. Instead it is taken to be the study of the outputs of that competence, where:

The theory of outputs has a certain epistemic and explanatory priority over the theory of competence (ibid. 2006)

The claim of independent accessibility to linguistic outputs matters because it is taken to support large conclusions about the nature of language and the objects of linguistic inquiry. It is the basis for rejecting Chomsky's claim that linguistics is the study of certain states of competent speakers—states of the mind/brain of the speaker—and for promoting Devitt's alternative vision of linguistics as the study of linguistic reality.

But what priority do outputs have? And what sort of independent access can we have to them? Can we establish facts about a language independently of facts about its speakers? Devitt seems to imagine we can, making the easy assumption that what we are talking about is just there anyway, knowable without knowing anything about the speakers who produce it. Of course, were there no competent speakers there would be no language, as Devitt readily admits: 'language is what the competence produces' (31). Language exists because we do: because creatures with minds like ours produce it and because it can be acquired from other such minded creatures that continue to use it. But it is as if, once we have

put bits of language into the air, so to speak, they stay there and take on a life of their own, rendering them independent of any, and perhaps all, speakers' minds, whatever dependence they had on minds to begin with. Competence is required to produce the outputs that linguistics studies, so Devitt is not a Platonist. But once the outputs have been produced they become readily available linguistic tokens and a source of evidence for the linguist, who need not allude to, or know anything about, the nature of the linguistic competence that produced them. The outputs which serve as the evidence for linguists can be observed and studied independently of speakers' psychology.

Accounts of this kind involve what Chomsky calls an E-language—that is, a language:

Externalised...in the sense that the construct is understood independently of the properties of the mind/brain ([1986], 20)

They can be studied and understood independently the mind brain, even though languages are not independent of speakers' mind/brains *in origin*. For Devitt, the study of competence must go via the study of languages: 'whenever there is a linguistic competence there *has* to be such a language...the language is what the speaker is competent in' (ibid. 31), and by studying the properties of the language we bring to light the object our linguistic competence has to grapple with. The idea seems to be that competence gets to grips with an already existing language—the language of our community, perhaps. Thus competence cannot constitute the linguistic data, as Chomsky thinks.

By my lights this is a strange metaphysical picture. What direct or unfiltered access do we have to linguistic reality? How do we select linguistically significant items for scrutiny without ourselves being competent speakers who perceive speech sounds as meaningful? The linguistic facts are not simply unproblematically accessible. Where exactly do we find facts about the complex arrangements of expressions standing in relations of, say, c-command to one another? Not *in* the sounds, that's for sure. But here, of course, one is meant to be contented with the uninformative platitude that in linguistics we are describing linguistic reality. Such a response immediately invites further questions. What does linguistic reality consist in, what kinds of things make up that reality, and in virtue of what does that language have the linguistic properties it has? However, we are prevented from asking these questions because, according to Devitt, they are *further* questions that need not concern the linguist. This is surprising, for surely methodological issues ought to be of concern both to the philosopher of language and the linguist. Surely, linguists are entitled to ask what the objects of inquiry are, and to expect an informative answer. After all, philosophers would hardly settle for being told in their discipline that ethics is the study of ethical reality, aesthetics the study of aesthetic reality, or, for that matter, that religious study is the study of religious reality? The nature of ethical and aesthetical properties is a highly contentious issue, and an important

topic in philosophy. And a similar set of issues arises for the nature of language and linguistic properties. Questions about the status of these properties cannot be evaded by an uninformative platitude.

However, Devitt has another way with the methodological issues. When push comes to shove he acknowledges that the syntactic properties of a language may be ‘largely determined by the mind/brain’ (26) but he thinks there is room to resist the view that linguistics is about the mental even if the linguistic properties of symbols are *determined* by the mind/brain. For:

Even if symbols had their properties in virtue of certain mental facts that would not make the theory of those symbols about those facts and so would not make the theory part of psychology. (ibid. 40)

Why is this? Well, to think a theory was *about* certain facts just because they determine its domain of inquiry would be like accepting the following general conclusion:

Every theory—economic, psychological, biological, etc.—would be about physical facts and part of physics because physical facts ultimately determine everything. (Ibid. 40)

But this trivialising response is wholly inadequate. It may be true that real entities have the properties they do in virtue of the layout and arrangement of all the physical facts there are. Nevertheless, generalisations in the special sciences concern different levels of organisation of the physical world and have their own proprietary levels of description of the phenomena. The special sciences are about groupings of entities the laws of those special sciences apply to. As Devitt himself acknowledges: ‘A special science does not lose its own domain because that domain supervenes on another’ (ibid. 40). That’s right, and it is still a good question for each special science what entities its domain encompasses. It turned out, for example, that chemistry could be reduced to physics, but so far as we know, the generalisations in biology cannot be captured without loss by the laws of physics. And this is not because biology is about non-physical stuff. Rather, it is because the phenomena biologists are interested in, and over which they generalise, lack integrity or unifying features at the level of description physicists address. Still, we can ask whether biology concerns properties of entities in the natural world, or whether its findings concern some mysterious *élan vital*.³ Similarly, we can ask whether linguistics concerns physical properties of noise emitted by speakers, mental representations of linguistic structure, or something else. This remains a good question.

We can and should ask what the generalisations of linguistics are about, [and] what facts make them true, when true. We need to address these questions in order to contest Devitt’s own speculations about the nature of linguistic phenomena. For in spite of his injunction not to ask *in virtue of what* questions, Devitt cannot refrain from giving his

³ I take it none of us thinks it is the latter.

own answer, in terms of spatio-temporal, physical entities. Linguistic entities, or symbols of the sort linguists study, are, according to Devitt, ‘parts of the spatio-temporal physical world’. But just which parts of the physical world are they? He goes on, ‘our theory can abstract ...from a range of properties of the outputs—for example, form of script and pitch of sound—focusing simply on *the syntactic properties* we are interested in’ (24, italics mine). The problem should by now be clear. Linguists are concerned with the syntactic properties of symbols. Symbols are spatio-temporal parts of physical reality. But syntactic properties are not to be found among the spatio-temporal parts of physical reality, as Devitt knows. Devitt is admirably clear about the importance of the theoretical entities posited in syntax:

A triumph of generative grammar has been to make us appreciate how much of the syntax of a sentence is not explicit, “how unrevealing surface structure may be as to underlying deep structure”; “surface similarities may hide underlying distinctions of a fundamental nature” (Chomsky 1965: 24). Consider empty categories, for example: “An empty category...is a constituent that has no phonological substance associated with it; it is inaudible in speech, and invisible in the standard orthography” (J. D. Fodor 1989: 156). Consider also the scope of quantifiers. This scope is often not explicit on the surface and yet is revealed at the level of logical form or LF, the level of syntactic structure that most concerns us (1.3)...Considerations of this sort lead Stephen Neale to go further: “rather than saying that (30) [‘Every poet respects some sculptor’] is an ambiguous sentence, really we should say that (30) is the surface representation of two distinct sentences...that share an S-Structure representation and in fact look and sound alike” (1993: 119).’ (ibid. 154)

Syntactic properties—the ones we are most interested in—should present difficulties for Devitt’s conception of linguistic entities as external, physical tokens. And yet he remains utterly blithe:

Syntactic properties are ones that reflect a token’s relations to other tokens in the language; they are functional properties and *extrinsic* to the representation. Although formally so different, a written and spoken token of (1) might share all their syntactic properties. Sentences that “look different” can be syntactically alike. Formal differences are one thing, syntactic differences, another. (ibid. 154)

Why is Devitt so optimistic that syntax can be easily accommodated by a brute physical view of language? Well it is because syntactic properties are not themselves intrinsic, or brute-physical properties of physical tokens. They are ‘extrinsic’, ‘non-obvious’, ‘abstract’, ‘higher-order’, or ‘relational’ properties; and properties of physical sound tokens, presumably. Yet despite this admission, Devitt claims that linguists study ‘physical sentence tokens’. The use of ‘sentence’ in the context of physical tokens is strange given that Devitt is willing to concede that a sentence is something with a syntactic structure: a hierarchical not a linear structure. Much of that structure is inexplicit: it may contain relations between displaced elements and phonetically null elements not detectable in the physical world. And while syntactic properties and relations

can be properties of public strings, or uttered sounds, they are not *public* properties of those strings or sounds. They are not observable or intrinsic properties of the speech sounds speakers produce. Nothing like PRO, trace or VP-ellipsis are realised in the surface string. Given that Devitt accepts the findings and generalisations of syntactic theory that make appeal to such entities necessary, his talk of ‘physical sentence tokens’ is at best misleading. And so too is his claim that linguists study tokens. The very idea that it is tokens—unrepeatable, dateable, physical events—that they study is so implausible that Devitt seeks to cover over this inconvenient fact by retreating to ‘as if’ talk:

It is often convenient to talk of the objects posited by these [linguistic] theories as if they were types not tokens, as if they were abstract Platonic objects, but this need be nothing more than a manner of speaking (ibid. 26)

We may wonder whether is it ‘talking about linguistic types’ or ‘talking “as if” we were talking about linguistic types’ that is the manner of speaking. What can it even mean to say linguists study tokens? Perhaps we can make sense of what Devitt has in mind by taking acoustic signals to be physical tokenings of the syntactically-structured sentence types that linguists actually study. If so, how are the particular acoustic tokens *typed*? (Note that some acoustic sequences can be typed by more than one sentence structure, but let us put this complication aside.) We are told that ‘physical sentence tokens’ are ‘governed by a system of rules’ (26):

Something counts as a particular sentence, has its particular syntactic structure, *in virtue of* the particular structure rules that govern it (24, italics mine.)

This won’t do. Linguistic structure rules do not apply to physical sound properties: they apply to elements in a structure; elements that carry information about grammatical roles and categories, including phonetically null elements in the syntax. It is these elements that are said to be represented in the minds of speakers. Devitt misses the thrust of Chomsky’s mentalist’s claim because he takes the mentalist to be saying that linguistic symbols have their properties *in virtue of certain mental facts*, as if mental facts gave rise to the fact of symbols having those properties, or as if the symbols got their properties by standing in relations to mental facts. But the mentalist’s claim is that the *linguistic symbols themselves* are mental representations. A sound is only a linguistic item (symbol) if speakers, or their underlying linguistic systems, assign phonological, syntactic and semantic properties to it: and these are not properties in the world but properties mentally represented in the minds of speakers. This is why a sound is heard *as* a piece of language. Without the cognitive wherewithal to represent sounds as symbols they would be heard as no more than inarticulate gurgling or throat clearing. Symbols are what we, as linguistically competent individuals, take sounds and marks to be, by hearing or reading into them the linguistic significance they have for us.

What is the alternative supposed to be? Where in the physical world could we locate a physical sentence token of (1) which has the structure represented by (2)?

- (1) John seemed to think Peter promised to leave
 (2) John₁ seemed [t₁ to think Peter₂ promised PRO₂ to leave]

The structure rules that govern the sentence we take (1) to express apply to the represented structure shown in (2), a structure remote from surface form. This same syntactic structure—a type—could be assigned to a variety of sounds made by different speakers, or to written marks, or to instances of sign language. In what sense, then, are linguists studying physical tokens rather than linguistic types?

Towards the end of the book, Devitt eventually addresses the right question by asking:

In virtue of what does a sentence have its syntactic structure, whether explicit or not? What makes it the case that a particular token, for example a token of (1), has the structure it has, perhaps the structure revealed by (2)...? This question does not seek a *description* of the structure that the sentence has—a description that linguists have been so successful at providing—but rather an *explanation* of its having that structure. Since this structure along with the sentence’s word meanings determines the sentence’s meaning, this explanation is part of an explanation of the sentence’s meaning. It is part of an explanation of in virtue of what the sentence has whatever meaning it has. (ibid. 55; numbers for examples changed)

The appeal to meaning is meant to help explain why a string has the structure it has. Devitt takes it that structures and meanings are made for each other. But sentence meaning depends systematically on form and, a fortiori, the meaning comprehended depends systematically on the syntactic structure of the string we assign to it. We see semantics *through* syntax, as David Wiggins once put it. So where is this perceived structure? In virtue of what does a string of sounds or marks have the structure the linguist describes? We are offered nothing more to go on. And where more detail is needed about how Devitt thinks about the outputs of linguistic competence—the putative objects of linguistic inquiry—he tells us:

The theory of them is as much concerned with real-world objects as the theories of horseshoes [and] bee dances (26).

Is there something here to cling to? Let us examine the ideas a little further.

2. Can Language Really Be Like Honey Bee Foraging or Horseshoes?

Unfortunately, much of the argument for separating the study of linguistic competence from the study of its outputs—and for the independent existence of its outputs—works by analogy, and so the claim is only as good as the analogies themselves. The first analogy concerns the blacksmith and the horseshoes he produces. The second is an analogy between language and the foraging Honey Bee’s dance. Let us look at these in turn.

The competent blacksmith makes horseshoes. These products lie around as parts of the world existing independently of the blacksmith. We could come across them without knowing how they were produced, or even what they were for. Once produced, horseshoes take on a life of their own and we can study them without knowing anything about the complex mental and physical properties the blacksmith needed to make them. For Devitt, the same is true of linguistic entities and the competence of those who produce them.

The analogy is deeply flawed. It is true that once horseshoes have been produced, *they* can exist independently of the blacksmith. We could come upon them without knowing what they were, and begin to fashion all sorts of hypotheses about them. Even if we did recognise them we could describe and study horseshoes without knowing how to make them. But it cannot be like that with the sentences of language. In order to come across them, or recognise them at all, we need to be linguistically competent users of a language who know what they are. Without linguistic competence such bits of ‘linguistic reality’ would be invisible or inaudible: they would not be encountered as *linguistic* but just as noise. We only recognise speech events as including sentences when we have the competence to speak and understand. Without understanding we could not tell the difference between inarticulate gurgling and language. Even if we speak a language, how could we tell, when listening to a foreign speaker, whether he was actually speaking Russian or merely making noises that sounded like Russian (or sounded as we imagined Russian to sound) unless we understood Russian? The data of linguistics—the bits of linguistic reality Devitt talks about—are only data for those who recognise them as such by exercising their linguistic competence. The sentences—syntactically structured strings—that we see or hear depend on the competence we have. Unlike horseshoes, linguistic data are not out there independently of the competence of those who seek them out. They are there for those with the linguistic competence to comprehend them. Study the competence and you will find the sentences it carves out of the ambient sound stream.

The second crucial analogy for Devitt is the Honey Bee Dance, described by von Frisch as a ‘language’. According to von Frisch, the foraging honey bee finds a food source and returns to the hive to perform a dance with wagging and turns of its body. The other bees observe this dance and use these movements as a set of instructions about the direction and distance of the food source they then fly to. Once again, Devitt is talking about events in the physical world that can be observed without knowing anything of their significance, what these movements are for, and how other bees make use of them. It is precisely this state of affairs that allowed von Frisch to speculate about the significance of what he saw and to describe the bee dance as a ‘language’.

Although Devitt lays great weight on this analogy, the “language” hypothesis is highly controversial and has very little empirical credibil-

ity these days. Work by Adrian Wenner repudiates many of von Frisch’s findings and strongly suggests that the bees are much more likely to use odour signals to direct each other to the food. The wagging may turn out to be an efficient method for the foraging bee to spread odour to the others. The lead bee most probably leaves odour trails along the route for the other bees to follow.⁴

There are many weaknesses in von Frisch’s experiments and a lack of adequate controls. Many of his predictions have turned out to be false, and in a recent survey paper discussing the controversy, Subhash Kak reports that:

The 1980’s saw a renewed challenge to the dance theory (Rosin 1980, 1988). In the face of evidence against the original dance language hypothesis of von Frisch, even its proponents now grant the locale-odor mechanism a role in the communication amongst bees. But, as argued by Rosin (1988), these proponents do not, amongst themselves, agree on a ‘single’ dance language. Says Rosin: ‘There are numerous contradictory versions of the ‘dance language’ hypothesis that concur only in the belief that somewhere, somehow, some honey bees use ‘dance language’ information, but disagree on practically anything else.’ (Kak [1991])

The availability of negative evidence, the lack of replication of Von Frisch’s results, the proliferation of different versions of the dance signals and their meanings, the counter-evidence for odour-based tests all suggest that Von Frisch’s ‘language’ hypothesis has little going for it. An odour mechanism that explained the data would not carry the heavy and implausible commitment to the bee’s brain having an unknown cognitive mechanism able to process the variety of information allegedly displayed in the different bee dance signals. Unlike language, the bee dance can be observed without any idea of the significance of the activity we are looking at. But equally, the best hypothesis of what we are looking at is unlikely to have anything to do with representational language. For Devitt’s purposes the analogy proves doubly unhelpful.

When called upon to give a positive construal of properties of linguistic structure such as gapping, c-command, the location of empty categories like trace and PRO, we saw that Devitt takes these to be non-obvious properties of physical sound tokens. The elements in linguistic structures and their relations are not to be equated with brute-physical entities. These non-obvious properties of sounds or signs are ‘higher-order’ relational properties (see 27). Resorting to another analogy, Devitt tells us they are like Michael Devitt’s (amusingly) ‘non-obvious’ property of being Australian. This is a relational, not an intrinsic, property of a person. Similarly, in the language case, physical sound tokens will be said to have non-obvious relational properties. However, there is a considerable difference between the cases. In the former there is no mystery. A spatio-temporal path can be traced between Michael

⁴ See Wenner [2002]; Wells and Wenner [1973]; Wenner [1967]; Wenner and Wells [1990]. I am very grateful to Ophelia Deroy for drawing this research to my attention.

Devitt and a large continent in the Pacific Ocean. But what are the relations in the analogous linguistic case? Linguistic properties are not intrinsic properties of physical tokens. A sound's having these linguistic properties consists in there being a relation between the sound and features of a speaker's mind/brain. For this is the only candidate for the other relatum, as Devitt himself seems to admit. As we saw, phonemic sequences, grammatical categories, and the dependence of some items on others in the string, are not to be found at the physical end of this relation. All the fine-grained linguistic detail and richness, which gives sounds their place in a language, are to be found in the representations in the speaker's mind/brain. Why not then concede that it is these domain-specific features of a speaker's cognitive organisation, rather than the brute-physical tokens, that the linguist most needs to focus on? Moreover, it is hardly an alternative to mentalism to hear that the non-obvious properties that make sounds parts of language are part of a speaker's psychology. Sure, we can say that physical tokens of sound *have* linguistic properties, but crucially they do so because they stand in important relations to the psychological states of language users. Which linguistic properties tokens have, and how speakers come to recognise them, concerns the performance and competence of the speaker-hearer, and not just the noises they emit at different places and times.

3. *Devitt's Difficulties With Competence*

Coming at things from a different angle, why does Devitt struggle to resist the inclusion of linguistic competence in the subject matter of linguistics? After all, if foregoing discussion is right, he will have to concern himself with speaker's competence in order to recognise physical sound tokens as having the linguistic properties that make them part of language. What is more, if the leading practitioners in linguistics take themselves to be studying speaker's competence, why not leave it to the experts to settle the subject matter of their own discipline? The anxiety is all about maintaining a view of language as external to the minds of speakers, and resisting Chomsky's conception of language as part of the mind/brain of speakers. According to Chomsky: 'language has no objective existence apart from its mental representation' in the mind/brain⁵; according to Devitt, 'linguistic competence is in the mind/brain, the language is not.' (23) However, as we have just seen, Devitt fails to make out the case that linguistics concerns spatio-temporal tokens of the physical world rather than being as Chomsky thinks 'mentalistic, since it is concerned with discovering a mental reality underlying actual behaviour.' (Chomsky [1965], 4). As we saw, Devitt is unable to characterise the outputs of linguistic competence as linguistic save in relation to a speaker's competence. And if it is conceded not only that 'language is what the competence produces' (31) but that language is inseparable

from that competence, the only remaining issue concerns the right account of that competence. For if Devitt were right and Chomsky was mistaken about the nature of linguistic competence it would not follow that he was mistaken in his mentalist view of the subject matter of linguistics. But is Chomsky's account of linguistic competence mistaken? Devitt certainly thinks so. The question is whether Chomsky holds the view that Devitt spends most of his time attacking.

4. *Devitt's View of Chomsky on Competence*

According to Devitt, Chomsky thinks that the rules of language linguists posit in syntactic theories are psychologically real: they are the rules speakers embody by mentally representing them. What is more, 'there is no significant evidence for this thesis, and given what else we know, it is implausible' (9) The key thesis concerns the representation of rules. According to Devitt:

Chomsky puts the claim about representation with characteristic firmness: "there can be little doubt that knowing a language involves internal representation of a generative procedure" (1991a: 206; see also 1965: 25; 1975a: 304; 1980a: 201; 1980b: 9; 2000a: 50). (Devitt [2006], 4)

In response, Devitt tells us that we should:

doubt that speakers must have *propositional* knowledge of the language or that they must have representations of linguistic rules in the language faculty or anywhere else in the mind. (ibid. 5)

Notice the discrepancy between Devitt's rejection and what Chomsky is actually saying. Chomsky says that knowing a language '*involves* internal representation of a generative procedure' (my italics), a generative procedure that determines the expressions of the language and their various linguistic properties. He does not say a speaker/hearer has *propositional knowledge* of the rules of the language. Chomsky is notoriously wary of talk of knowledge and the problems philosophers see it as giving rise to: that was why he proposed the use of the neologism 'cognize' instead of 'know' to indicate that linguists were concerned with a rather special cognitive state. To say, as Chomsky does, that some special part of the speaker's cognitive apparatus represents a generative procedure is a subtly different and noticeably less committal claim than the one Devitt attributes to Chomsky.

However, I sympathize with Devitt. Chomsky is notoriously difficult to pin down on what he means by a speaker having internalized a grammar for his language. At one stage, Chomsky says:

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)

The statement is misleading. For if items of conscious knowledge flowed from unconscious knowledge in a manner similar to deduction, it would suggest that unconscious knowledge was propositional in form:

⁵ Chomsky [1972], 169 fn.

something having the right format to sustain logical relations. Though Chomsky's picture is surely different. Unconscious, information-bearing states of the language faculty give rise to conscious knowledge that is immediately reflected in the speaker's intuitive linguistic judgements, but there is no reason to think, and it is surely implausible to think, that such unconscious knowledge takes exactly the same form that the theorists' pronouncements take. In the book, *Knowledge of Language*, from which the above quote comes, Chomsky is careful to distinguish between a speaker *knowing that rule R is a rule of his language*, and his *knowing rule R*. The mistake of supposing that the speaker represents the theory the linguist comes up with is due to an assumption that the speaker somehow embodies the information laid out by the theorist. This cannot be Chomsky's view. But it was Quine's. It is Quine's thought that the learner and the theorist must be in the same predicament and that therefore the learner will have the same (behaviourist) data to go on as the theorist who tries to provide an account of the language. Both face the same evidence from which they have to arrive at an understanding of the language. For Quine both must be behaviourists, and he suspects that mentalists assume that the language learner has knowledge, albeit unconscious, of the theory the linguist constructs. However, Chomsky distinguishes the project facing the linguist and facing the child acquiring a first language. The linguist is trying to work out what native endowment the child must be equipped with in order to acquire natural language on exposure to a course of experience. The child, notice, is doing no such thing. The child is simply exposed to that course of experience and on the basis of its native endowment, somehow maps the course of experience onto an internalised grammar that equips it to speak and understand. The theorist will try to identify the nature of that native endowment, the principles that govern the choice of resulting grammar, and attempt to understand the way the child's resulting grammar grows out of the initial structural properties the language faculty respects on exposure to the primary linguistic data. And notice, one again, the generalisations the linguist frames about the child's grammar need not be framed in the form in which they are observed by the grammar. The linguist's theory is a model of the grammar not a specification of the form of its encoding by the speaker.

It is easy to agree with Devitt that we should doubt whether speakers have 'propositional knowledge' of the rules of their language, or represent rules as formulated by linguists' theories of their languages. But why should anyone think this is what the Chomskian is committed to? The theorist tries to state (partly propositionally) what the speaker knows, but not in the form in which the speaker knows it. The theory does not attempt to show how such knowledge is put to use. If accounts which do try to do these things are the target of Devitt's book then it simply fails to address Chomsky's position. Be that as it may, the key issue is whether

a speaker really could 'be competent in a language without representing it or knowing anything about it'; whether she really could 'be totally ignorant of it' (5), as Devitt asserts. Were we to consider a speaker's knowledge as consisting in the representation of rules of her language—her internalisation of the rules as formulated by the linguist—we could agree that the speaker does not have knowledge in anything like that form. But what would that tell us about more plausible mentalist claims? Devitt's boldly stated conclusion is surely meant to apply to all mentalistic or Chomskian claims. If not, it's hard to see why anyone should get so excited about 'the ignorance claim'. I shall argue that Devitt's criticisms leave the central tenets of Chomsky's mentalist programme intact, and, as we have seen, Devitt's own views about the nature of language and the objects of linguistic inquiry go badly awry. Similarly, his thoughts about how a competence that involves neither knowledge nor representations of the language enables us to produce and comprehend indefinitely many sentences are wildly speculative.

5. *Devitt's Dialectical Strategy: the Real and the Unreal*

It must be said that Devitt's dialectical strategy is puzzling. It appears to be as follows: suppose Chomsky were saying that competent speakers of a language had propositional knowledge of its rules; look what a mess he would get into. What should we conclude from that? The natural report is to say: why suppose this is the right or the only way to interpret Chomsky, why attribute to him such an implausible view? Why is no real attempt made to make out a more plausible and less straw man version of Chomsky's position? I sympathise with Devitt's difficulties in trying to arrive at a consistent interpretation of Chomsky on the psychological reality of grammars. Nevertheless, Chomsky has made clarificatory comments designed to defuse criticism, and these remarks should be heeded. Devitt does concede that 'we cannot be confident that Chomsky is using 'represent' in the familiar and relatively clear sense that I have endorsed (which is also the sense relevant to RTM).' And he insists that 'interpreting Chomsky is not my major concern'; the major concern being 'to evaluate a variety of ways in which language might be psychologically real in the speaker' (7). The crucial questions should be whether linguistics is part of psychology and whether linguistic rules are psychologically real or linguistic facts are psychological facts. Yet, despite this broader set of concerns and the declaration that he takes 'no firm stand on this matter of interpretation' (7), Devitt goes on to assume that, 'the natural interpretation attributes RT to Chomsky' and he proceeds as though this were not only the correct but also the only position worth considering. Devitt does this not only to reject what he takes to be Chomsky's position but also to reject the claim that 'language might be psychologically real in the speaker'. The trouble is that the interpretation fixed upon is the least plausible construal of the mentalist's claim; one which even leads Devitt into confusion. For instance, things go very badly wrong in this passage:

The rules of the language are represented in the speaker's language faculty; the language is psychologically real because the theory of the language - its grammar - is psychologically real. Indeed, a language simply is this system of rules encoded in the mind. Those represented rules are the reality that a grammar is theorizing about. If that interpretation is wrong then Chomsky must hold that a language is a system of rules embodied somehow in the language faculty without being represented. Then those embodied rules are the reality that a grammar is theorizing about.' (7)

Let us suppose rules are mentally represented by the language faculty. We are then told that the language has psychological reality because the theory of it—its grammar—is psychologically real: i.e. mentally represented. But why should this follow? Even if the linguist's theory of a language were mentally represented that would not show that language was psychologically real; we could still be theorising about an external language. This is the view advocated by Gareth Evans, Martin Davies and Christopher Peacocke who argue that speakers had tacit knowledge of a semantic theory for their public languages.⁶ But the claim Devitt focuses on is the claim that language *is* the system of rules encoded in the mind. (If this is right, what are the rules rules of?) We are then told that the grammar (theorist's sense) is theorising about the represented rules. So the grammar is not a theory of the language unless we accept the equation of the language *with* the rules of the language. And since the grammar (theorist's sense) is also mentally represented, speakers are tacitly theorising about their own mentally represented rules. This is way off track—the rules as mentally represented comprise the mentally represented grammar (psychological sense), they are not the subject of another mentally represented grammar (theorist's sense). The embodied rules are not what a psychologically real grammar 'theorises' about, even if grammar is psychologically real. Devitt fails to heed Chomsky's explicit distinction between the theorist's notion of grammar and the speaker's internal grammar. The theorist's grammar is a model of the speaker's grammar, and only the later is internally represented.

We, the theorists, can represent the body of information encoded by the speaker, and *represent* (model) how it issues in intuitive judgements by displaying a derivation in a generative grammar. A derivation is a logical derivation not a matter of actual processing. The theoretical derivation displays the way information about the underlying structure of language can issue in pronouncements about particular strings: but it is a task for the psycholinguist to tell us how the actual processing story goes and how it enables the speaker to arrive at his judgements given the linguistic information he possesses.

The rules of a grammar, a syntactic theory, records, in finite form, the knowledge the speaker has about a language. Once again the theory states what the speaker knows but not the form in which speaker knows it. And it is not part of the mentalist's claim to say what form speaker's knowledge takes. The theory states the knowledge the speaker has but

it is not obliged to say how that information is encoded mentally by the speaker.

Notice, once again, Devitt assumption: 'According to Chomsky, *on our natural interpretation*, this competence involves representations of the rules of the language.' (italics mine). He continues:

So those representations determine what expressions the speaker produces and understands. According to the point about intuitions, those representations also determine what the speaker says about those expressions in her intuitive judgments. (ibid. 4)

The rules of the language determine the expressions the speaker produces and understands. Of course, they are not processing rules so 'determine' must mean 'fixes' which expressions count as ones the speaker could (given some further psychological facts) produce or understand. Either we represent *rules that fix the expressions of our language*, or it is the *representations of rules* that fix which expressions belong to our language. The representations of the rules also provide the speaker with knowledge of the language and enable her to arrive at intuitive judgements about expressions generated by those rules. Though, how does this square with Devitt's claim that the speaker 'has knowledge about her I-language'. For if the I-language is the set of rules (in Devitt's terminology) then the speaker doesn't have knowledge about that: her knowledge *consists in* representation of the I-language (set of rules). And what our knowledge consists in cannot also be what that knowledge is about. Once again, something is wildly amiss in Devitt's setting up of the picture he wants to attack. He tells us that

The linguist produces a "grammar" which is a theory of the I-language. That theory, hard-won by the linguist, is precisely what the speaker tacitly knows. (4)

But how can that be? The speaker is now said to tacitly know (mentally represent) a theory *about* the set of rules of the language. For even if the speaker is said to know the theory, this is a theory *about* the rules of the language. Yet, surely, what the speaker knows, if anything, is the language: which expressions count as part of the language as fixed by the rules. And yet, on the picture Devitt attacks, this consists of knowing the rules of the language. But in the picture Devitt gives us in the above quote, the speaker stands at one remove still, knowing not *the language* or *its rules*, but *a theory about the rules of the language*. A speaker's knowledge or competence cannot consist in knowing the rules if that knowledge is also about the rules. If it is knowledge about the rules of the language, analogous to the linguist's knowledge (about those rules), then speaker's knowledge cannot consist in represented rules. The mentalist's claim is that speakers somehow encode the rules of their language, and the question is: how do they encode them? But if we are now saying that speakers encode *a theory about the rules* of their language, we would find ourselves asking a different question about how speakers succeed in encoding the *theoretical statements* about rules: a question at one remove from the one about *rules* pressed earlier.

⁶ See Evans [1981]; Davies [1987]; Peacocke [1987].

A more ‘natural interpretation’ is available. The I-language studied by the linguist is the speaker’s grammar. And the grammar (in the theorist’s sense) models the grammar (in the psychological sense) of the speaker. When we say the speaker knows the grammar we are saying no more than this. Compare, Michael Dummett’s claim that a speaker has implicit knowledge of a theory of meaning when that theory is a theoretical representation of a practical capacity, a theoretical representation organised in deductive form as a set of propositions.⁷ According to Dummett, we don’t have to say *speakers* have propositional knowledge of the theory: the propositions of the theory *represent* the knowledge the speaker has which *consists in* a number of interrelated abilities. The theory segments the overall practical ability into component abilities and matches each proposition of the theory with an ability the speaker possesses. This is not the picture Chomsky favours but it shows yet another way of relating speakers and theories of what they know without supposing, implausibly, that speakers mentally represent these theories of their *own* linguistic knowledge.

6. Access

Consider Devitt’s further point about how knowledge determines speaker’s intuitions—their intuitive linguistic judgements. In ascribing to Chomsky the view that speakers have propositional knowledge of rules of their grammars, Devitt claims that this knowledge not only enables them to produce and understand expressions but ‘those representations also determine what the speaker says *about* those expressions in her intuitive judgments’. How are speakers able to do this? We hear that: ‘*simply* in virtue of being competent, speakers have propositional knowledge of syntactic facts; their competence gives them “privileged access” to this ‘reality’, and ‘intuitions reflect the underlying representations of the rules of the language.’ (4). Can this be right? To begin with, ‘privileged access’ is a matter of having first-person, conscious knowledge of our mental states, and as far as I know, *no-one* (not even Devitt, see 96) wants to claim that we have privileged access to the reality of the rules of our sub-personal linguistic system. So talk of intuitions reflecting ‘the underlying representations of the rules of the language’ cannot mean that we have privileged access to the rules. But nor do we have privileged access to ‘the syntactic facts of the language’ if these are meant to be in line with the immediate outputs of the language faculty. If we did the task of the linguist would be much easier. They could simply ask speakers what the syntactic facts of their language were, or engage in introspection about the syntax of their own languages. But linguistic theorising is not like this. Speakers’ intuitions do not give us access to what our internal linguistic system delivers—the syntactic facts. Instead, the linguist relies on speakers’ conscious intuitions about accept-

⁷ Dummett [1993], 132.

able or unacceptable forms of expression—acceptable or unacceptable for the speaker—that provide *evidence* about the syntactic facts of the speaker’s language, from which the linguist tries to formulate generalisations in terms of grammatical constraints or principles that will apply to all expressions of that language. These can be tested by consulting further intuitions of the speaker that either confirm or infirm the predictions about structure made by the theorist. The only privileged access speakers have is to their own linguistic intuitions about particular sentence forms — acceptable strings. They are generally authoritative (though not infallible) about the grammatical properties of their languages.

7. Respect

A further problem for the exposition of Chomsky is Devitt’s assumption that ‘To be competent in a language is to be able to produce and understand the expressions of that language’. This is to treat competence as an ability to speak and understand. But this is a mistake as Chomsky himself has frequently stressed. The use and understanding of expressions are performance matters. Competence issues in performance as mediated by other cognitive capacities, such as memory, attention, auditory perception and motor control. And since competence helps to *give rise* to performance, it cannot *consist in* it. Competence is just one factor among others responsible for our ability to produce and understand expressions, it informs our production and comprehension of speech.⁸ Devitt appears to acknowledge Chomsky’s view that competence leads to performance both in a footnote (5, fn.4) and when he says of a speaker ‘What role does her knowledge of the language play in understanding and producing expressions of the language?’ (4) Though it is of course much easier to reject the view that competence is knowledge-involving if you think that competence just *is* the ability to understand and produce expressions. What gives rise to this ability, why is it so systematic in its workings, how does it equip us to produce and understand indefinitely many new sentences just as easily as we use and understand familiar ones? Under pressure, Devitt, admits that the linguist’s grammar describes ‘*what* the competent speaker knows’ (not just the ability). Speakers’ knowledge is articulated in the grammar by linguistic structure rules, and ‘the speaker behaves *as if* those linguistic structure rules were psychologically real in her, *as if* she embodied them’ (ibid. 25). The ‘as if’ talk is necessary, according to Devitt, to register the fact that a correct theory (grammar) of the language is one that is respected by the speaker’s competence:

A theory of the competence must posit a psychological state that respects the

⁸ Chomsky has pointed out that certain speech impairments may lead one temporarily to lose the ability to produce and comprehend expressions without losing one’s competence. Competence may still be intact, and eventually issue once again in speech after a time.

rules governing the linguistic outputs. And a theory of the linguistic outputs must posit rules that are respected by the competence and its processing rules [parser] (ibid. 25)

This is the minimal position on psychological reality: ‘there is an uncontroversial minimal position that is committed only to there being a psychological reality that “respects” the linguistic structure rules’:

[We] ‘know something substantial about a person when we know there is something-we-know-not-what within her that respects the rich and complicated structure rules of a certain natural language’ (25)

But now the looming question is this: how does our cognitive organisation respect linguistic structure? How does competence—a mere ability of one who may be totally ignorant of the language—succeed in respecting the rules of the language? How can a psychological state in us be sensitive to linguistic structure and to the constraints on linguistic structure without in some way representing the linguistic properties and their distributions among the expressions it handles? After all, what we encounter perceptually are sounds. Whatever story we have about its sensitivity to, and respect for, structure, there must be a mechanism for finding those differences in structure between ‘John is easy to please’ and ‘John is eager to please’ and between ‘John promised Peter to go’ and ‘John persuaded Peter to go’ that treats as structured the strings of sounds or marks it responds to. And:

The only serious theory anyone has ever proposed about how a brain can be sensitive to grammatical, phonological and semantic categories is a computational/representational theory (a “CRTT”), which posits computations over representations of those categories. It’s very hard to see how this could be paraphrased away without loss of explanatory force. (Rey [2005], 471)

The only way in which we, or the unconscious apparatus at work in us that responds to other people’s speech sounds, can be sensitive to phonological, grammatical and semantic categories is due to our having the means to assign those categories mentally, or call upon them, in the course of cognitive information processing of (certain) speech sounds. Internal or intermediate objects of processing are required to explain our sensitivity to these unobservable linguistic entities. The linguistic phenomena the listener hears, depend on the linguistic object the listener’s internal apparatus supplies in response to speech sounds encountered. Competence in a language is having the means to make use of linguistic objects with the appropriate, related properties specified by the linguist’s theory of grammar. We automatically respond to certain speech sounds with a linguistic percept having properties not locatable in the sounds (acoustic signals). In response to the limited sensory input utterances of speech sounds provides, the mind constructs a rich, syntactically structured, interpreted sentence. How does our cognition—what we have within us—enable us to do this? Whatever it is, it must deal with information about recursive structures the grammar provides, for these are not found in the external environment. And yet, Devitt clings to the view that competence somehow engages with an independently

existing linguistic reality. The Respect Constraint tells us that:

The competence of the speaker/hearer must respect the rules ascribed to linguistic reality by the grammar.

But if the competence of the speaker/hearer must respect ‘the rules ascribed to linguistic reality’ then we could study the structure of linguistic reality by studying the speaker/hearer’s competence. Initially, this may seem to be an indirect route, but given the difficulty of locating linguistic reality, wherever it is supposed to be, or of gaining access to it unfiltered through the competence of the speaker/hearer that tells us what its available forms are, there seems to be no other choice. What would provide a more direct route to studying linguistic reality? Either one could use one’s own competence to select the acceptable forms of the language, or one would simply study actual use—linguistic performance. But there is a well-known reason why we don’t study actual output only. It is messy, full of false starts, uncompleted utterances, slips of the tongue, interruptions, and speaker/hearers would not always take what they uttered to be a true reflection of what they would say under good performance conditions. So, instead, we ask them to consult their intuitions about the data presented to them. These judgements give us cleaner data but also give us something more in accord with their competence. By getting at the constraints respected by competence, we get at the rules that govern the linguistic reality competence creates and makes available to them.

In the final analysis, accepting the Respect Constraint is a way of admitting that the study of outputs will be perfectly aligned with deliverances of a competence theory while remaining lily-livered about the psychological reality commitments—as the ‘as if’ talk shows. However, the important point is that since there is no direct access to *the* language, no way of getting at the syntactic properties of a language that its speakers must respect without eliciting the data from enough linguistic intuitions to fashion a theory of the full extent and nature of speakers’ competence, the Respect Constraint will be trivially met and the study of outputs will be the study of the psychological properties of speakers’ mind/brains that assign form and content to the signs and sounds they encounter.

In what way can Devitt avoid triviality? What can he tell us about the psychological states and what it respects? In a fair and concessive passage, he allows that perhaps:

The grammar is descriptive of competence in a way stronger than simply positing rules that are respected by competence but weaker than positing rules that govern processing or are used as data in processing (ibid. 83)

I would agree. The mechanisms in the speaker/hearer must not only conform to the linguistic generalisations given by the theory of the grammar. But the stronger position of the mentalist is that the cognitive organisation and mechanisms of the speaker/hearer do not only satisfy the generalisations, but by making them available they are the source of the linguistic generalisations. The generalisations hold in virtue of these cog-

nitive psychological facts about the speaker/hearer about the representation not of rules but of the structures the rules determine. The facts about languages cannot be identified independently of their speaker/hearers. They are dependent on the generalisations speaker-hearers' representations of structure can sustain. Without this dependence why would the syntactic properties of natural language expressions be respected by psychological states of speaker-hearers? In answer, Devitt speculates:

Humans are predisposed to learn languages that conform to the rules specified by UG because those rules are, largely if not entirely, innate structure rules of thought. (273)

But unlike the rigorously specified principles of UG, which, together with an account of the possibilities of parametric variation left open by the principles, explain the differences between languages, there is simply no account of the structure rules of thought. There is no explanation of why they would have to take the special, idiosyncratic and largely unchanging course linguistic structure rules do, or of how they operate in the same way across people with different levels of intelligence, exposure to speech, and background cultures. Devitt's unsupported assumption that we should replace the innate principles of UG which govern the initial configuration of the language faculty in all speakers by some (unspecified) set of innate structure rules of thought appears to be a desperate attempt to avoid the Chomskian conclusions at all cost. Alas, the savings are far from obvious.

7. Chomsky's Reality

There is considerable distance between Chomsky's official view and the position Devitt attributes to him, but ironically the real Chomsky is dangerously close to the ignorance claim since it denies that speakers have knowledge of a language.

As we have seen, according to Chomsky, language is in the mind:

Linguistics is simply that part of psychology that is concerned with one specific class of steady states, the cognitive structures that are employed in speaking and understanding. (Chomsky [1975b], 160)

Languages are not represented by our internal system, nor do we have to pair speakers to languages. Speakers have a linguistic capacity or infinite competence because of their innate endowment of universal grammar and their initial exposure to data: data that don't determine their 'language' or I-language, as the poverty of stimulus arguments tell us. Language, as now understood by the linguist, is the internal mechanism that enables us speak and understand: a very different notion of language from that Devitt works with. Consider the following from Chomsky:

...what should we take as a language...? The natural choice is *g*, the generative procedure; thus a person who knows language *L* has a specified method for interpreting arbitrary expressions, such as ['Who do you think that John saw' and 'What do you wonder who saw', 'Who do you wonder what saw', 'He

likes John', 'His mother likes John', 'John likes him', 'John's mother likes him', *J* (a sentence of Japanese)] Let us call g_E the I-language that some particular speaker of English (Jones) has acquired. (Chomsky [1987], 181)

The generative procedure applies to representations in the mind, but these are representations in the sense of data structures, not representations with contents about something external to the speaker.

[W]hat could such contents be possibly true of? There is, as it were, nothing to get right. Languages are not external objects we can go right or wrong about....The languages we end up with will reflect, along certain narrow parameters, our initial experience, but we do not end up *representing* (or *misrepresenting*) that experience. The complex grammatical principles and features are just not in the data; that's the very point of the poverty of stimulus arguments' (Collins [2004], 514)

The speaker's language faculty does not represent the speaker's language, or its grammatical structure. Gone is the view of 'the language faculty generat[ing] grammars that are true of languages', and 'As go E-languages, so go grammars as things which might represent them':

Chomsky understands 'language' to be the internal finite generative procedure that accounts for our linguistic competence.' So construed I-languages are 'internal states of our individual language faculties'.An I-language is not an independent 'object'—a set of propositions—that is represented/known by a speaker-hearer; it just is a state of the speaker/hearer....The language faculty [and grammar] is not represented by the mind/brain; it is an aspect of the mind/brain. (ibid. 517)

This surely puts Chomsky beyond the target range of Devitt's criticisms. But notice now we are in danger of losing our grip on there being any such thing as a speaker's knowledge of language either on Devitt's view or on the orthodox Chomskian's view.⁹ Of course we do not end up representing our initial linguistic experience, nor are the complex grammatical phenomena in the data. And I agree with Collins that languages are not external objects, however this does not licence the idea that there is nothing linguistically speaking that we can go right or wrong about, or that we cannot have knowledge of language. The linguistic structures we deal with are internally generated and assigned to sounds and marks which otherwise carry no linguistic information.¹⁰ So what can we be said to have knowledge of?

⁹ I take the orthodox Chomskian view to be ably represented by Collins.

¹⁰ Even words are not, so to speak, 'out there'. Each item counts as a word only when it carries its full freight of phonological, syntactic, and semantic properties, and these are assigned by the language faculty of the speaker, not found in the world as the poverty of stimulus arguments establish.

9. Keeping Knowledge of Language

If knowledge of language is a state of the person and the language a person speaks is *constituted* by that state it is hardly knowledge that is at issue: for how can a genuine state of knowledge constitute the thing known? Knowledge requires there to be a subject matter we can be right or wrong about. And the trouble with the linguist's notion of knowledge-of-language, or I-language, is that it fixes the properties of the language rather than conforming to them. The problem for the mentalist in a nutshell is this: if the psychological states that constitute one's competence determine *the facts about one's language*, how can those same psychological states be at the same time answerable to those facts in the way knowledge requires? This is what I have called elsewhere 'The Missing Object of Knowledge Problem'.¹¹ And there had better be a solution to it if we are to claim that speakers *do* have knowledge of language and that this domain of knowledge constitutes the subject matter of linguistics.

Briefly, the solution I favour can be given along the following lines. Our linguistic intuitions are immediate judgements, not based on observation or evidence; they are simply internal responses to the sounds or signs we are presented with. They are first-personal conscious opinions about the acceptability or naturalness of certain strings of words (or phonemes). These linguistic intuitions provide evidence of linguistic facts. But how can judgements not based on any evidence give us objective linguistic knowledge? The answer is that there will be sub-personal facts about the language faculty involved in producing (giving form and character to) our intuitions: the language faculty dictates the structures assigned. And when speakers' intuitions are in step with these internal assignments they will count as knowledge. It is when our linguistic intuitions track the workings of the underlying linguistic system and the syntactic structures it assigns, that our intuitions amount to knowledge. But what do they give us knowledge of? Although it may seem as if they give us knowledge of what is produced and going on around us in our linguistic environment, actually, linguistic experience, like linguistic intuition, is only giving us knowledge of our own internal system's workings: the system which fixes the properties of all expressions of our languages. It can seem as though we were judging the acceptability ('okay for me') of sentences 'out there', as part of *the* language. But in fact we are responding to something in our own breast. There is nothing 'out there' but sounds and signs. The language that I speak is constituted by the internal system in me that assigns structure to the sounds and signs I encounter and produce. Language is essentially structured and the structure is internal to speakers' minds. And although this is not the face value phenomenology of our intuitive linguistic judgements (and it is good question what our phenomenology *does* present us with either with respect to the location of language or speech sounds) there

is an object of knowledge in the case of grammar. And since linguistic intuitions provide each of us with objective knowledge they furnish us with a subject matter for linguistics: the study of the individual's grammatical knowledge; the grammar of that person's idiolect.

10. Understanding Other People's Use of Language

A problem we still face with the view of language as individual and internal is explaining how we succeed in communicating with others. How can the assignments *I* make to the sounds *I* encounter, and *what I hear them as saying* ensure knowledge of *another's* linguistic meaning? I rely (inter alia) on my knowledge of grammar and my knowledge of word meaning. And while I am authoritative about what I mean by words, and about which word strings are grammatical, I am not authoritative about what you mean or what is grammatical for you. Notice, though, that I do not typically have trouble understanding others. I do not first hear sounds and then have to hypothesize about the linguistic significance others attach to those sounds. It cannot be like that: we simply speak and expect to be understood; we listen and respond without hypothesis. We hear sounds *as* meaningful. I hear the sound of what is being said as structured, and if it is not a structure of my grammar, it will be heard as deviant but intelligible—if I can re-arrange the words heard. The way communication works, *I* have to make something of the sounds *you* emit. But what meanings do I hear your words as having? This will not be a matter of sub-personal assignments by my language faculty. Word meanings do not belong in the language faculty: knowledge of word meaning is conscious and first-personal. And the meanings I hear your words as having are the meanings those words have consciously for me; they are the meanings I have managed to endow those sounds with in my early acquisition of word meaning. The default assumption—not even explicitly entertained—will be that you mean by your words what I mean by them; indeed the default is that this is what anyone means by this word.¹² (Later experience will provide plenty of evidence that this is not always the case and that the default must be overridden.) Because the conditions of learning take place in a common, shared environment, the meaning we attach to a word will be what others from whom we learned the word mean by it. And provided there is enough overlap within and between linguistic communities this will secure the usefulness of the default assumption and the entitlement to claim to know what others mean. The default assumption can be overridden, of course: others can use words with meanings different from the ones we attach to these sounds. Nevertheless, at first, we hear sounds as meaningful according to the language *we* each speak. That language is fixed by our knowledge of language. This consists (inter alia) in knowledge of grammar and knowledge of word meaning. Knowledge of grammar is knowledge of linguistic facts internal to our language faculty. Knowledge of

¹¹ See Smith [2006a] in Lepore and Smith [2006].

¹² For more on this see Smith [2006b] and [2006c].

word meaning is default knowledge of meanings that others too attach to these words; everything else is deliberative interpretation. We have knowledge of language but what we know when we know a language is not a single thing. It consists of different kinds of knowledge and different objects of knowledge. Some objects of our linguistic knowledge are mental and internal, some depend on our dealings with others and aspects of the environment. There is no single locus of linguistic significance, but through the joint exercise of our cognitive abilities we come to have a unified experience of the linguistic significance of speech. It may be that experience misleads Devitt about the nature and location of language and encourages him to attempt to dispense with knowledge of language. However, we will still need knowledge of language in the picture in order to understand our experience of language. The science of language may discard our ordinary notions but it still needs to retain the notion of knowledge of language.

Bibliography

- Chomsky, N. [1968], *Language and Mind* (New York: Harcourt, Brace & World).
- _____[1975], *Reflections on Language* (New York: Pantheon Book).
- _____[1986], *Knowledge of Language: Its Nature, Origin, and Use* (New York: Praeger).
- _____[1987], 'Reply to George', *Mind and Language*, 2, 178–197.
- Collins, J. [2004], 'Faculty Disputes', *Mind and Language*, 19, 503–533.
- Davies, M. [1987], 'Tacit Knowledge and Semantic Theory: can 5% make a difference?' *Mind*, 96, 441–62.
- Devitt, M. [2006], *Ignorance of Language* (Oxford: Clarendon Press).
- Dummett, M. [1993], *The Seas of Language* (Oxford University Press)
- Evans, G. [1981], 'Semantic Theory and Tacit Knowledge' in S. Holtzman and G. Leich (eds.), *Wittgenstein on Following a Rule* (London: Routledge), 118–37.
- Kak, S. C. [1991], 'The Honey Bee Dance Language Controversy', *The Mankind Quarterly*, Summer 1991, 357–365.
- Peacocke, C. [1986], 'Computational Psychology and Level 1.5', *Mind and Language*, 1, 101–123.
- Rey, G. [2005], 'Replies to Critics', *Croatian Journal of Philosophy*, 5, 465–481.
- Smith, B. C. [2006a], 'What I Know When I Know a Language' in E. Lepore and B. C. Smith (eds.), *The Oxford Handbook of Philosophy of Language* (Oxford: Oxford University Press), 941–982.
- _____[2006b], 'Davidson, Interpretation and First-Person Constraints on Meaning', *The International Journal of Philosophical Studies*, 14, 385–406.
- _____[2006c], 'Publicity, Externalism and Inner States' in T. Marvan (ed.), *What Determines Content: The Internalism/Externalism Dispute* (Cambridge: Cambridge Scholars Press), 1–27.
- Wells, P. H. and Wenner, A. M. [1973], 'Do bees have a language?' *Nature*, 241, 171–174.

- Wenner, A. M. [1967], 'Honey bees: do they use the distance information contained in their dance maneuver?', *Science*, 155, 847–849.
- _____[2002], 'The Elusive Honey Bee Dance "Language Hypothesis"', *Journal of Insect Behaviour*, 15, 859–878.
- _____[1990], and Wells, P. H. [1990], *Anatomy of a Controversy: The Question of a Language Among Bees* (New York: Columbia University Press).