

CHAPTER 37

WHAT I KNOW  
WHEN I KNOW  
A LANGUAGE

BARRY C. SMITH

EVERY speaker of a language knows a bewildering variety of linguistic facts and will come to know many more. It is knowledge that connects sound and meaning. Questions about the nature of this knowledge cannot be separated from fundamental questions about the nature of language. The conception of language we should adopt depends on the part it plays in explaining our knowledge of language. This chapter explores options in accounting for language and our knowledge of language and defends the view that individuals' languages are constituted by the standing knowledge they carry from one speech situation to another.

The title of this chapter alludes, of course, to Michael Dummett's seminal paper, 'What Do I Know When I Know a Language?' in which he raises many fundamental questions for the philosopher of language. More than twenty-five years later, I hope we are at last beginning to see how to address some of the important foundational issues Dummett first brought to light. The chapter builds on issues first raised in Smith, 1992, 1998, and 2001. I should like to thank graduate students from the University of London and the University of California at Berkeley for their many helpful responses to this material; in particular Cheng-Hung Tsai, Julian Dutant, and April Jones. Thanks to Stephen Schiffer for discussion of his paper to which I responded at the GLOW conference in Geneva 2005 and to the audience at that session. And for their helpful conversations on many of these issues I should especially like to thank Paul Pietroski, Donald Davidson, John McDowell, John Searle, Jennifer Hudin, Stephen Neale, Guy Longworth, Peter Pagin, Kathrin Gluer, Asa Wikforss, Jim Higginbotham, Jason Stanley, Peter Ludlow, Georges Rey, Michael Devitt, Ernie Lepore, and John Collins.

## 37.1 LANGUAGE AS A SOURCE OF KNOWLEDGE

Language enables us to acquire knowledge of the world and of other people. We learn what people think by what they say, and, in turn, we speak our minds to them. Words give us immediate entry to the minds of others. Just by using these words I can inform you or amuse you, excite you or insult you. I get straight through to your mind, perhaps uninvited. Similarly, your words have immediate and unexpected effects on my thinking. This is possible because we hear people's emission of sound as meaningful speech, and cannot but hear it that way when the words uttered are familiar. In this way, language establishes intimate connections between minds and shows how easily the sanctity of individual minds is violated.

In addition to being an interface between minds, language gives us much of our access to the wider world. Through what we are told and what we read we come to acquire a vast range of world knowledge. Language is our means of learning about science and culture, mathematics and history: information that makes up our much of our vision of the wider world. To gain access to this knowledge we must first have access to language. So what gives us access to language? Do we first need to know a language? Is this what equips us to produce and understand utterances?

These questions go to the heart of our ability to make intelligible sense of certain sounds people utter, and to give meaning to the sounds we utter. How are we able to attach linguistic form and significance to certain speech sounds we and other humans produce? Michael Dummett frames the philosophical issue as follows:

The central task of the philosopher of language is to explain what *meaning* is, that is, what makes a language *language*. Consider two speakers engaged in conversation. To immediate inspection, all that is happening is that sounds of a certain kind issue from the mouths of each alternately. But we know that there is a deeper significance: they are expressing thoughts, putting forward arguments, stating conjectures, asking questions, etc. What the philosopher of language has to explain is what gives this character to the sounds they utter: what makes their utterances expressions of thought and all these other things? (Dummett, 1978, p. 96)

The needed explanation has to cover a potential infinity of cases. For unless people live dull and repetitive lives, everyday they will produce and hear utterances of sentences they have never heard before. They will hear news, read interesting books, and engage in distracting conversations. What explains this capacity to share their states of minds linguistically in new but comprehensible ways? Whatever it is, the remarkable fact is that we understand utterances of sentences we have never heard before just as easily as we understand those which are already familiar, suggesting the same system is at work in our handling of both the novel sentence and the sentences already understood. What is this underlying system and how do we exploit it to form and recognize new sentences?

## 37.2 LANGUAGE AS AN OBJECT OF KNOWLEDGE

After the passage just quoted, Dummett goes on to ask a key question: ‘Is the significance of language to be explained in terms of a speaker’s knowledge of his language?’ (1978, p. 97). We want to know what gives certain sounds their linguistic character, and we also want to know how we are able to hear those sounds as the utterances of meaningful sentences (in indefinitely many cases). One approach would be to say why we were entitled to treat those sounds as the articulations of expressions from a particular language and how we were able to recognize them as belonging to that language. Another approach would be to suggest that it is our ability to hear those sounds as meaningful sentences that confers linguistic character on them. Either way, we are only able to pair up a potential infinity of sounds and meanings by knowing a language. We need to know indefinitely many sentences, and know for any sounds we hear, which sentences are being uttered. So the difference between merely hearing people issuing noises and hearing what they are saying resides in one’s knowing the language being spoken. If so, we need to account for this knowledge—and how it enables us to hear and utter speech sounds as particular meaningful sentences and thus know the minds of others:

- (i) What form does this knowledge take?
- (ii) How do we acquire it?
- (iii) How does it enable us to hear *sounds* as meaningful sentences?
- (iv) How does it enable us to know what other people mean?

## 37.3 TWO CONCEPTIONS OF LANGUAGE AND KNOWLEDGE OF LANGUAGE

In addressing these questions it will be useful to consider two competing conceptions of language and our knowledge of language. On one traditional story, individuals in a linguistic community are able to use their language to express and convey their thoughts because of their participation in a common practice.<sup>1</sup> To understand other people’s utterances, to hear what they are saying, one must know the meanings of the words and sentences they use. These will be the meanings those words and sentences have in the common practice, or public language. To know these meanings one must belong to that linguistic community and participate in its practices. But what enables

<sup>1</sup> See Dummett, ‘What Do I Know When I Know a Language?’, p. 102 and McDowell, ‘In Defence of Modesty’, especially p. 94, and ‘Antirealism and the Epistemology of Understanding’, p. 314, and pp. 332–3.

one as an individual to participate in those shared practices? What is the extent of the knowledge one thereby acquires? What does it give us knowledge of? And how exactly does it enable one to hear sounds as part of a public language? The traditional story is short on detailed answers to these questions. We are told that participation in a practice is a gradual matter; that ‘Light dawns gradually over the whole’;<sup>2</sup> that ‘one hears more, in speech in a language, when one has learned the language’;<sup>3</sup> that one cannot explain the linguistic significance of items by reference to anything outside our linguistic practices; that the meanings of words and sentences lie in open view to the surface of our practice but are only visible or audible to participants of that practice.

It is crucial to see that *the* language, or the common practice that displays and sustains it, has to extend beyond individual speakers, and beyond the actual point any speaker has reached since it must make room for the potential infinity of meaningful sentences yet to be encountered that count as part of the language. The expressive possibilities open to us are already mapped out by the language and must somehow be contained, or latent, in the practice. An individual’s power of expression rests on how much of that language or practice has been mastered and how effectively it is put to use. The system underlying familiar and novel sentences is part of the workings of the language. Whatever gives us access to the language gives us access to that system. But what does give us access to the public language and how are individuals able to exploit the workings of the language for their own purposes? As Dummett points out, to understand a novel sentence we rely on our recognition of familiar words and the methods of sentence construction: for this we need knowledge of the meaning of words, of the grammatical means of combining them and some awareness of the significance of so combining them. It would be hard to overestimate what an achievement it is for a child to acquire this knowledge starting out by simply confronting sounds. Through language acquisition we succeed in putting our minds into our words in a way dictated by the requirements of the common language, so that others who do likewise can find out what we think.

A radical alternative sees the psychological states of speaker-hearers—their knowledge of language—as endowing speech sounds with the linguistic significance they have for us. Although we perceive sounds as linguistic items, on this view, the linguistic items are internal to the mind of a speaker-hearer (see ‘*Language as Internal*’).<sup>4</sup> Knowledge of language is a state of the speaker-hearer that fixes which sounds count as expressions, which arrangement of expressions counts as grammatical, and what those expressions are taken to mean. It does not provide knowledge of anything external to the mind of speaker-hearer: it is not knowledge of an external system of linguistically significant sounds, or signs, speaker-hearers have to master. Instead, all there is in the world are sounds and marks: it is we who give them their linguistic

<sup>2</sup> The phrase is of course Wittgenstein’s from *On Certainty*, §141, and is appealed to by McDowell (1998a, p. 333).

<sup>3</sup> McDowell (1998a, p. 333)

<sup>4</sup> This is the view advocated by Noam Chomsky. As he has put it: ‘language has no objective existence apart from its mental representation’ in the mind of the speaker, (1972, p. 169 fn.).

form and meaning. So it is not the sounds and signs people produce that constitute the subject matter of linguistics, but the *linguistic forms* people impose on those sounds and signs as a result of their internal states. It is because there are creatures like us, with the distinctive linguistic capacities we have, that signs and sounds come to be assigned a meaning and structure at all. On this view, the focus of linguistic inquiry shifts from the actual and potential behaviour of speakers to the internal organization of speakers' minds. The study of language becomes part of the study of mind, and so linguistics—the science of language—is seen as a branch of cognitive psychology.

As speakers, whether we are producing or perceiving speech sounds, it is we who supply whatever linguistic significance they have. The producer's awareness of the linguistic significance of the noises he emits is much like his awareness of his tapping out a tune for another with his fingers. The other may recognize the tune by the rhythmic tapping, or merely hear it as the drumming of fingers. But in the experience of the agent, the rhythm tapped out is an integral part of a whole musical score running through his head.<sup>5</sup> In a similar way, the speaker's psychological states give linguistic form and character to the sounds he experiences himself as producing. His utterances of sounds depend, for their particular linguistic character, on the precise psychological states that give rise to them, just as a bodily movement depends for its identity as an action on the states of mind of the agent performing it.<sup>6</sup>

Whatever experiences one has in producing speech sounds, comprehension of them by listeners is always due to what *they*, in turn, and in virtue of *their* internal apparatus, can make of the sounds uttered. Knowledge of language is a state that enables us to produce and understand a certain range of human speech sounds, as determined by our internal linguistic systems.<sup>7</sup> The difficulty for this account will be to explain how people succeed in communicating with one another, how language appears to put us immediately in touch with the minds of others, and why we often appear to speak the same language.

A correct view of language and our knowledge of language will need to account for our capacity to hear complex meaning in speech sounds and to invest sounds with such meanings; an account that explains our immediate readiness to produce and comprehend utterances of sentences we have never used or heard before. It will also have to explain how, by these means, we succeed in making our minds available to one another.

The traditional conception of language, and its accompanying view of knowledge of language, can be called, following Chomsky, an E-conception. 'E-language' was the term invented by Chomsky for languages conceived of as external, as sets of

<sup>5</sup> People are notoriously bad in their confidence judgements about how likely others are to recognize the tune they are tapping.

<sup>6</sup> This is not the view that linguistic productions are reducible to, or fully analysed in terms of, the intentional states such as beliefs and intentions. The mental states involved may be cognitive psychological states of the language faculty or, at any rate psychological states with dedicated linguistic contents.

<sup>7</sup> Chomsky conceives of a language: it is a 'way to speak and understand' rather than that which we make use of in speaking and understanding. See Chomsky, 1993, p. 49.

sentences extensionally characterized, and, we might add, as extended beyond the current reach of their speakers.<sup>8</sup> There are different notions of E-language. On a Platonist view, languages are abstract objects consisting of infinitely many sentences, each of which has at least one meaning in the language. Sentences of a language have a meaning for a speaker if and only if they are part of the language the speaker knows or uses.<sup>9</sup> The difficulty here will be to say what makes one rather than another of the infinitely possible languages the *actual language* of a given speaker, or community of speakers (see below).

By contrast, a social view of E-languages sees them as extrapolations from sets of common practices.<sup>10</sup> Rules enshrined in such public practices extend beyond the furthest reach of existing practice to determine which new combinations of words count as meaningful sentences of the language. In this way, rules dictate the precise contours of the language—the language spoken by those who participate in the practice.

The difficulty for the social view will be to say what counts as participation in a given practice as following rules for one rather than another language. Where the Platonist appeals to an infinity of *meaning facts*, the social theorist appeals to *meaning norms* governing infinitely many applications of the items found in the practice. (See ‘*Meaning, Rule-Following and Normativity*’.) On either of these E conceptions, a language is independent of any individual speaker; on the Platonist version, language is independent of all speakers.

The alternative conception of language, and knowledge of language, takes languages to be individual and internal to the minds of speakers. Let us call it a Cognitive Conception to evoke cognitive–psychological conceptions of language grounded in the psychology of individual language users.<sup>11</sup> On this view languages have no existence independent of human cognition. A change in a speaker’s cognitive organization can change his language. Selective loss of cognitive function can lead to loss of the language; and such permanent losses in all human language users would extinguish language altogether. The illusion of language continuing to exist in the absence of appropriately organized human minds would be due, perhaps, to the continued existence of recordings, written documents and signs that would no longer have linguistic significance for anyone.<sup>12</sup> One version of the Cognitive Conception, Chomsky’s, admits of no epistemic relation between language and knowledge of language: what is meant by ‘language’ is just an I-language, a finite part of the

<sup>8</sup> ‘The standard approaches [in philosophy] to developing a more technical concept [of language] take a language to be a variety of what I called ‘E-language’, where ‘E’ is to suggest ‘extensional’ and ‘externalised’: for example, a characterization of language as a set of utterance types, or a set of (utterance, meaning) pairs, where meanings are construed in set-theoretic terms. This general approach, however, leads to innumerable problems and is best abandoned . . .’ (Chomsky, *Mind and Language*, 1987, p. 179)

<sup>9</sup> Different varieties of Platonism have been advocated by George (1989); Higginbotham (1983); Katz (1990); Lewis, and Schiffer (1994).

<sup>10</sup> Advocates of this view include Dummett (1978); McDowell (1998); and Wiggins (1991)

<sup>11</sup> I avoid the use of ‘I-conception’ here with its association with I-languages since what Chomsky means by an I-language is a finite state of the mind/brain of a speaker; something that is definitely not an object of the speaker’s knowledge.

<sup>12</sup> There is nothing enduring about our prose without anyone left to appreciate it.

speaker's mind/brain: a generative procedure for assigning structural descriptions to expressions.<sup>13</sup> Language is just a means of speaking and understanding. It provides us with an infinite competence but the I-language itself is finite.<sup>14</sup> And since one's knowledge of language is just a state, amounting to possession of an I-language, it does not amount to knowledge of anything, certainly not to knowledge of the I-language. On Chomsky's view, nothing much remains of the ordinary notion of a language, nor of speakers' knowledge of language.

What this view downplays is a perfectly respectable notion of linguistic knowledge had by individual speakers. There is a vast amount of specific knowledge the speaker has and about which he is authoritative: including knowledge of what his words mean, knowledge of which arrangements of his words are sentences, and of how utterances of them can and cannot be understood.<sup>15</sup> Of course, knowledge involving sentences has to be derived since we don't carry around knowledge of individual sentences as part of our standing knowledge. Sentences are transient and ephemeral, no sooner produced than replaced by others, and there is evidence that in conversation we are not able to remember the sentences we have just heard even though we keep track of the conversation. Instead our standing knowledge provides the words and means by which we fashion indefinitely many sentences on the hoof and come to have knowledge of them. These are items of conscious knowledge that can be elicited by speaker's linguistic intuitions: judgements that give a speaker immediate and authoritative knowledge of linguistic facts. The states Chomsky is interested, however, are states that fix the facts of speakers' languages. But there is still a vast amount of conscious knowledge the speaker has about expressions of his language fixed by these states. It is an open question whether the states that give rise to the facts we know are themselves states of knowledge.

If so, we need a Cognitivist Conception that can do justice to the experience of the language user, to his having knowledge and being authoritative about what his words mean, about which arrangements of words are sentences, and about how utterances of them can and cannot be understood. To make room for such a view we need to make sense of there being *objects* for such pieces of speaker's knowledge to be about. And yet if the psychological states that determine *the facts* about one's language also *constitute* knowledge of language, how can they also be answerable to those facts as knowledge demands? Can there be something independent of the speaker for the speaker's intuitions to get right or wrong? Call this the problem of the Missing Object of Knowledge. It is about the putative object of knowledge our linguistic intuitions are thought to concern.

Speakers are authoritative, not infallible, in their native speaker intuitions. This authoritative knowledge is special because, although it purports to be about a range

<sup>13</sup> See Chomsky, 1987.

<sup>14</sup> See John Collins, 'Faculty Disputes' for a clear account of this point.

<sup>15</sup> A further question which I cannot address here is whether a speaker arrives at knowledge of the meaning of sentences or simply uses knowledge of word meaning and knowledge of syntax to constrain understanding of an utterance of that sentence in context. See 'The Distinction Between Semantics and Pragmatics'.

of objective facts (—the linguistic facts about one’s language—), it is based on nothing more than what one takes one’s words to mean, and which strings one takes to be grammatical. The problem, then, is to show how there can be a genuine subject matter for these judgements to concern—a range of objective facts about one’s language—while at the same time accommodating authoritative knowledge of them. The objectivity of linguistics requires there to be objective facts to which a speaker’s linguistic intuitions are answerable—there should be an intelligible gap between linguistic facts and our opinions about them. However, first-person authority requires the linguistic facts to be, pretty much, as we take them to be, for our linguistic intuitions to be largely correct, thus minimizing the gap between opinion and fact. This tension between the objectivity of linguistic facts and first-person authority of linguistic knowledge is a problem for the Cognitive Conception. Chomsky tries to reduce it by giving up claims to knowledge. But this is to deny not just the problem but the phenomena that give rise to it. In what follows, I will argue for a Cognitive Conception of language and knowledge of language that offers a solution to the Missing Object of Knowledge Problem, thereby safeguarding a genuine notion of speaker’s knowledge.

Speakers’ intuitions deliver information about the meaning of words and acceptability, or grammaticality, of word strings. We hear certain word strings as sentences, and where strings are ambiguous, what we hear is the utterance of one or other sentence.<sup>16</sup> We do not hear both readings of a string at once, or something neutral between the two. Our hearing a string as structured is why we hear it as the utterance of one sentence and not another. To do this we seem to draw on knowledge of word meaning and knowledge of grammar to extract key linguistic information from the string of words we hear uttered: we hear and interpret it as we do because of the joint exercise of these two kinds of knowledge. And although the product is experienced as a unified, conscious experience of what was said, we are not obliged to suppose that these two kinds of knowledge target the same aspects of linguistic reality. There is no reason to suppose there is a single locus of linguistic significance corresponding to what we are aware of in conscious speech. Meaning and structure play a joint role in creating linguistic objects but may occupy different locations in the world or the mind: they may be properties of quite different things.

There is an interesting puzzle when we compare our understanding of the notions of meaning and knowledge of meaning with our understand of the notions of grammar and our knowledge of grammar. In the case of grammar we know a great deal and we have developed well-attested theories of the syntax of natural languages that make detailed and specific claims about the structure of language. However, when we turn to our *knowledge* of syntax things are much less clear. Can speakers

<sup>16</sup> Strictly speaking, sentences are not ambiguous, only strings are ambiguous. A sentence has a structural organization, its constituents stand in grammatical relations to one another, it can be interpreted in a certain way. A string of words can have more than one internal organization, its elements can stand in different grammatical relations to one another and it can support different interpretations. Sentences type utterances. Utterances have tokens but sentences do not. Utterances can be ambiguous, sentences cannot.

be credited with tacit knowledge of the syntactic properties described by theories of their language? What is the nature of this knowledge? Despite well-advanced theories of syntax, accounts of our knowledge of syntax are problematic and controversial. However, when it comes to word meaning it is just the reverse. Our knowledge of word meaning is not in doubt: it is far from controversial to say that we know what a word means. If I ask you whether you know what ‘pleached’ means, you can immediately tell me. However, we have very little idea what a correct theory of word meaning would look like. We are not even sure what materials or primitives a theory of word meaning should employ. This puzzle alone should make us wonder whether meaning properties and syntactic properties are properties of the same thing.

Returning to Dummett’s way of posing the problem, it is important to understand how knowledge of language *effects* the difference between hearing speech sounds as noises and hearing them as meaningful speech. For we have seen that there are two very different directions of explanation. Does that knowledge equip us to speak and understand *the* language spoken by those around us? Or does the presence of this kind of knowledge in human minds explain the very existence of language: is it knowledge in the minds of individual speakers that makes language *language*? As we shall see, this view comports best with empirical findings. But a fully satisfying account must meet several desiderata. We must respect the phenomenological datum that we hear more in the speech sounds when we have learned a language. We need an account of what knowledge of language gives us knowledge *of*. We need to respect the differences between knowledge of word meaning and knowledge of grammar. And we need to explain how our possessing such knowledge equips us to speak our minds and know the minds of others?

Let us start with our ability to experience sounds as speech. For before we can even talk about words, grammatical relations, and sentences we have to remember that speech episodes start as mere encounters with sounds, and that sounds by themselves are not identical with words, grammatical structures or sentences. The acoustic properties of speech and the linguistic material hearers perceive in it are not so easily aligned. Finding words in a sound stream is difficult for a learner of a foreign language. Grammatical relations occur in the gaps between the words in sentences. The syntactic arrangement of expressions is hierarchical and not linear or temporal. So the gap between words or syntax and sounds is vast. To understand how the gap is bridged we must begin by asking, as Dummett does, what gives the sounds people utter their linguistic form and character: the linguistic form and character they have *for us*.

### 37.4 EXPERIENCING SOUNDS AS LANGUAGE

Acquisition of a first language should not be modelled on second language learning. Nonetheless, the experience of listening to speakers of an utterly foreign language is instructive. It reminds us that at one level all that goes on in the environment when we speak is the issuing of certain sounds. Even though this is not how speakers

perceive the sounds of their language: to others they are just sounds. And far from being a philosophical distortion of the phenomena: we can actually perceive things this way in the foreign case. When listening to people speaking French or Japanese all that some people hear is a continuous sound stream with occasional pauses. Others will hear people engaging in intelligible speech. Speakers talk in a continuous flow of sounds, joining one sound to another and breaking off only when they need breath. So we have difficulty even recognizing the word boundaries in a foreign language since they do not correlate with breaks in the acoustic signal. Several properties we appear to perceive in speech do not correlate with acoustic properties of speech sounds at all. (We shall consider these below at 37.5.) So how does knowing a language enable us to perceive more in the speech sounds we listen to? How does knowing French or Japanese help us hear or see something that others miss?<sup>17</sup> In the case of a second language, we have good answers to questions (i) and (ii) in 37.2 about our knowledge of language. We can set out what the speaker has come to know by charting the explicit linguistic training he or she received in their laborious efforts to learn the language. But things are very different with a first language and these answers won't do. For a start, there is little evidence of explicit training in syntax. Some language learners receive no training of any kind at all. Secondly, it is obvious that in learning a second language one is relying on the fact that one already speaks a language and much of the learning is at first a matter of translation between one language and the other. It cannot be like this for the infant acquiring its first language.<sup>18</sup>

Although much more is experienced in speech sounds by those who know the language than by those who don't, the patterns the former 'perceive *in*' the sounds produced by speakers cannot be identified with, or taken to inhere in, the physical properties of the sound waves hearers encounter. The conscious auditory images of words, phrases or sentences, are inner mental objects, 'not physical phenomena inhering in sound waves' (Harris and Lindsay, 2003, p. 203). The same is true of most of the linguistic information we glean from speech. Such information consists of phonological patterns, word patterns, grammatical patterns, including hierarchical configurations of elements and dependencies between them, word meanings and logical scope. Information about such phenomena must be drawn upon in order to hear the sound stream as an unfolding sequence of meaningful words that makes up a particular sentence. But how can all the relevant information be extracted from

<sup>17</sup> The seeing case is subsidiary if it is reading the language, though in the case of congenitally deaf language learners who use signing to communicate in the language, seeing may be the right modality from which to extent information about sentences.

<sup>18</sup> The idea that *we* 'translate' into thought is a non-starter. There is no such conscious experience in the home language: speech perception is fast and automatic and requires no conscious conversion of sounds first heard without any meaning. Nor do we have transparent access to the form and character of thought, so it is not clear how we would go about correlating an item of thought with a linguistic item. None of this rules out the idea of mechanisms that automatically map sounds or signs of the public medium into a language of thought. But explaining how this works is no easier than explaining how our knowledge of language makes possible our hearing of significant speech in the sounds we encounter. The two accounts may end up as notational variants.

the speech sounds one encounters? What is there in the much richer experience of one who knows the language is not *there* in the sounds emitted by the speaker: 'linguistic information is projected by means of articulations but is not embodied in them' (*Ibid.*, p. 203). The linguistic information read into, or onto, those sounds is simply part of 'the specifically human way with sounds' (*Ibid.*, p. 203).

However, according to John McDowell, what we perceive in speech, in virtue of having learned a language, is something lying open to view of the surface of linguistic practice. These are linguistic phenomena that we come to perceive as a result of acquiring knowledge of the language: a range of facts that were not previously perceptible and that come into view as we find our way into the language. But how does what goes on in us when we learn a language enable us to perceive these supposedly linguistic facts? McDowell thinks there is no way to explain language acquisition, or to account for what we know in terms available to people outside the practice. However, despite McDowell's self-imposed quietist embargo on giving explanations, he cannot resist giving hints about how the process might go. The acquisition of linguistic knowledge, he tells us, is a matter of drilling in certain behavioural routines, using sentences at first without fully understanding them.<sup>19</sup> So the question for him is:

How can drilling in a behavioural repertoire extend one's perceptual capacities—cause one to be directly aware of facts of which one would not otherwise be aware? (McDowell, 1998a, p. 333)

McDowell admits that this is a very difficult question to answer but one could be forgiven for thinking that, on the contrary, it was impossible because it is the wrong question. No behavioural drilling could extend one's capacities for perceptual experience of meaningful sentences, nor could a repertoire of behavioural routines for the use of particular sentences yield any insight into why the learner extends his use and understanding of utterances, in potentially infinitely many ways, to some arrangements of words but not others.<sup>20</sup> McDowell suggests that we cannot explain how people arrive at their knowledge of sentence meaning in indefinitely many cases, but we can describe what contents people would hear assertoric utterances of those sentences as expressing.<sup>21</sup> We do this by means of a Tarskian truth-theory for the

<sup>19</sup> See McDowell, 'Meaning, Communication and Knowledge', pp. 47–9.

<sup>20</sup> Chomsky's devastating 1959 critique of Skinner's account of verbal behaviour as conditioned learning showed how little of what was needed could be accomplished by a behaviourist account of language acquisition. Readers are advised to consult the details.

<sup>21</sup> See McDowell, 1998, p. 180. Note he does say, 'The ability to comprehend heard speech is an information-processing capacity, and the theory would describe it by articulating in detail the relation, which defines the capacity, between input information and output information', p. 179. However, to describe is not to explain and whatever mechanisms are responsible for this information processing capacity they do not amount to dealings with content. Later in the same article, McDowell tells us that; 'There is no merit in a conception of the mind that permits us to speculate about its states, conceived as states of a hypothesised mechanism, with a breezy lack of concern for facts about explicit awareness.' This way of thinking is described as 'philistine' and leads to darkness within. But one may wonder where the charge of philistinism belongs if there are explanations of linguistic intuitions to be had, by reference to underlying states, about why states of explicit awareness take the linguistic form and character they do and we refuse to avail ourselves of them and opt instead for mystery.

language of the linguistic community. The T-sentences will be homophonic but there will be infinitely many of them and so the question remains how speakers in a linguistic practice succeed in conforming to the theory and arrive at indefinitely many pieces of knowledge described by the theory. The theory will identify constituent structure in sentences, and see the meaning of each sentence as having been assembled by the legitimate syntactic combining of meaningful parts. But talk of meaningful parts and wholes will be confined to the workings of the theory. Whenever we have such knowledge, however acquired, we simply perceive the meanings that lie open to view on the surface of our practice.

What precisely is the nature of the linguistic facts that become directly perceptually available to us as a result of learning the language? And what relations do such facts bear to the acoustic properties of the sounds uttered? McDowell says we hear the meaning in people's words? And while there is an authentic phenomenological insight here, we do hear meaning in *words*, but how do we hear *words* and *sentences* in the sounds people utter? It is necessary to identify the linguistic items we hear when we hear sounds as linguistically meaningful. Yet nothing in the phenomenological claim settles the issue of the location of the linguistic properties linguists describe. McDowell thinks of the words as out there, and he even supposes the same about syntax, as if it were to be found in the facts about behaviour we are presented with. He talks of a 'match between theoretical syntax and actual utterance-events' (1998, p. 146). The hard physical facts are thought to constrain 'the structural properties of physical utterance-events that permit the language to be given a syntactic description' (*Ibid.*). He tells us that a relation 'must hold between the structures assigned to sentences by the syntax with which the theory operates . . . and configurations observable in physical utterance-events' (*Ibid.*). But the overwhelming data from linguistics teach us to expect no such match between the syntax our best theories assign to sentences and the physical arrangements of uttered sounds (see Section 37.5 below). Words are only marginally easier to identify with particular ranges of sound. But if we reject McDowell's talk of a match between syntax and utterance-events, and reject the account of behavioural drilling in the use of whole, but at first unintelligible, sentences,<sup>22</sup> how do we come to have knowledge of the public language, and how does it help us secure the transition from facts about sound to facts about linguistic meaning?

According to McDowell, speaker-hearers directly perceive the linguistic facts displayed in the practice, however to know *the* language thereby displayed both speaker hearers and theorists of that language need a way to identify the relevant range of facts that belong to *the* language. What is that range of facts, and how do the relevant linguistic properties of words and sentences help speaker-hearers to identify them?

<sup>22</sup> This idea seems empirically flawed too. Children between the ages of 12 and 20 months are at the one-word and two-word stage. They show an understanding of these words and acquire many before they are able to use sentences. The use of sentences suddenly occurs at the syntax spurt at around 20–24 months.

To say more about what enables us to experience a certain range of human sounds as part of a *language* we first need to say something more specific about what it is for something to be a *language*: to say more about what makes language *language*.

At the very least, by *language* we mean a system for building complex structures combinatorially out of a stock of meaningful items—words or signs (in the case of a sign language)—and non-meaningful items—morpheme endings, phonemes, specific speech gestures. The structures humans build are recursive and potentially infinite in number despite employing only finitely many discrete elements and rules to generate them. While these syntactically complex combinatorial structures vary across human languages they do so within strict parametric limits, and have many properties in common.

The discrete infinity this system provides is built out of lexical items from a stock of between 50,000 and 120,000 items and appears to be species specific. Apes trained to use sign language or a keyboard may develop between a hundred and two hundred items at most and may combine them in two and three word combinations, without extending naturally and spontaneously to larger and larger combinations, as human children do. Moreover, chimpanzees do not go on acquiring vocabulary at the fast rate of the infant, nor is there sufficient evidence (in so far as they show elementary combining of symbols) of a capacity for *syntactic* combining of linguistic elements.

Whatever we should say of apes' word or symbol combining there is no evidence of the recursive structure building that characterizes the human linguistic capacity. Many animals show only sequence and signalling in fixed, alternating and repeated patterns. However, syntax requires hierarchical structuring. The requests, commands and assertions children make are syntactically structured. According to Hauser, Chomsky and Fitch (2002) the language faculty, narrowly conceived (FLN), is what equips us with the capacity for recursive structure building and should be distinguished from the faculty of language broadly conceived (FLB), involved in the totality of our linguistic communication. Hauser, Chomsky, and Fitch say:

most, if not all, of FLB is based on mechanisms shared with nonhuman animals. . . In contrast, we suggest that FLN—the computational mechanism of recursion—is recently evolved and unique to our species (p. 1573). We propose in this hypothesis that FLN comprises only the core computational mechanisms of recursion as they appear in narrow syntax and the mappings to the interfaces (p. 1573), the interfaces with mechanisms for speech perception and production, and the cognitive system for conceptual thinking and intention (p. 1573)

Could animals recruit recursion from elsewhere? There is evidence that many animals have numerosity and even some arithmetic but this is usually limited to between 4 and 6 items. Even monkeys taught to count up to 4 and do some subtraction and addition, do not naturally progress to 5 but have to be taught to deal with this new number all over again through repeated trials. They show no ability, of the sort the child has, to extend the series indefinitely; that is, they have no concept of successor.

Non-human primates do not have the capacity for language just characterized. What they do have is a variety of systems for (non-linguistic) communication. We

are unique both in our handling of recursive structure and of our capacity to encode this in a limited range of speech sounds. And this ability to apprehend and integrate, so rapidly, the phonetic, syntactic and semantic information in virtue of which sounds events are recognized as linguistically significant speech is a staggering achievement and requires explanation. Nothing we are offered in terms of drilling in behavioural routines, or of matches between syntax and utterance-events begins to touch this problem or even to indicate how what is going on amounts to *linguistic* activity.

The perceptual experience of one who has learned the language is rich and complex but McDowell is missing much of the complexity involved, even in talk of our use of sentences. A sentence is a linguistic structured string of meaningful items (some with endings) that play certain grammatical roles in the sentence. The syntactic configurations will constrain in various ways the semantic interpretations we can give to a sentence, and how the sentence uttered will be perceived. Syntax plays a vital role in just which sentence is perceived and in what we take the uttered sounds to mean. How do we become apprised of these syntactic properties, how do we take them into account, and how does our knowledge of the language gained through participation in a common practice connect us with them?

What enables a child who at first hears only sounds to come to hear them as meaningful speech? How does it make the transition from one who lacks this capacity to someone who has it? And if it is distinctive linguistic knowledge, along with other cognitive, perceptual and motor skills, that explains our capacities to produce and comprehend meaningful speech, then a full explanation demands an account of the nature of that linguistic knowledge and what precisely it gives us knowledge *of*. We have been looking in vein for some way for the supposedly recruited public language to be used to make the difference between hearing what people say in speaking the language and merely hearing human noises but so far we have made no progress. Exploring an idea he ultimately rejects, Michael Dummett suggests:

The natural answer is what makes the difference is the fact that both speakers understand or know the language. Each has, so to speak, the same piece of internal (mental) equipment, which enables each to interpret the utterances of the other as an expression of thought, and to convert his own thoughts into sentences that the other can likewise understand. It thus seems as though the key to the explanation of the expressive power which makes a language a language is an individual speaker's mastery of the language; and this mastery . . . requires the notion of knowledge for its explication. (Dummett, 1978, p. 97)

Dummett's suggestion looks at first to be offering a Cognitive Conception of language and our knowledge of language, but this is not what he intends. It is the language speaker-hearers know and share that explains their ability to express and convey indefinitely many things. Talk of internal equipment may be important for explaining how people keep in touch with, or keep track of, the language, but on this view it is still *the language itself* that is mastered and that affords the possibility of indefinitely many meaning possibilities. It is through mastery of the public language, albeit by means of a piece of internal equipment, (or for Dummett, the speaker's having certain practical and theoretical abilities) that the individual is

able to take advantage of the expressive possibilities the language provides. Despite talk of internal mechanisms, speakers and hearers are still here seen as reaching for the same language, and the worry for Dummett is just how they can each be sure that the hypothesized mechanisms relate them to the same words and meanings. This picture has a powerful grip on our imaginations. What we are calling *the* language and *the* meanings of words is supposed to depend on the workings of a common language mastered by several individuals. But just how do they master the same public language? What is its nature and what does it provide? them with knowledge of? An what enables individuals to acquire knowledge of precisely that language and put it to use themselves?

### 37.5 ACQUIRING KNOWLEDGE OF LANGUAGE

Given the complexity of language and our knowledge of it, there is little or no reason to suppose that children learn language by inductively generalizing from their linguistic experience. Such inductive explanations would have to show how exposure to others' use of language, first experienced as sound, leads the child to establish a highly complex and elaborate system for producing and understanding speech. Even if we had a ready explanation of how they recognize words from the sound stream itself, inductive explanations of their knowledge of syntax are implausible. Why, for instance, do learners never treat (1) and (2) as similar, despite having the same number of words of the same grammatical category in the same linear order?<sup>23</sup>

- (1) John is easy to please
- (2) John is eager to please

The fact that they don't can be shown by learners accepting the re-arrangement of (1) as (3), but never attempting the re-arrangement of (2) as (4).

- (3) It is easy to please John
- (4) \*It is eager to please John

Also, why is it that no speaker interprets (8) in relation to (7) in the same way they interpret (6) in relation to (5)? Instead speakers suppose (8) means that John is too clever for anyone to catch him.

- (5) John ate an apple
- (6) John ate
- (7) John is too clever to catch Peter
- (8) John is too clever to catch

<sup>23</sup> The examples come from Chomsky (1965). [left justify]

And why do no children attempt at first to form questions from statements by reversing the first two words of the sentence as in (11) when they have seen this done in cases like (9) and (10)?

- (9) Peter is asleep
- (10) Is Peter asleep?
- (11) \*Man the is laughing?

Why *don't* they use the 'is' as the point around which to attempt the needed inversion? If they formed the generalization that the first 'is' mattered we would expect children to say (13) as the question form of (12). But they never do.

- (12) The man who is laughing is silly
- (13) \*Is the man who laughing is silly?

We need to explain why *no* children attempt these most natural inductive generalizations. The negative facts about what speakers do *not* do are as important to explain here as the positive facts, if we are to formulate the correct explanatory generalizations. Notice too that children are not explicitly taught these regularities and yet as speaker-hearers they all conform to structure-dependent generalizations that must be given in terms of, not surface arrangements, but an underlying level of logical form (see 'Logical Form and LF'). How would one learn generalizations about logical form from inductive generalizations over the surface forms from which they can depart so significantly?

Moreover, there are properties of grammatical structure common to all human languages, exemplified in the sentences produced by children around the ages of three or four years of age. Consider the Binding Principles that explain why we treat the pronoun 'him' and proper name 'John' as disjoint in reference in (14) and the reflexive pronoun or anaphor 'himself' as referentially dependent on 'John' in (15). By contrast, it is left open in (16) that 'John' and 'him' could take the same reference. That the position is not referentially dependent on 'John' can be seen in the unacceptable (17) where the anaphor cannot depend for its referent on 'John'.

- (14) [John shaved him]
- (15) [John shaved himself]
- (16) [[John's mother] shaved him]
- (17) \*[[John's mother] shaved himself]
- (18) Peter said that [John shaved him]
- (19) Peter said that [John shaved himself]

The asymmetric relations between the subject 'John' in the noun phrase (NP) and the object 'him' or 'himself' in the verb phrase (VP) in (14) and (15) ensures a syntactic configuration known as *c-command*.<sup>24</sup> An item in the syntax can referentially bind another when it *c-commands* it. Universal Principles of Binding tell us that:

<sup>24</sup> In tree-geometric terms, an expression  $\alpha$  *c-commands* an expression  $\beta$  when the first branching node dominating  $\alpha$  dominates  $\beta$ . For more on *c-command* and binding see Chomsky, 1995.

- (20) Principle A: An anaphor is bound in its domain  
 (21) Principle B: A pronoun is free in its domain

The domain of anaphors ('himself', 'herself', 'each other') and pronouns ('him' 'she', 'them') is indicated by bracketing in the above examples. In (16) and (17), the possessive noun 'John's' does not c-command the pronoun or the anaphor and cannot bind them. Notice, that a pronoun must not be 'too near' to its referential antecedent and an anaphor must not be 'too far away'. Hence, 'him' cannot be referentially bound by 'John' in (18) and 'himself' cannot be referentially bound by 'Peter' in (19). The semantic facts about how we use and understand (14) to (19) referentially are known to all speakers. They are also determined by configurational facts about syntactic structure. Were sentences not structured in an asymmetric and hierarchical NP-VP structure but simply linearly ordered as Noun-Verb-Noun we could not explain the abundant data: data that obtain in all human languages regardless of the surface order of subjects and verbs. Children all over the world in their early uses of sentences conform to these universal properties of structure despite the differences in the their experience, intelligence, and background culture.<sup>25</sup> How do all of them succeed, in pretty much the same time course, in closing the gap between what experience provides them with and the vastly greater knowledge they end up having about the speech sounds around them? The absence of any adequate explanation of how they do this on the basis of inductive learning strongly indicates that part of the cognitive system must be dedicated to arriving at this highly intricate form of knowledge on the basis of exposure to random and limited amounts of data.<sup>26</sup> As Noam Chomsky puts it:

The problem, then, is to determine the innate endowment that serves to bridge the gap between experience and knowledge attained. (Chomsky, 1986)

The special, dedicated component that enables the normally functioning infant to acquire any of the possible human languages, on the basis of exposure to a course of experience, is the Language Faculty (narrowly described).<sup>27</sup> As mentioned above, it is a component of the mind/brain with which humans are uniquely endowed. And so far this is the only serious proposal we have about how humans acquire grammatical knowledge of their first language. The forlorn idea that we learn to do all this by analogy with the repetitious learning of a manual skill is a non-starter and does not merit serious discussion. There is no evidence that practice takes place or that mistakes of the kind expected in such training actually occur. What needs to be explained is why

<sup>25</sup> See Stephen Crain, 'Language Acquisition in the Absence of Experience' Behavioural and Brain Sciences, 1991.

<sup>26</sup> Limited success has been achieved in capturing individual patterns of structure or particular grammatical regularities through computational techniques for capturing statistical similarities across a large corpus (see Elman, *et al.*). But it is the bewildering number of patterns and highly interdependent regularities that have to be captured and so far the only way to encode such a system is through a finite set of highly interactive grammatical principles governing those structures that permit a certain flexibility within a strict parametric range.

<sup>27</sup> Empirical support for this view is to be found Chomsky, 1955, 1965, 1980, 1986.

a speaker, who has had exposure to a limited set of utterances, can produce indefinitely many new ones and ‘distinguish a certain set of “grammatical” utterances, among utterances [of sentences] that he has never heard and might never produce’ and in doing so ‘projects his past linguistic experience to include certain new utterances and excludes others’. (Chomsky, 1955, p. 61) Conformity to innate principles of grammar would explain how we do this, but not much else would.

Humans are the only creatures able to engage in such *linguistic* communication, because they alone are innately endowed with a language faculty that gives them the capacity to acquire any one of the possible human languages on exposure to a particular course of experience).<sup>28,29</sup> These findings strongly suggest that the existence of languages is due to the existence of humans with minds like ours. If this is the case, language would not exist without us. However, the dependence of language on minds is not enough to establish claims about the location of linguistic phenomena and some philosophers remain wedded to the idea that languages exist external to, or even independently of, the minds of human language users, supposing that it is for these genetic reasons, or for some other non-linguistic reasons, that we may be the only creatures capable of accessing and exploiting the intricacies of languages.

Such philosophers still need to accommodate the facts about language that make language *language*. The essential properties of a language, those on which its identity as language depends—properties of syntactic structure, case and theta-role assignment, c-command, binding—are not brute properties of the physical sounds uttered, nor are they all properties experienced in our conscious reception of speech; although they can have precisely predictable effects on speakers’ conscious linguistic intuitions.

An argument can now be given for a Cognitivist Conception of language and knowledge of language as follows. Structure is essential to language, to the very identity of sentences that count as part of our language, but we cannot locate these structural properties in the physical or phenomenal world of speaker-hearers. The other option is in a part of the mind, a special cognitive component, dedicated to handling language. According to the Cognitive Conception, the structural properties in question are sub-personally represented in the language faculty: a species-specific and largely innate component of the human mind. Note that these mental representations of linguistic structure play a content-involving causal role in shaping the first-person linguistic experience of the speaker-hearer, in giving one the linguistic experience one has. We can formulate hypotheses about the precise character of

<sup>28</sup> Speakers come equipped with universal grammar, which prescribes a certain structural organization but leaves open certain permitted variations within strict parametric ranges. For example, Italians can use null subject sentences: sentences with a phonetically null subject position (e.g. *\_\_ha parlata*/Gianni ha parlata), but English speakers cannot. Part of a child’s acquisition and in which it does need information from its environment is to help it set the parameter for a null or non-null subject for its language. Much of syntactic acquisition is parameter setting, triggered by the child’s linguistic environment.

<sup>29</sup> This is not to deny the fact that non-human animals have rich and complex systems of communication. It is simply to point out that they are non-linguistic. Humans are also capable of a good deal of non-linguistic communication, from a nod and a wink to raising one’s eyebrows and looking at one’s watch to signal to another it is time to go.

those syntactic representations by observing their impact on the form and character of a speaker-hearer's linguistic intuitions. What must her internalized grammar be like, we can ask, in order for her to find these arrangements of words acceptable but not those; for her to be able to interpret a sentence in this way but not in that. To arrive at specific hypotheses about the internalized grammar we reason counterfactually: had the grammar been different, had it not respected a particular constraint then it would have been possible to hear certain utterances differently. For example, take the following ambiguous string:

(22) I saw her duck and swallow.

We can hear it in two different ways according to the readings shown in the following syntactic disambiguations:

(22a) I saw her [<sub>NP</sub> [<sub>N</sub>duck] and [<sub>N</sub>swallow]]

(22b) I saw her [<sub>VP</sub> [<sub>V</sub>duck] and [<sub>V</sub>swallow]]

But the question is why can't we hear it as four-ways ambiguous? We cannot even have the conscious experience of either of the two other readings. As theorists we can see they are logically possible:

(22c) ?I saw her [[<sub>N</sub>duck] and [<sub>V</sub>swallow]]

(22d) ?I saw her [[<sub>V</sub>duck] and [<sub>N</sub>swallow]]

but as speaker-hearers we cannot even consciously experience these ways of construing (22).<sup>30</sup> Why not? The explanation is that the internalized grammar that shapes and conditions our conscious experience of speech respects a co-ordination constraint. The co-ordination constraint says that we can conjoin expressions of the same syntactic category: e.g. NP and NP, VP and VP, PP and PP, S and S, etc, but not PP and VP, or VP and NP.<sup>31</sup> This is why we do not accept as grammatical, or intelligible, putative sentences like:

(23) \*She coughed and the boy.

There is no reason why our grammars had to obey a co-ordination constraint but evidence from the conscious experience of speaker-hearers, elicited by their linguistic intuitions, confirm that the grammars that make up our linguistic competences do satisfy such a constraint. The methodology here shows why the conscious experiences of the language user—immediate and authoritative knowledge of language in the form of linguistic intuitions—do play a role in constructing explanatory adequate theories of our linguistic, and in particular, syntactic competence, pace Chomsky.<sup>32</sup>

Notice that the underlying states of our linguistic system are *content-involving* states. They *represent* the syntactic structure of sentences. They are not merely syntactic: formal symbols whose syntactic shape has a causal effect on our conscious

<sup>30</sup> My thanks to Emma Borg for pointing out this example.

<sup>31</sup> We may need to consider the syntactic category of a predicate to accommodate more difficult cases. The details can be finessed by a proper syntactic theory.

<sup>32</sup> See Chomsky, 2000.

states. When talking about syntax, philosophers are prone to make this vehicle-content conflation about such representations. Saying that we, or more accurately our language faculty, represents syntactic structure is not the same as saying that these are just syntactic representations, as though they had to take the syntactic structure of the sentence whose structure they represent. There is no reason why a *representation* of sentence structure has itself to have the *structure represented*, any more than a red idea has to be red. It is *what* those underlying representations represent, not the *way* they represent it, that plays the crucial role in fixing the linguistic character and content we experience in perceived speech.

What alternative is there to locating the essential properties of linguistic structure that make language *language* in the mind of speakers and listeners? The prospects for the traditional view of language as a social practice look bleak. It may be said that as participants in the practice we can just tell which strings of words are grammatical and which ones are not. But our being able to tell, to perceive which sentence was uttered, already takes into account what it is in us that shapes these perceptions and gives us our way of finding the utterance intelligible, permitting some interpretations and not others, hearing ambiguous strings one way and not another. In the case of the Cognitive Conception we have detailed explanations of why we hear some arrangements of words as grammatical sentences and others as not. On the social view we have none. It is also false to say that we don't need explanations because we can just tell which sentences are grammatical.<sup>33</sup> We are not infallible. When first asked whether sentence (24) is ambiguous most speakers will say no.

(24) I almost had my wallet stolen.

However, if we point out that they could have had their wallet stolen and were trying to steal it back when they were spotted and withdrew, or were in the process of having a third party try to steal it back from the thief when the attempt was foiled, we can see that (24) would do to report what happened. The sentence is three-ways ambiguous: I almost had my wallet stolen from me/for me/by me. An inverse case is the illusion of grammatical well-formedness:

(25) Many more people have been to Paris than I have.

At first most people will treat (25) as well formed and interpretable, but further reflection will show them there is nothing it could possibly mean.<sup>34</sup> Speakers' immediate intuitions are authoritative but not infallible. They can be corrected and the speaker can come to find out more about the character of *his* or *her* linguistic system.<sup>35</sup> We

<sup>33</sup> This would be like saying we do not need a theory of gravity since if you bring me objects I will tell you which ones fall.

<sup>34</sup> My thanks to Paul Pietroski for the use of this example.

<sup>35</sup> Even if it was possible to say we don't need explanations because as participants in the practice we know which arrangements of words are grammatical, since explanations *can* be given, this would be like saying to Newton, we don't need an account of gravity because if you bring me the objects I will tell you which ones fall. I take it the moral is if genuine explanations are available we should avail ourselves of them.

need an explanation of these facts and of the mistakes we make in performance. Such an explanation draws a distinction between competence and performance.

An alternative remains. For perhaps a Platonist E-conception could treat structural properties as part of the abstract realm. But now we need some account of how these properties of abstract objects have an impact on a speaker-hearer's intuitions and how speaker-hearers conform to generalizations framed in terms of those structural properties. If languages are either social or platonic abstract objects, existing beyond individual speakers or outside space and time, there will have to be some finite fact about speakers' knowledge, or use of language, or some other property or practice, that relates them to one rather than another of these social, or abstract objects. The difficulty will be to say what makes one of these languages the actual language of a given speaker or a group, and just what is encompassed by that speaker's language. What language does the speaker speak and what is the precise extent of a speaker's language? So far we have had little by way of answers on the social E-conception of languages sustained by common practices. But Platonism provides a more robust idea of language, viewing it as a system of meaningful expressions existing independently of us: a system we must be related to in order to express and communicate our thoughts. The Platonist option makes it clear how languages can be characterized in full but it leaves the problem of what it is for speakers to use and understand these languages.<sup>36</sup>

In studying language the philosopher or linguist needs some way to delineate the whole language of a given speaker, or community, taking in the whole range of legitimate expressions in that language. This requires us to delineate all and only the meaningful and grammatical expressions of the language, even though all we have observed to date are the finitely many uses of the language speakers have displayed. How are we to extrapolate correctly beyond these observed uses to the case of as yet unused sentences, which speakers could, if presented with them, easily recognize as belonging to their language? How do we fix the full extent of a natural language?

### 37.5.1 Must the Study of Language Involve Reference to Speakers?

If we simply define languages without reference to speakers—as we do in purely formal languages—we can, using whatever formal means are at our disposal, construct infinitely many possible languages, each described in terms of an ordered pair  $L = (\Sigma, M)$  where  $\Sigma$  is the set of sentences and  $M$  is the set of meaning specifications, one for each of the sentences in  $\Sigma$ . (See M. Davies, 1981, pp. 5–6, D. Lewis 1983.) Assuming there is no ambiguity, we can treat  $L$  as a function from sentences in

<sup>36</sup> A quick way with the Platonist options would be to point out we need an internal way to recapitulate the structures of the abstract domain in order to keep our sights fixed on one of these abstract objects. So why not settle for an account of the character of a speaker's representations in order to account for the properties of the speaker's language? We could then just slough-off the abstract structures, treating them as mere projections of the speaker's inner linguistic systems.

$\Sigma$  to meaning-specifications in  $M$ , so  $L(s) = p$ , where  $p$  is the proposition that is the meaning of sentence  $s$ .

Each specified language  $L$  is conceived as an abstract object owing nothing to human nature. But the question is how (some of) these formally described languages can be the *natural languages* spoken by *actual speakers* or populations? If we assume an exhaustive specification of the infinitely many combinations of sentence-meaning pairs, we can perhaps assume that the language of a given speaker or population is among the set of abstract objects of type  $\langle \Sigma, M \rangle$ . But which of the infinitely many possible languages is the *actual language* of a given set of speakers or speaker? This is the problem of defining the *actual language relation* between a speaker (or population) and a language  $L$ .<sup>37</sup>

A language  $L$  is the actual language of a given population  $P$  iff  $R(L,P)$

The problem is to solve for  $R$ . What is the relation in question? If we can find out we will learn, it is supposed, in what way meaning and other linguistic properties supervene on the psychological states or social practices of language users.

The difficulty is that each language is itself infinitely large, permitting the construction of infinitely many well-formed meaningful sentences. But speakers will only ever produce finitely many utterances. How are we to know which of these infinitely large objects is the language spoken by a speaker, or set of speakers, if they can only ever exhibit finitely many uses of the language? The fragment they produce could be extended or enlarged in infinitely many different ways thus leaving it open which  $L$  is the language they are actually speaking.

(P) What is it for a given language  $L$  to be the actual language a person uses?

However that relation is defined we will still want an answer to the question:

( $\Lambda$ ) What enables a person to understand indefinitely many sentences of his language?

Question ( $\Lambda$ ) asks what equips us to do what we do linguistically? It is subtly different from the question:

( $\Psi$ ) How do we actually produce and comprehend utterances?

which asks how we make use of our equipment to do what we do. Attempting answers to ( $\Psi$ ) is a proper task for the psychology or psycholinguistics, while attempting answers to ( $\Lambda$ ) is the proper task for generative linguistics. The linguist's question targets a particular conception of language and knowledge of language that proves theoretically fruitful and that diverges considerably from the Platonist's conception.

<sup>37</sup> The problem was first discussed by David Lewis in response to challenge by Stephen Schiffer. I follow Schiffer, more or less in the presentation of the problem, though I reject his definitions of language that uses finite sequences of 'types of marks or sounds' as belonging to the set in the first member of the pairs. We cannot begin with sounds for the reasons that have just been rehearsed.

The Platonist philosopher of language has to say what secures the relation between a speaker and the language he speaks, where a language is itself a relation between infinitely many finite sentences and an infinite class of meanings. In that way we'll be able to say what it is for infinitely many sentences to mean something for a person since:

(S) If L is x's actual language then what a sentence means for x is what it means in L.

Fix the language the person is using, and since that language has infinitely many sentences with meanings the person has a language with infinitely many sentences that are meaningful *for him*. What of his knowledge of language? He will not be able to understand or know the meanings of any but a minority of the sentences of his language: nevertheless, if this is the language he uses then, according to Stephen Schiffer, he knows the language 'in the sense of knowing a language in which you ipso facto know the languages you use' (2005, p. 16). But what sense is that? Is knowledge of language to be equated with a person's use of language, or ability to use a language? This deflated (or deflationist) conception of linguistic knowledge converts the problem of accounting for our knowledge of language into the problem of explaining what it is for us to use a language. What is it about our use of language that secures the *actual language relation* to the abstract object L as the actual language we use?

In his paper 'Actual Language Relations', Schiffer considers a range of problems we face. The finite use of language doesn't seem to determine which infinite language we are using. For even if there was a practice of speakers uttering sentences of L with the meanings they have in L, this would only account for their conformity to L within a finite fragment of it. And given that there are infinitely many continuations of sentence-meaning pairs diverging beyond that point, all consistent with the sentence-meanings pairings in the fragment, what connects us to one rather than another of these languages? The problem is to understand how the used fragment of language can uniquely relate us to the unused part of our language. The problem is how to extend what we regard as the meanings and structures of the sentences used to the unused parts of the same language: what David Lewis calls *the meaning-without-use problem*.

A natural thought is that there are principles at work governing the grammaticality and meaning of the sentences in the used fragment that extend to sentences in the language as a whole. If we can discover the principles or rules governing the workings of the language fragment we can extrapolate to the properties of sentences in the whole of the language. To do so we need to figure out the rules of syntax and semantics that generate the sentences of the fragment that will also generate the rest of the sentences in L. This suggestion gives rise to  $L_G$ .

[ $L_G$ ] L is the actual language of a given population P iff every adequate grammar that generates the fragment used by P is a grammar of L.

Lewis opted for this solution and after abandoning a number of attempts to solve the problem of the actual language relation in terms of conventions of trust and truthfulness. It requires extrapolation from the used fragment by a grammar for the whole language that covers that fragment:

First use *somehow* determines meaning for the fragment of the language that is actually used. There are rules of syntax and semantics that generate the right sentences with the right meanings within the used fragment. These rules also generate other, longer sentences, with meanings, outside the used fragment. Use determines some meanings, those meanings determine the rules, and the rule determine the rest of the meanings. Thus use determines meaning, in part directly and in part indirectly, for the entire language. ('Meaning without Use: Reply to Hawthorne', p. 149, *italics mine*.)

But how do we know that *a* grammar that generates the sentences in the fragment is *the* grammar of the language as a whole and will continue generating only sentences of the language? Sensibly realizing that the finite facts of language use don't, and couldn't, determine the language we speak, Schiffer has come to accept that we must appeal to the linguist's conception of knowledge of language, which he describes as consisting in an internally represented grammar used in language processing. Only that will provide an explanatorily adequate grammar. The Chomskian idea is that the theorist's grammar *G* that generates *L* must be a grammar that generates the fragment used by *P* and must be a model of the internal grammar used by the speech processing mechanisms of speakers in *P* to produce and comprehend utterances in their language. The grammar (or the information it encodes) must be internally represented in the mind/brain of the speakers. Assuming that this grammar generates not only the linguistic forms and meaning of sentences in the used fragment, but also has the capacity to generate forms and meanings for all other possible expressions of our language, we can use the linguist's account of knowledge of language to connect the used and the unused part of the language so as to identify the infinitely large abstract object *L* that the speaker speaks.<sup>38</sup>

What is it for a sentence to mean something for a person, given that there are infinitely many meaningful sentences of the language? Knowing a language is thought to put infinitely many sentences at one's disposal and, if, following Schiffer, knowing a language is equated with (or deflated to) using a language, we can now ask what is it for a person *x* to use a language *L* with infinitely many sentences? Every sentence in *L* has a meaning (by the notion of language given), so if *S* means *M* in *L* and *L* is *x*'s language then *S* means *M* in *x*'s language. We relate the person to an infinity of meaningful sentences by showing why this is the language he uses. For *L* to be the language person *x* uses, is for the set of sentence-meanings pair meanings of *L* to match the infinity of linguistic forms and meanings generated by that person's internally represented grammar. The language generated by a speaker's internally represented grammar will be the language the speaker uses, the language the speaker knows 'in that sense of knowing a language in which you ipso facto know the languages you use'.<sup>39</sup>

<sup>38</sup> The search space is already constrained considerably by relating speakers to *sentences* rather than *sounds* as part of these abstract languages, and this will further constrain the semantic interpretations that human language users, whose language faculties are configured in accordance with the principles of universal grammar, will be able to give to those sentences, thus significantly reducing the possible pairings.

<sup>39</sup> Stephen Schiffer, 'Two Perspectives on Knowledge of Language', paper given at the GLOW Conference, Geneva, 2005, p. 16.

Notice that if we take this way with the actual language relation we are unable to identify natural languages *without* mention of the psychological states or practices (or knowledge) of speakers. Shouldn't we then just look for an account of a speaker's language via an account of his or her knowledge of language?

Why give such a roundabout account of our knowledge of language and of the language known? By first invoking the notion of language as an abstract object we then have to define the actual language relation to relate a person to the language L he speaks. This is done via the forms and meanings generated by the person's internal grammar. These line up with the sentence-meaning pairs of L and as a result we then claim that what the person means by his use of language is what the corresponding items mean in the abstract object L. His knowledge of L is then taken to be knowledge "in that sense of knowing a language in which you ipso facto know the language you use". But why go this roundabout route when the linguist offers a more direct account by addressing ( $\Delta$ ) without reference to (P):

- (a) What makes it the case that there are infinitely many meaningful sentences of a person's language?
- (b') It is for the person's internal grammar to assign a form and meaning to each expression that features in his use and understanding of language.

Therefore:

- (c') The language L the person uses and understands just is the set of expressions and sentences the internal equipment generates.

For the linguist, the order of determination is reversed. We fix the person's language, and the meaning of his sentences, by means of his knowledge of language. We specify the meanings and structures the person's internalized grammar assigns to his sentences. We don't start by specifying a language and try to work out the speaker's relation to it: we specify the language he knows by finding out which internalized grammar he has.

### 37.5.2 The Need to Study Speaker's Knowledge of Language

The problem created by treating languages as abstract objects—the problem of defining the actual language relation between languages and speakers—led us to conclude that the best (or perhaps the only) way to identify the actual language of a given speaker, or community of speakers, was via their knowledge of language or internalized grammar.

We have to delineate the whole language of a given speaker, or community, taking in the range of legitimate expressions in that language, and it is most likely that the only way to do this is to appeal to the knowledge speakers have of a language that equips them to produce or comprehend indefinitely many sentences. For this we need an account of what speakers know in advance of producing or encountering particular utterances that extends to new cases and enables them, effortlessly, to recognize such newly encountered expressions as part of their language.

The range of speech sounds we will produce or respond to, the meanings we attach to them, the structures we take them to have, all this depends entirely on our linguistic competence: it is the extent of our knowledge of language and so the extent of the language we know. So if we could find out precisely what range of linguistic knowledge people have and settle the boundaries of linguistic production and comprehension, we could find out what is and what is not part of *their* language. This would give us one important motive for studying speakers' knowledge of language as part of the study of language, of what makes language *language*.

A second, weaker motive for appealing to speakers' knowledge would simply be that once we have furnished an account of a given language we shall need some account of what it is for people to know that language. Notice, that this too provides an important epistemological constraint, since any account of a natural language that rendered it implausible that speakers should know the language in question would cast doubt on the cogency of the account as an account of *their* language.

A third, stronger motivation comes from viewing the relation between *language* and *knowledge of language* in terms of the latter constituting and exhausting the former. The range of speech sounds we can produce or respond to, the meanings we attach to them, the structures we take those expressions to have, depends entirely on and has no existence apart from our knowledge of language. In this way, the extent of our language is the extent of our knowledge of language. So if we could find out more precisely what range of expressions people had knowledge of—the scope of their potential linguistic production and comprehension—we would know what was and what was not part of *their* language.

Advocates of the second and third positions see it as an important adequacy constraint on a correct theoretical account of a given language that it respects the way that language is understood by its speakers. We could not ascribe grammatical or meaning properties to their language that diverged from the properties of the language as they understand it: their understanding of the language is fixed by their knowledge of meaning and grammar. The knowledge speakers have is our guide to the meanings of their words and sentences.

This view is advocated by Dummett, who claims that a theory of meaning for a language must be a theory of understanding: the understanding speakers have of that language. However Dummett stops short of the stronger, third position, advocated by Chomsky, that language is entirely constituted and exhausted by speaker's linguistic knowledge. For Dummett, the language known is a social language, shared by a community of language users. Individual speakers may have only a partial knowledge of their language, but it is important to see that there are no properties of the language that are beyond the epistemic reach of all speakers of the language.

On either Chomsky's or Dummett's view, the answer to the latter's key question: 'Is the significance of language to be explained in terms of a speaker's knowledge of his language?' (1978, p. 97) is yes. It is the only way to settle the matter. For it is what people know, linguistically speaking, that individuates the precise domain of the language they speak. The words and phrases they use, the significance they attach to them, the arrangements they take to be grammatical, all depend on their

linguistic competence and this settles the identity of their languages. The hope is that by tracing out the full extent and character of a speaker's knowledge we should be able to fashion the precise contours of his or her language.<sup>40</sup>

To pursue such a strategy we need to know what we are out to study when we aim to investigate a *speaker's knowledge*. What is the nature of speakers' linguistic knowledge and how do we set about characterizing it accurately? And once we have answered this, what would such an account of a speaker's knowledge tell us about 'what makes language *language*' as Dummett puts it?

### 37.5.3 The Elusiveness of Speakers' Knowledge

The trouble is that it is unlikely to strike us as any easier to get at the facts of a speaker's knowledge than it is to study her language directly, and, in fact, it is apt to strike us as potentially a harder problem. How do we identify it? And isn't it more likely that it is via someone's use of language that we find out about their knowledge of language? The priority may seem to be exactly the reverse of what's being proposed. To overcome this we need a clear sense of how we capture facts about what speakers know linguistically.

But now the accessible evidence seems to be the expressions speakers use, how they use them, which combinations of them they find acceptable, and so on. And this is to suppose that we can identify the items of their knowledge before discovering anything about the nature of that knowledge. Doesn't this simply return us to the original problem of delineating someone's language with the detour through knowledge of language amounting to a redundant step?

### 37.5.4 Objection to Treating the Study of Language as the Study of Speakers' Knowledge of Language

One construal of this objection is faced by Dummett in 'What Do I Know When I Know a Language?'. If the only evidence of speaker's knowledge of language is the use he makes of words and sentences, then either there is a systematic description of that use or there isn't. If not, we can learn nothing of his linguistic understanding. On the other hand, if there is a systematic description of use then we don't need to appeal to a speaker's knowledge to account for the facts of his language. This dilemma is intended to show that knowledge plays no role in accounting for language.

A number of responses are available. Dummett resists the dilemma by arguing that no matter what regularities we can point to in linguistic use (in the course of giving a systematic description of someone's linguistic behaviour or practice), we cannot

<sup>40</sup> The tendency to think language extends beyond a speaker's knowledge of it, and to see the speaker as having partial knowledge of his (the) language depends on phenomena such as the existence of books, and sign-posts and other static objects that exist independently of us. Whether all the properties we are inclined to attribute to them extend beyond us is another matter.

neglect the fact that language is a conscious rational activity, and we are aware of the regularities we make use of. After all, we don't just use language, we know how to use it, and use rests on understanding. For Dummett, it is no good just cataloguing regularities in the observable use of speech. The only regularities that can count as part of someone's language are the ones that speakers consciously choose and subscribe to. As Dummett puts it, we must distinguish,

between those regularities of which a language speaker, acting as a rational agent engaged in conscious, voluntary action, makes use from those that may be hidden from him and uncovered by a psychologist or neurologist: only those regularities of which, in speaking, he makes use characterize the language as a language (Dummett 1993, p. 104)

This response very properly recognizes a conscious and first-personal dimension to human speech. But Dummett is surely wrong, on empirical or evidential grounds alone, to insist that the only generalizations or regularities that feature as part of our language are those of which we have conscious apprehension. That would be a hopeless move in the face of the generalizations linguistic theory gives rise to, and for which there is evidence in the conscious intuitive judgements (linguistic intuitions) of speakers. Consider the following:

- (24) Mary expects to feed herself
- (25) Mary expects the woman sitting up in bed to feed herself
- (26) I wonder who Mary expects to feed herself

In (24), we take the reference of the reflexive pronoun to depend on the subject 'Mary'. The reflexive pronoun appears to depend for its reference on the nearest noun-phrase that agrees with it in number and gender, as the Binding Principles require (cf. 'Mary expects to feed her' where the non-reflexive pronoun cannot depend on the nearest noun-phrase). Thus in (25), the reflexive pronoun depends referentially on the noun-phrase, 'the woman sitting up in bed'. However, in (26), the reflexive pronoun cannot be referentially dependent on 'Mary', and speakers never assume that it is, despite the very same expression (24) being contained in (26). Speakers know these facts but they do not know how they know them? The relevant generalizations governing this and countless other examples are that 'who', the subject of 'to feed' has moved and has left an NP trace or empty category after 'expects' and that NP referentially binds 'herself' as Principle A requires. Does the speaker make use of the regularities that explain these data? She certainly conforms to them. Does she consciously know the generalizations? The answer is obviously, no. But this does not mean that the psychologist or even less plausibly, the neurologist can explain the judgements we make in (24) to (26). The relevant generalizations are expressed in linguistic terms and cannot be expressed in neurological terms. The linguistic facts in question can only be explained in terms of the hierarchical structures of strings: structures assigned by the language faculty. These syntactic structures are derived from interacting principles governing the grammatical relations of the lexical items combined: we do not carry whole sentences

around from one speech episode to another. Somehow we must deploy information about these structures, and so speakers have some way of conforming to, cognizing or heeding the principles of syntactic structure without being consciously aware of them. What is more, Dummett is mistaken in supposing that the consciously made generalizations will correctly predict the data about the speaker's own intuitions. Speakers reflecting on cases like (27) and (28) will usually predict that the use of the complementizer 'that' after 'believes' is optional. It can be put in or not.

(27) Bill believes that George is intelligent

(28) Bill believes George is intelligent

But this is not the case, as speakers themselves come to see in the unacceptability of (29):

(29) \*Hillary believes that Bill to be intelligent

This is because of Exceptional Case Marking where 'believes' assigns accusative case to 'him' in structures like (30) and must be adjacent to it:

(30) Hillary believes him to be intelligent.

On this picture, the structure of language, on which the identity of language depends, is due to the internalized grammar, not consciously accessible to the speaker, but which is part of the language faculty. On Chomsky's cognitive conception of language there is *only* the internalized system or I-language:

a person who knows a language has mastered a particular way of interpreting expressions . . . the person has acquired a generative procedure  $g$  [I-language] which associates a structural description (SD) with every possible expression. (Chomsky 1987, p. 179)

The identity of the language—the linguistic facts the theory of language aims to capture—depends on what goes on solely in the mind of language users. The cognitive states that embody this knowledge of linguistic structure becomes the domain of linguistic inquiry. A speaker's knowledge of language, conceived as some specialized internal cognitive state, becomes the object of inquiry, the proper domain of study in linguistics.

This is Chomsky's mentalist assumption about the status of the linguistic facts under discussion, which Chomsky sees as belonging wholly within the mind of the individual language user. According to this view, linguistic theory is not about linguistic behaviour or use: ' . . . linguistic theory is mentalistic, since it is concerned with discovering a mental reality underlying actual behaviour.' (Chomsky, 1965, p. 4)

This alternative response to Dummett's problem of documenting regularities in language use, relies on the argued for assumption there is no such project since the facts of natural language are not to be found in external or performance features of speakers' use, but in facts about speakers' minds—facts about the competence that consist in states of knowledge of language not always consciously accessible. These postulated states concern unobservable entities—levels of linguistic representation

postulated by the theory that lie behind the behavioural data, giving it the form and character it has (cf. McDowell's talk of facts lying open to view of the surface of practice). The postulated entities are required to explain patterns in data, but they are not reducible to it. The states of knowledge should explain how we can be apprised of certain consciously apprehensible linguistic facts and of how we conform to the generalizations framed in terms of the underlying structures.

But why should we accept this wholesale move away from linguistic behaviour and use in linguistic theory and turn towards a cognitive domain of in-head psychological states?

Notice that Chomsky makes little of Dummett's appeal to language as a conscious, rational activity. According to Chomsky, 'consciousness forms a scattered and largely irrelevant subpart of the domain of cognitive states', it plays no role as far as the study of language is concerned. But this has been contested in examples like (22) above where data about speaker's conscious and immediate intuitions provides evidential grounds about how speakers can and cannot hear certain strings as structured, from which we draw conclusions about the underlying grammars. So we must modify Chomsky's picture to accommodate such phenomena. However, Chomsky's view of language as 'a system represented in the mind/brain of a particular individual' does provide a useful corrective to Dummett's view since it is not just (and sometimes not even) what the speaker consciously selects as regularities that constitutes his language. What matters is the phonological, lexical, and syntactic information mentally represented in the mind of the speaker. For Chomsky, this represented information exhausts the reality of language: 'language has no objective existence apart from its mental representation' (1972, p. 169, fn.). Here the concern is not with sets of sentences, nor with speech production and comprehension, or any other performance-related phenomena, but with the knowledge of language that makes all such phenomena possible.<sup>41</sup> Chomsky sees languages as epiphenomena. Linguistic theories are the study of I-language (internal, individual, functions in intension). It is the mental mechanisms responsible for our linguistic capacities that interest Chomsky, rather than the products of this capacity. For him, it is the particular configurations of the language faculty in each of us that determine the languages we speak: we can read off the language, or better, structure of the language, from the set of structures and meanings the I-language provides.

On this conception language is firmly located within the language faculty. Such a conception leaves no room to accommodate the *person's* knowledge of the language he or she speaks. The account is strictly sub-personal, for facts about the internal configurations or cognitive mechanisms are not *first-personally* accessible. And yet we have first-personal and authoritative knowledge of what our words mean and which arrangements of those words are grammatical or deviant. So we must augment Chomsky's story to make room for a notion of a speaker's conscious and

<sup>41</sup> Language—a matter of linguistic competence—should be distinguished from speech—a matter of linguistic performance. Competence is just one cognitive factor among others responsible for speech production and comprehension. Other factors include memory, attention, perception, etc.

authoritative knowledge, as well as saying what it gives us knowledge of. In one sense Chomsky's cognitive conception of language may relieve us of the problem of relating speakers to the actual languages they speak; their languages are finite states of their mind/brains. But it doesn't relieve us of the epistemological burden of saying what it is for *a person* to know a language. (Waiting in the wings is the further problem of how we understand *other people's* languages, or what they are saying.) Must we accept even a modified Chomskian account of where our knowledge of language—and indeed the facts of our language—reside?

The move to mentalism in linguistic theory is based on powerful arguments designed to counter the mere suggestion that appeals to speakers' knowledge can be based solely on a characterization of language arrived at through the study of speakers' uses of their languages. The counter-arguments trade on a battery of empirical considerations about the impossibility of finding linguistic facts on the surface of speech.

Summarizing countless many pieces of empirical evidence of the sort mentioned above we see that:

the crucial properties of sentences are not revealed by thinking of them as they are outwardly presented to us, namely as strings of signs, but rather by their unobservable grammatical structure. (Higginbotham, 1991, p. 555)

### 37.5.5 The Elusiveness of a Speaker's Language

These counter-arguments enable us to reject the all too hastily made assumption that the facts about speakers' languages are uncontroversially available to us as theorists. It is just this assumption that we have been questioning all along. What we want to know is how, from observing finite stretches of people's linguistic behaviour, characterized (at first) in non-question begging physical terms, as mere sounds or gestures, we can extract an account of the linguistic significance of those exchanges, revealing the meanings people attach to their expressions, the structures they assign them, leading us to extrapolate safely beyond any point so far reached, in just the right ways to compound words into sentences of their languages. As yet, we have seen no way to do that until we know something of what speakers know about the linguistic forms and meanings they attach to the sounds they use and respond to and what gives rise to this knowledge.

The Chomskian arguments also counter the McDowellian alternative: that of phenomenological presented linguistic facts, taken to be directly perceptually available on the surface of practice—though audible only to those who are part of the practice. No one disputes the phenomenological datum that in a language we understand we hear people's words as meaningful. But what we can and cannot hear, and why only these things, needs to be explained. As perceivers of meaningful speech, how do we arrive at these specific bits of knowledge of novel sentences in countless many cases. Part of that perceptual experience depends on one's knowledge of, and sensitivity to, syntactic structure described at a level remote from what is

phenomenological accessible on the surface. The only empirically supportable—that is descriptively adequate—theories of grammar for natural languages are those that postulate underlying, or unobservable, structure as *the* structure of a sentence. But in doing so, we are postulating an structure that does not occur in the sounds or marks that make up the external or phenomenally accessible properties of speech. Just where, then, are we to locate the levels of linguistic structure that linguistic theory postulate to explain the patterns in the data? The only place is in the mind of the speaker: syntactic structure is the structure a speaker, or rather her linguistic system, imposes on or assigns to the sounds and signs she encounters in order to hear those sounds *as* meaningful sentences. Words are not words until the noises, or marks, by which we indicate them, are seen as carrying their full freight of semantic, syntactic and phonological properties. The same set of sounds can count as different linguistic expressions, belonging to different languages, or different sentences in the same language, depending on what structure the speaker/hearer assigns to them. It is the language faculty that assigns structure to the sounds people hear, it is the language faculty's contribution to speech events that makes up a large part of the domain of linguistic inquiry.

### 37.6 CAN MENTALISM BE RESISTED?

Despite overwhelming empirical evidence, mentalism in linguistics has been resisted and there are two sources of pressure. The first comes from Quine when considering the choice of grammar to characterize a speaker's language. This is the issue of the indeterminacy of grammar. But the indeterminacy in syntax is not so troubling, Quine claims, as the indeterminacy of meaning or translation, since there are no facts about syntax that we should find problematic. For Quine:

The business of syntax is the demarcation of strings of phonemes proper to the language. More than one battery of grammatical constructions and vocabulary will probably be capable of generating the same total output of strings, but in this freedom there is no indeterminacy analogous to that of translation. Indeterminacy of translation consists rather in conflict in the outputs themselves. (Quine 1990, p. 49)

The difference between grammars is only verbal, according to Quine, or at most 'a choice of one syntactic structure rather than another for generating one and the same total output of [linguistic] strings.' (Quine 1990, p. 50) This contrasts, thinks Quine, with the case of meaning where two translation manuals for a given language can yield different results—different outputs—for given sentences of the object language, even though the two manuals are empirically equivalent, each being compatible with the same total set of behavioural observations.

But Quine's understanding of the purpose of grammar allows him to relegate differences between grammars to matters of the choice of delivery of the testable output, the empirical content, of the theories. The output of different grammars can

be the same if they deliver (provide a means of recursively enumerating) the same set of well-formed strings. In this sense two incompatible grammars will be extensionally equivalent—generating the same set of strings—and empirically equivalent—accounting for all the same empirical evidence.

The trouble with this argument is that Quine's criterion of correctness is merely *observational adequacy*, i.e. that a grammar for a language L is correct if and only if it generates all and only the strings acceptable to speakers of L. (Such grammars only *weakly generate* strings, as Chomsky puts it.) Thus if a grammar  $G_1$  and a grammar  $G_2$  generates the same total output strings—the strings of the language acceptable to its speakers—there is no indeterminacy because the two theories (grammars) do not differ in their syntactic subject matter: the proper strings of the language.

This is not the notion of *descriptive* or *explanatory* adequacy that concerns the linguist. The output of grammars is not *strings* but *syntactically structured sentences*. Strings do not cut finely enough to say which, or how many, sentences belong to a speaker's language, since the same string can count as ambiguous and be assigned two sentences structures within one language, but not count as ambiguous and be assigned only one sentence structure in another. The two languages do not contain the same sentences. Grammars postulated by generative grammarians capturing generalizations about all human languages *strongly generate* strings by assigning them a *structural description*. None of these considerations are addressed by Quine, who simply doesn't engage with the real subject matter of linguistics: the properties that make language *language*.<sup>42</sup>

A second source of pressure against mentalism is the missing object of linguistic knowledge—after all how can the facts of language we study be identified with a speaker's knowledge of language? In what sense do speakers have knowledge of something if the linguistic facts they know are determined by the very states that are meant to provide knowledge of them? This is what we called the *Missing Object of Knowledge Problem*.

We do not find out about the speaker's knowledge of language by going via his language. The linguist studies the speaker's knowledge-of-language (a state of the speaker/hearer) and this tells her what a speaker is linguistically able to do: which expressions will count as well-formed, and which will not. This study yields the linguistic properties of all the expressions generated by the speaker's I-language: what we might be tempted to call the speaker's language or idiolect—even though it is not external to the states that constitutes the speaker's knowledge of language, and partly belongs to the language faculty.

If languages are no longer seen as 'out there' but are conceived as internal to speakers, what *is* the object of speakers' knowledge of language? It is not their I-language. The I-language a speaker using is discoverable by a theorist. So which object is *a person* to when he or she knows a language? There is less difficulty in conceiving a relation between speakers and their languages on this picture since in some sense people embody their languages: languages do not exist independently of them, and so

<sup>42</sup> For a thorough treatment of these issues see Neale, (1987). [left justify]

there is no need for an elaborate account of the actual language relation. But we seem to have swapped one problem for another since now there appears to be nothing to relate them to: no *object* of knowledge. There is, however, still an important question about what notion of knowledge relates speakers to languages, and what we meant here by the language they know. Chomsky has always insisted on departing from the ordinary notion of knowledge and has talked about a speaker *cognizing* his grammar or I-language.

### 37.6.1 What Notion of *Language* Should we Adopt?

It is one thing to give an account of Chomsky's notion of language and knowledge-of-language, it is another to decide whether it is philosophically adequate in addressing the issues with which we started. So now we must ask whether we should endorse the linguist's conception of what makes language *language*.

On the linguist's conception, where languages are not external objects, we do not represent them. We need initial experience (exposure to the data<sup>43</sup>) to trigger the settings of a narrow range of parameters, but we do not learn languages from our surroundings. Thus the complex structures we compute from selected lexical items, and that count as sentence structures, are not to be found in the data we are exposed to: they are *internal* to the speaker/hearer. That was the upshot of poverty of stimulus arguments.

Language in the sense of I-language are steady states of individuals' language faculties; faculties whose initial states are genetically determined in accordance with the principles of UG. I-languages are functions in intension, and internal states of the mind/brains of individuals. Our knowledge-of-language is just such a state. I-languages generate infinite pieces of knowledge but I-languages are not to be identified with what is infinitely generated: they are states of the speaker/hearer, and thus finite. So knowledge-of-language is a property of speakers rather than a relation to some independently existing object: it is not really knowledge *of* something independent of what we are calling the knowledge state itself. It *is* the state that (in part) enables speakers to produce and comprehend indefinitely many expressions. But now we need to ask: do we really have knowledge of language at all on the linguist's conception of language?

## 37.7 THE MISSING OBJECT OF KNOWLEDGE PROBLEM

It is one thing to say that speakers do not stand in a relation of representation or knowledge to their internal grammars, it is another to say that they do not

<sup>43</sup> Although what these data are is somewhat controversial and the subject of continuing empirical research.

have *knowledge of* a language. And yet without a notion of language as something independently known it is merely a *façon de parler* to talk of a person possessing *knowledge of* language. If knowledge-of-language is a state of the person and the language a person speaks is determined by that state, it is hardly knowledge that is at issue. For how can a genuine state of knowledge constitute the thing known?<sup>44</sup> Knowledge requires there to be a subject matter to be right or wrong about. And the trouble with the linguist's notion of knowledge-of-language is that it fixes the facts of the language rather than conforming to them. The problem in a nutshell is this: if the psychological states that constitute one's competence determine *the facts about one's language*, how can those psychological states be at the same time answerable to those facts in the way knowledge requires? They can't, of course, and this is the Missing Object of Knowledge Problem. But if there is nothing independent of the internal representations to be represented, there is nothing objective to know. So how can the linguist be getting at objective linguistic properties of expressions via the study of speaker's knowledge?

We seem forced to reject the idea of an object of knowledge on the Chomskian conception of knowledge-of-language, and this is exactly what John Collins confirms:

There is, as it were, nothing to *get right*. Languages are not external objects we can go right or wrong about . . . It turns out, as part a matter of discovery, part methodology, that we do not *know* languages (better: I-languages). (Collins, 2005, p. 514–16)<sup>45</sup>

But there is still room for genuine items of *linguistic* knowledge (and Collins acknowledges as much), like the *knowledge that* certain nominal expressions interpretatively depend on others in examples like (1)–(27).

Speakers know that the pronouns in (14) and (26) cannot be construed as referring to the same person the proper names refer to respectively in those sentences. In (15) the pronoun may bear the same interpretation as the proper name 'John' and so can the possessive pronoun in (16). The reflective pronoun (or anaphor) in (15) and (19) must be construed as indicating the same person as the preceding proper name, whereas it cannot in (17). A person's I-language will generate all these items of knowledge (in fact, infinitely many of them). And all I-languages will generate equivalent structures, since these exemplify universals of humanly acquired natural languages. But the known facts are determined entirely by the internal state. They are not facts 'about the language' so much as a reflection of it. The items themselves surely count as knowledge, but the knowledge-of-language that gives rise to them is not *knowledge-of* anything. How then can what we have just said about (14) to (19) count as knowledge? What is it knowledge *of*? If we cannot think of the products generated by a person's internal I-language as the person's language, how can we

<sup>44</sup> Even psychological self-knowledge requires there to be a something known. In reflective or cogito-like cases the knowing and the known coincide but there is still a psychological state to have knowledge-of.

<sup>45</sup> Collins 2005 presents an exceptionally clear account of the linguist's conception of language along the lines discussed here.

have *anything* that counts as knowledge of language, even in our intuitive judgements? And when we take our intuitive judgements in (1)–(27) to count as items of *knowledge that*, with the subject matter being the properties we have just described, how can the intuitive judgements that present such properties of sentences also be about them? If the linguistic judgements we form constitute their own subject matter, this raises the question of whether such judgements can concern objective facts, whether such judgements have objective correctness conditions. The problem is that the linguistic knowledge claims about (1)–(27) appear to assure their own success and depend on nothing but thinkers' being in the states of having made those claims. And unlike cases of psychological self-knowledge these knowledge claims do not even appear to be about the states of knowledge themselves.

What are we to make, then, of the ordinary point that as speakers we have indefinitely many pieces of individual knowledge about the words and sentences we use? Our production and comprehension of speech results in knowledge that someone is saying such and such by hearing that someone is saying such and such. These pieces of knowledge, the output of our linguistic system, operating in concert with many other cognitive systems, need to be respected. Yet the contemporary Chomskian linguist has no satisfactory way to do so.

### 37.8 A SOLUTION TO THE MISSING OBJECT OF KNOWLEDGE PROBLEM

How did the problem arise? For the cognitivist, a grammar provides a model of a speaker's knowledge-of-language. At the same time it describes facts about the structure of the speaker's language. It can do both simultaneously because a speaker's language is individuated by the knowledge-of-language that determines it. A speaker's knowledge-of-language fixes the language she speaks (and hence the properties of her language). So a linguistic theory that characterizes a person's knowledge-of-language can specify the properties of a person's language: the structured expressions produced and recognized. However, because the psychological states that constitute one's linguistic competence determine *the facts about one's language*, such states cannot at the same time be answerable to those facts.

So how can the person count as knowing the meanings and forms of the linguistic expressions generated?

The problem lies with the particular linguistic knowledge ascribed to speakers. Speakers typically do know, without evidence or inference, what they mean by their words and which configurations of their words are grammatical. However, their intuitive linguistic judgements are not *ipso facto* correct. Speakers are authoritative, not infallible, in their native speaker intuitions. This authoritative knowledge is special

because, although it purports to be about a range of objective facts —the linguistic facts about one’s language—, it is based on nothing more than what, *prima facie*, one takes one’s words to mean, and what strings one takes to be grammatical. The problem, then, is to show how there can be a range of objective facts about language while accommodating authoritative knowledge of them. The objectivity of linguistics requires there to be objective facts to which a speaker’s linguistic intuitions are answerable—there should be a gap between the linguistic facts studied and our opinions about them. On the other hand, first-person authority requires the linguistic facts to be, pretty much, as we take them to be,—for our linguistic intuitions to be largely correct. This tension between the objectivity and first-person authority of our knowledge of language is the real problem facing a Chomskian account. (For more see Smith 1998 and 2001)

The right way to secure the objectivity of linguistic knowledge is by finding a way to make room for a distinction between *how things seem*, linguistically, and *how they are*, even when the two coincide. We need not say that the linguistic facts are independent of the speaker’s view of them. The Chomskian view that facts about languages are settled by facts about speakers’ minds requires that linguistic phenomena are *dependent* on speakers’ mental states. But the speaker’s intuitions or judgements count as knowledge only if we say these intuitions are usually the reliable upshot of speakers’ underlying states of competence. Thus the tension can be eased by paying more attention to the different levels in a speaker’s linguistic knowledge. A speaker knows a vast amount about his language as a matter of ordinary conscious reflection, as examples like (1) to (27) demonstrates.

His occurrent and conscious knowledge of these, and a welter of other, facts depends on the systematic body of unconscious, standing knowledge that he carries from one occasion to another and which provides the means of generating structures from the lexicon. Each conscious item of linguistic knowledge about a particular string or expression is derived from more general knowledge governing the structure and content of the I-language. It is this underlying linguist’s sense of knowledge-of-language, some of which is innate and shared by all language users, that fixes many of the facts an individual knows. The speaker may have no idea of the precise extent and nature of his standing knowledge; it is this knowledge the linguist is trying to model by means of a grammar. By contrast, when we say a speaker is authoritative about expressions in his language we are talking at the first-person level about his knowledge of particular linguistic facts such as which nominal expressions must be interpreted as referring to the same thing—facts accounted for by a descriptively adequate grammar. In cases of spontaneous linguistic intuitions, how the speaker takes things to be is usually how they are. But not always. The speaker can be out of step with his own linguistic system. However, the facts about the structure of his language will always be settled by reference to the underlying facts dictated by the I-language. And it is only when the effortless and groundless intuitive judgements speakers make flow from (in whatever way they do), and conform, to assignment of

structure made by the underlying I-language that the speakers' intuitions count as knowledge. There will be cases where a string will appear, *prima facie*, meaningful to our immediate intuitions but on reflection will be uninterpretable, such as (25). And equally some strings will appear ungrammatical even though one's I-language can generate permissible structures, as reflection and some coaching with centre-embedding sentences like 'The girl the cat the dog bit scratched cried.', reveals.

What this shows is that our native speaker intuitions do *not* ensure their own success but when produced *in the right way*, so as to be in conformity with the underlying facts dictated by the I-language, they count as knowledge. In just those circumstances, how you take things to be linguistically *is* how they are. This preserves both authority and the claim to knowledge. But what is it knowledge of? What are we getting objectively right under these conditions? The answer will be something about the structure of our generative procedure or I-language. So although it may appear to us phenomenologically as though we were reacting to facts about the language, out there and external to us, in fact we are reacting to something within our own breasts and not consciously accessible to us at all. It is still psychological facts about the speaker that determine the linguistic properties of expressions but there is room for objective facts at the personal level, about which, our native speaker intuitions can be right or wrong. The full range of the facts we can know by this means makes up a substantial body of linguistic information about us as speakers—information the linguist has to account for—and there is some reason to call this linguistic knowledge, in the genuine philosophical sense of knowledge. It is both theoretically characterizable and first-personally available.

### 37.9 KNOWLEDGE OF WORD MEANING

Unlike our knowledge of grammar, there is no reason to think that our knowledge of word meaning is inaccessible or sub-personal. In fact, there are good empirical reasons to think it cannot be found in the language faculty, narrowly construed (see Hauser, Chomsky, and Fitch 2002).

We first learn words at the age of around 12 months, before we are able to use syntactic structure. We acquire words and their uses one by one, only later combining them in pairs. We do not move from two words, to three, to four, etc. But as mentioned above, we simply start using whole sentences in the 'syntax spurt' at around 20–24 months. This would appear to be the combination of a very different system with our facility to use sounds in a representational way prior to the onset of language. It is possible that word meanings have, prior to the syntax spurt, nothing to do with language, at least as the linguist understands it. What is needed, then, is an account of how we learn words, what meanings we give them, and how we use them to understand others.

Our knowledge of word meaning is conscious and first-personal. There is such an experience as the meaning of a word being all there at once, or of bringing the

meaning of a word to mind as when one decides whether the use of a particular word is more apt than another. These experiences of meaning belong at the personal level. How do we acquire them and how can we use them to understand others? The quick answer, that can only be sketched here, is that we learn to have experiences with words in the context of learning words from others. The early conditions for word learning typically happen best under conditions of *joint attention* where the child and the parent are jointly attending to a commonly perceived object. The sharing of their experience of that object can be commemorated by introducing a sound label that saturates the experience. The resulting state will bring to mind, on each subsequent presentation, a recreation of the experience of a commonly perceived thing. These cases of early acquisition show us that when word meanings are introduced the experience of two subjects is co-ordinated and the involvement of an object and another person are not negligible. The traditional story may think of the child acquiring the meaning of the word from someone but more plausibly, the child is *endowing* or *investing* the sound with meaning.<sup>46</sup> These experiences set in place a way of focusing attention on a thing at a very early stage. The experiences of meaning will then attach to that word and subsequent uses of it and there will be no need for the child to entertain the idea that the word means anything different to anyone else. Were they to reflect, and there is no reason to expect them to do so, they would think: that's just what the word means. By hearing uses of the word accompanied by the understanding of the words they have, they come to attach semantic significance to what they hear others as saying. At first the meanings they hear other people's words as having are just the meanings the words have for them. This is the default case where we have no reason to suppose others are using words in a different way. This reliance by children on what *they* mean by the word will serve them well. In the default case, using words in the company of those from whom one learned them, or in sufficiently overlapping groups with whom one shares vocabulary communication will go well. It is a shock at first to a child in a foreign linguistic environment to discover that not everyone uses these sounds with the meanings they have for him. Eventually, there will be adjustments at a later and more sophisticated stage of learning. Also, it will be by dint of the grammatical and logical relations between words that the syntax spurt makes available that we come to acquire meanings of functional category words like 'of' and words less directly connected to the immediate environment.

The combination of these two systems—for word meaning and for syntax—brings about a dimension shift in the expressive power of the language user. Combining such knowledge is necessary for full language acquisition. But what we see is that the experience of meaning and the experience of hearing strings as structured respond to different parts of cognition and despite the experience of hearing what you say as there in the words uttered, the sources and objects of these two kinds of knowledge are quite different.

<sup>46</sup> See Davidson 2001 p.14. [left justify]

### 37.10 RELATIONS BETWEEN A THEORY OF LANGUAGE AND A SPEAKER'S KNOWLEDGE OF LANGUAGE

Competent speakers of a language L know which strings of words are grammatical sentences and know what those sentences mean. A correct theory for a language L specifies the grammatical structure and meaning of each well-formed sentence of that language. What relation, if any, is there between the theory and the speaker's knowledge of a language? They both concern the same object: the language L used and understood by the speaker. The theory states what the speaker knows but not necessarily in the *form* in which the speaker knows it. Hence the speaker does not have knowledge (even unconscious or tacit) of the theory. Nevertheless, the correct theory of his language will be a theory of what the speaker knows, or of what determines all the facts he knows.

The theory will invoke certain linguistically relevant properties to describe the linguistic form and meaning of the expressions speakers use and recognize. What relation is there between speakers and these theoretically described properties? Is there some epistemic relation between the two that must be respected by any satisfactory theory of their language? Or can speakers be blind to the properties recorded by the theory? We have seen that the linguistic properties of structure uncovered by the theory of syntax are properties speakers must be sensitive to, or respect, in combining words into sentences. The structural properties of the strings they produce and respond to will be assigned by their underlying linguistic system and will have impact on their conscious experience of speech. In the case of word meaning, the best theory may have to make prior use of the meanings one hears words as having. There may be no other terms, and certainly no properties to be found in the language faculty, in terms of which to capture the meanings of words in the speaker's language. The meanings the speaker grasps are consciously experienced as part of the phenomenology of speaking and understanding. How close the theory gets to capturing the meanings of his terms depends on how close the theorist is in his understanding of those words to the speaker.

To end, there are two faulty assumptions in the study of language we are now in a position to give up. The first is Quine's behaviourist assumption, shared by many, that the task facing the child learning the language and the task facing the theorist may be the same. Quine is wrong on this point. As Chomsky's arguments show:

- (i) the learner cannot learn a language L unless she knows P antecedently, (where P is some set of domain-specific constraints on the structure of possible human languages)
- (ii) P is innately known (because it could not be learned on the basis of impoverished linguistic data available to the learner as the poverty of stimulus arguments show)

- (iii) Universal Grammar aims to specify the domain-specific constraints P that speakers innately know and respect.

The linguist's task is to figure out the value for P. The learner's task is to construct a grammar (acquire any of the class of humanly possible natural languages) on the basis of knowing P and being exposed to the primary linguistic data. The tasks facing each are quite different.

Secondly, Dummett rejects the model of speaker and hearer having a similar piece of internal apparatus because it would make our understanding of others a matter of hypothesis about what is going on in the other, which is both risky and phenomenologically distorting. The phenomenology of understanding speech can be immediate precisely because no such hypothesis is entertained. What happens is that my internal equipment—my language faculty—automatically assigns a structure to what is perceived thus giving rise to my hearing a sentence as structured. That together with the default case in which the meanings I hear someone else's words as having are the meanings those words have for me, ensures that I arrive at an understanding of what I hear as a matter of the fast and mandatory workings of my linguistic capacities and conscious apprehension of meaning. The job of those capacities is to interpret the sound stream and present a product to consciousness: it is not to hypothesize about what is going on in some one else's mind. And when you use words with meanings similar to mine, words arranged according to the constraints grammar places on combining words of those categories in that order, the meaningful sentence I end up hearing is probably close enough to what you experienced in producing it for us to count as communicating. Beyond that, reflective interpretation may be needed to understand one another, but in the normal case nothing said here about knowledge of language shows understanding to be anything other than immediate, inner and relatively secure. But it is knowledge of different kinds with different sources that give shape and character to our use and understanding of language.

## REFERENCES

- Chomsky, N. (1955). *Syntactic Structures* (Mouton).  
 — (1980). 'Review of Skinner's *Verbal Behaviour*, 1959, reprinted in *Philosophy of Psychology*, Vol.1 edited by N. Block (Methuen).  
 — (1965). *Aspects of The Theory of Syntax* (MIT).  
 — (1972). *Language and Mind* (Harcourt, Brace, Johanovich).  
 — (1980). *Rules and Representations* (Blackwell).  
 — (1986). *Knowledge of Language* (Praeger).  
 — (1987). 'Reply to Alex George', *Mind and Language*.  
 — (1993). *Language and Thought* (Moyer Bell).  
 — (1995). *The Minimalist Program* (MIT).  
 — (2000). *New Horizons in the Study of Language and Mind* (Cambridge University Press).  
 Collins, J. (2004). 'Faculty Disputes' in *Mind and Language*.  
 Crain, S. (1991). 'Language Acquisition in the Absence of Experience' in *Behavioural and Brain Sciences*.

- Davidson, D. (2001). 'Externalisms' in *Interpreting Davidson*, eds. P. Kotatko, P. Pagin, G. Segal (Stanford CSLI).
- Davies, M. (1981). *Meaning, Quantification, Necessity* (Routledge).
- Dummett, M. (1993). 'What Do I Know When I Know a Language' 1978 reprinted in his *Seas of Language* (Oxford University Press).
- Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., Plunkett, K. (1996). *Rethinking Innateness: A Connectionist Perspective on Development*, (MIT Press).
- Fodor, J. (1990). *In Critical Condition* (MIT).
- George, A. (1989). 'Whose Language is it Anyway?' in *Philosophical Quarterly*.
- Harris, J. and Lindsay, G. (2003). 'Vowel Patterns in Mind and Sound' in *Phonological Knowledge*, edited by Burton-Roberts, Carr & Docherty (Oxford University Press).
- Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). 'The Faculty of Language: What is it, Who has it, and How did it Evolve?', *Science*, 298, 1569–79.
- Higginbotham, J. (1983). 'The Psychological Reality of Grammar', in *How Many Questions*.
- (1991). 'Remarks on the Metaphysics of Linguistics', in *Linguistics and Philosophy*.
- Katz, J. (1990). *The Metaphysics of Meaning*, (Cambridge University Press).
- Lewis, D. (1983). 'Language and Languages' in his *Collected Papers*, Vol. 1 (Oxford University Press).
- (1995). 'Meaning without Use: a reply to Hawthorne' *Australasian Journal of Philosophy*.
- McDowell, J. (1998a). 'Anti-realist Semantic and the Epistemology of Understanding' in his *Meaning, Knowledge and Reality*, (Harvard).
- (1998b). 'In Defence of Modesty' in his *Meaning, Knowledge and Reality*, (Harvard).
- (1998c). 'On the Sense and Reference of a Proper Name' in his *Meaning, Knowledge and Reality* (Harvard).
- (1998d). 'Physicalism and Primitive Denotation: Field on Tarski' in his *Meaning, Knowledge and Reality* (Harvard).
- Neale, S. (1987). 'Meaning, Grammar and Indeterminacy' *Dialectica*.
- Quine, W. (1990). *The Pursuit of Truth* (Harvard).
- Schiffer, S. (1994). 'Actual Language Relations' in *Philosophical Perspectives*, 9.
- Schiffer, S. (2005). 'Two Perspectives on Knowledge of Language', Manuscript for GLOW.
- Smith, B. C. (1992). 'Understanding Language', in *Proceedings of the Aristotelian Society*.
- (1998). 'On Knowing One's Own Language', in *Knowing Our Own Minds*, eds. C. Wright, B. C. Smith and C. Macdonald (Oxford University Press).
- (2001). 'Understanding Idiolects: a Reply to Barber in *Mind and Language*.
- Wiggins, D. (1991). 'Language as a Social Object' in *Philosophy*.
- Wittgenstein, L. (1969). *On Certainty* (Blackwell).

**Queries in Chapter 37**

- Q1. Please note that, the author had inserted the footnote number as 35 for this footnote. We have renumbered this footnote sequentially to maintain consistency. Kindly check and advise.

