

What Remains of Our Knowledge of Language? Reply to Collins

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

*The new Chomskian orthodoxy denies that our linguistic competence gives us knowledge *of* a language, and that the representations in the language faculty are representations *of* anything. In reply, I have argued that through their intuitions speaker/hearers, (but not their language faculties) have knowledge of language, though not of any externally existing language. In order to count as knowledge, these intuitions must track linguistic facts represented in the language faculty. I defend this idea against the objections Collins has raised to such an account.*

Key words: knowledge of language, representations of language, language faculty, intuitions, linguistic self-knowledge

What equips us to use and understand a language? The full story of how we experience sounds linguistically, how we produce and hear sentence forms, is both complex and at odds with the everyday impression speakers and hearers have that they are participating in an external, common and conventionally sustained linguistic practice. As speakers and listeners we enjoy conscious experiences of speech that unify quite diverse sources of information giving us the illusion that we are dealing with a single thing: *The* language. The assumption that there is such a single thing as the language we share has no empirical credibility but remains a firm fixture of our naïve world-view, and is popular with many philosophers. Theoretical progress in generative linguistics, however, has significantly undermined the naïve view, and made us focus instead on non-conventional and largely innate principles of grammatical structuring in the human mind; principles that shape the acquisition of individual grammars. The linguistic elements we make conscious use of when speaking, or of which we are apprised when listening, are not to be found in a shared, public language. There are several sources of an individual's knowledge of language: some of it first-personal knowledge of word meaning, some of it third-personal knowledge of what other speakers mean by their words, and much of it due to the sub-personal

influence the language faculty has on what we find an acceptable syntactic arrangement of words. With our knowledge of language depending on first- third- and sub-personal aspects of cognition, there is no reason to suppose there is single locus of linguistic significance. Though we need not deny this is how things *appear* to us. Michael Dummett is right to describe speech as a conscious, rational activity: this is certainly how we experience it. But there is no reason to think, as Dummett does, that all the regularities of our speech are rationally chosen. The grammars that constrain the way our words come out in the right order when we speak are not consciously adopted or rationally adhered to: they are not the result of conventions, but of specialized, cognitive mechanisms to which we have no conscious access. To discover *their* properties we need to turn to generative linguistics: a branch of cognitive psychology.

This shift in focus from studying supposedly shared, public languages—usually by attending to facts about speakers' behaviour¹—to the study of *linguistic forms* in the minds of speakers—the forms their internal apparatus assigns to the sounds, signs or marks they produce and encounter—is due of course to the work of Noam Chomsky (1959, 1965, 1968, 1980, 1986, 2000) who has always stressed that it is *knowledge of language*, or the speaker/hearer's *linguistic competence* that linguists are studying when they study language.

There is a considerable irony here since Dummett has long maintained that the notion of knowledge plays a crucial role the study of language, and that an account of the significance of a speaker's language must go via an account of his knowledge of language.² Dummett is certainly right on this point. Without adverting to speakers' knowledge we could not be sure we were actually characterising *their* language: the languages speakers use and understand. Where Dummett goes wrong is in characterising knowledge of language as a complex practical ability, found at a single level of organisation in the speaker, and emerging from a single source. Dummett sees linguistic knowledge as a practical ability³ because language is conceived as a communal practice that has to be mastered by individuals in the community: a conception Dummett retains by not having seen all the way through to the consequences of his insight that we should be studying the individual speaker's knowledge of language. His residual, almost ancestral, commitment to an external entity—*The* language—prevents him from recognizing the di-

¹ Michael Devitt is no different to Quine in this respect. Even though the former embraces representations that the later would abjure, it is done on the basis of a theory about behaviour not a theory of the cognitive states of speaker/hearers underlying their behaviour.

² According to Dummett, a theory of meaning for a language should be a theory of understanding: a theory of what speakers know when they know the language; otherwise the theory we come up with may not capture the properties of *their* language, the language they use and understand. See Dummett 1976 and 1978.

³ Though Dummett does see this practical ability as having an irreducible theoretical component. None of this, however, explains why speech should be a conscious activity, since our practical abilities need not be part of consciousness.

verse sources and different cognitive levels on which a speaker's knowledge of language depends. As a result, the challenge to characterize the individual speaker's knowledge of language passes to Chomsky and the generative linguistics programme.

But here we discover a further irony. For now Chomsky, has recently taken to denying that our knowledge of language gives us knowledge of anything. (See Chomsky 2000) Whereas in *Aspects of the Theory Of Syntax*, Chomsky claimed that knowledge of language—the knowledge that fixes the properties of a speaker's language—was unlike ordinary knowledge in being in principle inaccessible to consciousness. Other significant differences between the notion of knowledge Chomsky was trying to capture and the notions of knowledge-that and knowledge-how that philosophers typically focus on led Chomsky to propose the neologism 'cognize' in place of 'know' when talking about a speaker's relation to his or her grammar. But these days little remains of the idea that a speaker stands in a cognitive relation to the grammar of his or her language, or even that he or she has *knowledge* of language.

The move is in many ways a rational one. For Chomsky, 'language has no existence apart from its mental representation' in the mind of the speaker⁴, so there is no object of knowledge independent of the states of the mind/brain for the speaker/hearer's linguistic competence to be right or wrong about. The facts about a language are fixed by the psychological states of the speaker that constitute his or her competence. So in what sense can we talk about those states as providing the speaker/hearer with *knowledge* of the linguistic facts, when those states are not answerable to the facts—as knowledge would require—but rather determine them? Why, therefore, not give up all claim to be talking about knowledge and simply concentrate on the states of the internal mechanism that determine the linguistic properties? What remains is a piece of rather misleading piece of terminology, 'knowledge of language', which we should now see as a technical term, if necessary hyphenated as 'knowledge-of-language', stripped of both its everyday and philosophical connotations.

Such is the new Chomskian orthodoxy, and perhaps no philosopher has done more to bring it to other philosophers' attention and explain it cogently than John Collins. Not only has he taken great care to uncover and render perspicuous the real motivations of Chomsky's thought but he has added a much needed philosophical framework to current linguistic theorizing, enabling him to produce a battery of impressive arguments in favour of the new picture. I greatly admire the insights he has brought to this area, and the way he has engaged philosophers who have not yet seen the light, but like all ideologies there will be persisting struggles about the right course for the struggle to take.⁵ So while Collins describes Robert Matthews and myself as being on the side of the

⁴ Chomsky 1972, 169 fn.

⁵ See Collins 2004, 2006, 2007.

angels he sees us as departing from the true faith and seeks to correct our wayward thinking. Like Matthews, I remain unconvinced of any need for correction. But I want to set out clearly the points of disagreement between Collins and me over the new orthodoxy and make some of the things I stated previously a little clearer. I realise that for some this will amount to ‘in-fighting on the left’, but I think real progress can be made by setting out a range of possible positions on knowledge of language, and I suspect, as Collins does too, that when the dust settles we will not be so far apart in our thinking.

The position I will outline aims to respect Dummett’s view that speech is a conscious activity, admitting personal level knowledge of language. However, it also respects Chomsky’s cognitivist conception of language as internal to the mind of the speaker. Reconciling these two positions is possible by departing from Dummett, in admitting different levels of cognitive organisation within the speaker and jettisoning the idea of an external, public language. It also requires one to acknowledge, as Chomsky is unwilling to, that we do have knowledge of linguistic facts, and that (most of) our linguistic intuitions are answerable to constraints imposed by the language faculty. The difficulty is seeing how there can be a range of objective facts for a speaker to know when these facts are determined by the internal states of speaker.

1. *Speaker/hearer’s intuitions as knowledge*

Speaker/hearers mostly know what they are saying, and what they take others to be saying. And when explicitly asked to consider presented strings of words, the immediate intuitive responses they give serve as good guides to certain linguistic facts. It is these immediate, intuitive, conscious judgements that the linguist takes to be a central source of linguistic data. And I agree with Matthews, 2006a, when he says:

[L]inguists need some justification for treating these putative sources of data as

alia, the upshot of underlying states of the subject's sub-personal linguistic system. Thus for the linguist they usefully reflect internal facts about the organisation of the language faculty. Sometimes, however, the immediately intuitive awareness a speaker has of how a string sounds may be out of step with what the language faculty requires because the response is subject to the limitations of the parser whose workings are affected by real-time constraints and memory limitations that prevent a speaker/hearer's apprehension of the permissibility of the string licensed by the real structure assigned to the string by the language faculty. By separating the conscious personal level states of the intuiting subject from the underlying states of the faculty, we can credit the intuitions of the *speaker* as delivering linguistic knowledge just in case his intuitive judgements of which strings are *acceptable* conform to the requirements of the *language faculty* about which arrangements of items are *grammatical*. In this way, the speaker's intuitions are answerable for their correctness to conditions imposed by the underlying language faculty. This picture enables a speaker's linguistic intuitions to count as knowledge even though there are no external linguistic objects for the speaker for her to have knowledge of. Thus it solves what I have called 'the missing object of knowledge' problem (Smith 2006a, 2006b).

Collins doubts whether this can be right because, according to him, it would appear to make the speaker have knowledge about the internal workings of his or her language faculty. But on the account I give, although the speaker's intuitions must be in some sense *aligned* with the facts settled by the underlying linguistic system, and *answerable* to them in order to count as knowledge, such intuitive judgements do not have to be *about* states of the faculty. A speaker's intuitions are *about* whether a presented arrangement of word-like items is acceptable and these are taken to be judgements of whether certain arrangements are permissible. Whether or not they are syntactically permissible is a matter determined by the speaker's internal grammar. I shall say more about these issues below.

2. *Collins' Complaint*

Collins makes a fair objection to Matthews' knowledge requirement, by pointing out that linguistic intuitions are data, and as such do not need to be justified, but merely correctly gathered, respected and explained by the theory. That's right. Intuitions are data and they can guide us in constructing linguistic theories. But there is something special about the data of linguistic theory and of course of the subject matter of linguistic theory: they belong to a speaker's psychology. What kind of psychological states are these? What contents do they have? Are they about or correlated with something external to the speaker's mind? All these questions are up for grabs in a way they are not in the data of physics,

or cognitive ethology, say.⁶

What are the linguistic facts the theorist targets and why are a speaker's intuitions so relevant? I take the story to go as follows. The linguist wants to explain certain facts about the grammatical organization of the language—for the Chomskian, facts about the speaker's cognition. But how does the linguist get at the facts to be explained? The answer must be largely via a speaker's intuitions. Of course, it is not linguistic intuitions *per se* that the linguist is trying to explain: linguistics is not *about* linguistic intuitions, it is not a theory of *them*. It is a theory of the linguistic structure of human languages. However, in order to get at the linguistic facts, which the appeal to underlying structure is meant to explain, the theorist will rely on the data of speakers' linguistic intuitions.

The reason linguists consider a speaker's intuitions as providing better data than the odd snippet of conversation, or hastily typed email, is that linguistic intuitions are thought to more accurately reflect what the speaker considers as acceptable strings than the half-finished phrases, false starts, slips of the tongue, or shortcuts, that crop up in a corpus of use. It is facts about the speaker's language *as she takes them to be*, rather than her, at times, aberrant use of language, that linguists seek to explain by positing level of linguistic structure constrained by an internal grammar or I-language. Thus speakers' intuitions are the best guide to the linguistic facts, containing less noise and more targeted information than corpus data. But they can be so only because they mostly amount to knowledge of linguistic facts; because they are mostly accurate guides to the permissible strings and construals of strings of a speaker's language. In this way they reflect the workings of competence.

Collins says: "It escapes me why we should also be interested in *justifying* a speaker/hearer's judgements to be a reflection of competence" ('Redux', 17). I agree that *the linguist* need not justify that such data usually reveal the linguistic facts to be explained. However, linguistic intuitions usually do amount to knowledge and are a record of the linguistic facts. Knowledge of this kind is spontaneously available to speaker/hearers as immediate conscious judgements, and the linguist is simply entitled to take these judgements to be mostly good guides to the facts. The task for the philosopher is to explain why the linguist is so entitled; i.e., why the immediate conscious responses of the speaker to presented samples typically counts as knowledge of facts about the permissible structures of speaker's language.

The story I just told about reliably produced linguistic phenomenology involves no supposition that the speaker has to *justify* the respons-

⁶ Closer analogies are those Collins mentions concerning the Sally-Ann tests for children's acquisition of a theory of mind, or the verbal reports of visual experience used as data for theories of visual perception. Children's judgements are not about the TOM module, they are however data for psychological theories about children's mentalizing capacities. We shall discuss these later.

es she gives, nor that the theorist can *justify* the assumption that the speaker's intuitions reveal a form of knowledge. The underlying cognitive states provide neither reasons nor causes alone for the verdicts a speaker effortlessly arrives at. Nevertheless, it is because the underlying grammatical facts are as they are that, in most cases, the forms the speaker finds acceptable are as they are. Linguistic intuitions reflect and show the influence of linguistic competence; though they will show the influence of much more besides. Collins is right when he says, "Our theory of the faculty is not simply a descriptive encapsulation of justified judgements". (Collins, 'Redux', 17). Agreed. However, he goes on to say:

The common minimal core is that the informant offers an *interpretation* or *construal*. Whatever form this takes, it is the theorist's task to discern the contributions of the different systems that give rise to the judgement. As far as I can see, in no case does an issue of justification or warrant arise. (19)

But of course the absence of justification or warrant does not mean the absence of *knowledge*. Though how intuitions come to count as knowledge, and what they give us knowledge of Collins has still to explain.

Collins' sums up his complaint against my view as follows:

I shall take Smith's point to be simply that intentionality must come into view because the faculty must ultimately bear on the first-person states, not that the faculty is itself contentful. Now, transparently, a non-intentional conception of the language faculty does *not* involve a denial of our knowledge of language, if such knowledge amounts to intuitive judgement of the kind that serves as data for the linguist. Smith would agree with this. What is denied according the conception I am commending is that such 'knowledge' requires any kind of epistemic grounding, justification, or authority for it to serve as evidence for linguistic hypotheses, i.e., the knowledge at issue, qua serviceable as linguistic evidence, raises no epistemological concerns of objectivity or answerability at all. ('Redux', 20)

I conceive matters differently since I take states of the faculty to be contentful. Notice, Collins doesn't deny that *people* have knowledge, what he rightly denies is that the sub-personal *language faculty* has knowledge. But why suppose that to deny the faculty has knowledge one must deny its states have *content*? I shall say more in a moment about the difference between states amounting to *knowledge of* something and their being *representations of* something.

But first, the important thing to note is that speakers' intuitions are usually partial products of the underlying system. Being products of that system is what makes them of interest to the linguist. Unless they typically presented the linguistic facts to be explained they would be at best very indirect and distracting evidence for the linguist to work on. However, being partial products of the faculty (along with other factors) we cannot automatically take them to reflect competence. In those cases where the speaker's intuitions about what is acceptable are at odds with what the theorist predict as grammatical, the theorist need not revise the generalisations that work in the main and attempt to accommodate the

new data: instead, she can attempt to explain the speaker's recalcitrant judgement in terms of the bias or limitations of the parser. This will be the strategy with centre-embedding constructions and garden-path sentences because, which according to the theorist have structures analogous the structures of strings the speaker already takes to be acceptable and on which the theorist's current grammatical hypotheses depend.

Grammar and Parser

Systematic divergence between what a speaker's finds *acceptable* and what the best theory of his I-language pronounces *grammatical* will be due, in many cases, to processing limitations of the parser. Such cases provide the linguist with evidence of the correct way to draw the distinction between linguistic competence and performance. In so far as there is an issue of answerability: of speakers' intuitions to facts about their grammar; of *acceptability* to *grammaticality*, there will be reason to say the recalcitrant intuitions are misleading as to the grammaticality of strings, or are erroneous. Of course, there may be nothing wrong with the speaker arriving at these verdicts because of way the parser works, just as there is nothing wrong with the visual system's delivery of a visual impression that the two lines in the Muller-Lyer illusion are the same length. Nevertheless, in normal cases, the workings of the system have such systematic errors and biases built into them.⁷

Thus we see that the range of evidence the linguist consults can extend well beyond knowledge-involving intuitions. But were none of the speaker's intuitions to record knowledge of linguistic facts there would be no knowing for sure what the properties of the speaker's internally constituted language were, and there would be no clear way to draw the competence/performance distinction.

So what makes these intuitions knowledge and what do they give us knowledge of? Collins notes that: 'Smith, 2006b, uses such terms as "response" and "tracking", which suggest a non-normative relation between personal and underling states. "Reflex" would surely do as well.' Not quite; though it is interesting that Fodor cites, in his epigraph to *The Modularity of Mind*, Marcus's remark that parsing is like a reflex. As Fodor points out in a later article, he is interested in computations, which, like reflexes, are automatic and mandatory, but, unlike the reflexes of behaviourists, are smart processes that go beyond the information given. These smart, fast, mandatory, domain-specific computations are of course how Fodor's encapsulated input modules work. In the case of speaker's intuitions, just like reflexes, the products that appear spontaneously in consciousness can be given no justification; nor need they be; this is simply how things appear to the speaker. But unlike reflexes, how things appear to a speaker, linguistically, in her immediate conscious

⁷ By my lights, it is in cases of this sort, and not in the case of correct intuitions, where acceptability coincides with grammaticality, that we should talk of illusions, pace Rey 2006.

responses, is mostly how they are. Being without grounds or justification these intuitions would appear to count as a form of basic knowledge.

Linguistic Self-Knowledge

There are other well-known cases of immediate judgements lacking justification and still amounting to knowledge. Judgements of our current psychological states provide just such cases. Each of us is effortlessly authoritative about what we are currently thinking, without basing these claims on any kind of inference or evidence. Unlike our knowledge of other people's mental states, we have no grounds on which we base our claim to know our own minds. The demand that someone justify her claim to know how she thinks and feels is barely intelligible: she just knows. The philosophical problem of self-knowledge is to explain why groundless judgements about one's own mental states are typically more reliable and authoritative than claims for which justification *can* be given. Why are opinions about one's current psychological states, arrived without appeal to inference or evidence, largely right, when they also answerable for their correctness to objective facts about the mental states one is in; facts others can know on the basis of inference and evidence.⁸

The parallels should be clear, and indeed linguistic self-knowledge seems like a special case of psychological self-knowledge: namely, how do we reconcile the effortless authority of immediate linguistic intuitions with their answerability and objectivity?⁹ Pace Collins, were they not authoritative and reliable, linguistic intuitions would not be a good source of evidence for grammatical theory.

The threatened lack of objectivity arises because the states that constitute the speaker's knowledge of language (in Chomsky's sense) determine the facts about the speaker's language, and so cannot be answerable to them in the way knowledge requires. The story would be, in one way, metaphysically simple if there *were* a set of externally given linguistic facts (as Michael Devitt supposes) for a speaker's intuitive judgements to concern. However, there would be a corresponding epistemological cost in trying to provide a credible account of how a speaker's spontaneous verdicts succeeded in aligning him with how things stood in the external realm.¹⁰ On the other hand, if there is nothing external for the speaker's judgements to latch onto how a speaker can have knowledge of language? What happens to the objective linguistic facts to which a speaker's intuitions are meant to be good guides?

In order to restore the idea of speakers' intuitions providing the lin-

⁸ For more on the problem of self-knowledge, see Smith 1998.

⁹ In the case of linguistic self-knowledge our immediate judgements are answerable in not to facts at the personal but sub-personal level.

¹⁰ Of course, Devitt has a (curious) view about intuitions to accompany his metaphysics, namely that they are quickly formed theoretical judgements, with linguists having more reliable intuitions than their informants.

guist with a good guide to linguistic facts we need to say how a speaker's intuitions can typically amount to knowledge: i.e. how intuitions can be objectively answerable to something more than the speaker's impression of what is linguistically acceptable. Impressions, after all, are cheap.

The account I previously gave (Smith 2006a and 2000b) says that a speaker's intuitions in part depend on the underlying linguistic system: they are correct when they classify as permissible strings licensed by the internal grammar which violate no constraints and where all features of the word-like items have been checked. The speaker's intuitions are answerable for their correctness to the constraints and requirements the language faculty puts on the grammatical arrangement of items in the speaker's lexicon. And when the speaker's intuitions—how things seem to him—give a verdict compatible with the faculty's constraints and requirements, they amount to knowledge of linguistic facts: facts the theorist must try to explain.¹¹

When a speaker's intuitions are not reliable—though correctly arrived at—how things seem to him will not be how they are. Speakers' intuitions have objectivity and can count as knowledge only if there is a distinction between how things seem to the speaker and how they are. And we secure that distinction by stating conditions where which how things seem to the speaker coincides with how they are. This can be done by advertent to facts at the underlying level of linguistic structure dictated by the faculty. To be knowledge, our judgements must be suitably *in tune with* how things are at the level below—to which arrangements and grammatical dependencies among the word-like items are permissible. It is this suggestion that troubles Collins most:

The idea appears to be that if intuitions—what seems right—are to serve as evidence, then they must be distinguished from what is right: 'there should be a gap between linguistic facts and our opinions about them' (*op cit.*). The problem with this thought is that we don't take intuitions to be 'opinions' about the language faculty anymore than we take first-person reports of the visual field to be 'opinions' about the visual system or reports from Sally-Ann experiments to be 'opinions' about a ToM module. ('Redux', 20, fn. 29)

First, it is not that what seems right must be distinguished from what is right if intuitions are 'to serve as evidence': It is that there must be room for such a distinction if they are *to serve as knowledge*. Second, Collins objects that on my account the speaker's intuitive judgements would appear to be *about* the underlying states of the language faculty to which they are objectively answerable. However, no such problem threatens.

¹¹ The nature of the compatibility in question needs more careful spelling out than I can undertake here. On the one hand we need to know how much structure a speaker is consciously aware of, and on the other, we need a detailed structural description for each string. Clearly, one is aware of far less than the structure posited by syntactic theory, yet it is clear that one hears a sentence as structured—think of ambiguous strings. People also have a sense of what the constituents of a sentence are. For a recent attempt to address the question of what linguistic information is and is not consciously accessible to speakers, see Kent Johnson 2007.

We can say everything I said above without thereby claiming that the speaker's immediate opinions concerning which arrangements of words are acceptable are about states of the faculty. The opinions of a speaker are about *strings*, and *not* about the language faculty, just as a child's immediate opinions in the Sally-Ann task are *about* where they think Sally will look for the marble, *not* about the Theory Of Mind module that the Sally-Ann task is meant to probe.

In both cases, it is because the subject's internal mechanisms are as they are that subjects' opinions take the shape they do. In the linguistic case it is the configurations that the internal grammar or I-language permit that makes the particular opinions the speaker comes up with, *about which strings are acceptable*, correct or incorrect. The facts about the linguistic properties of strings are determined by what the speaker's internal grammar licenses, and judgements made in accordance with configurations of the grammar amount to knowledge. These items of knowledge provide the theorist with information (evidence) of the workings of that internal grammar.¹² Similarly, although the three-year old child's pronouncements are ostensibly about where Sally will look for her marble; they reflect how she conceives things at her current stage of cognitive development and thus provide information about the development of a possible Theory of Mind module. Although the child's judgements are an accurate reflection of how she conceives others' mental states at that age, unlike the language case, these judgements are in addition required to get matters right concerning the external situation regarding the mental states of others. Thus the judgements properly arrived at for a child of age three are not accurate judgements about others' states of minds, while the judgements of four year olds typically will be. By contrast, when a speaker's linguistic judgements are informed without interfering factors they count as correct only because they are produced by and reflect the requirements of the I-language. They are answerable to nothing else; i.e. nothing external to the speaker. They simply record facts determined by whether strings of expressions meet or fail to meet the relations of grammatical dependence and other conditions the grammar places on word-like items. Thus, a speaker's intuition that (1) is acceptable and (2) is not, reflects the workings of the speaker's internal grammar that a theory of that grammar has to explain; namely, that (1) is grammatical and (2) is not:

- (1) The man I saw shaved himself
- (2) *John's mother loved himself

The linguist must construct a theory that models the speaker's grammar or I-language, that determines which arrangements of expressions, and which construals of them are permissible. Facts about the faculty are not known to the speaker, but facts about strings determined by the faculty

¹² Other intuitions provide the linguist with information about the interactions of competence and performance, thus indirect evidence about the workings of the grammar.

are typically known. We have seen how intuitions that do not amount to knowledge can still be of use in formulating a more comprehensive theory of production and comprehension of language involving a distinction between the grammar and the parser. Intermediate cases where our intuitions do not render a clear verdict one way or another, but which give rise to claims about which strings are better or worse than others may still give the linguist important clues about what *is* grammatical and what is not. So evidence comes in many shapes and sizes. Nevertheless, if the speaker's intuitions did not record knowledge of linguistic facts the linguist would have great difficulty in targeting the facts to be accounted for by a descriptive and explanatorily adequate theory.

The philosopher's task is to say why a speaker's intuitions mostly provide him or her with knowledge of linguistic facts given that these groundless opinions are not answerable for their truth to anything external to the speaker. The speaker's competence fixes the facts of the speaker's language and thus are not answerable to the linguistic facts in the way knowledge requires. How then can the speaker have *knowledge* of linguistic facts? I answered this by saying that the speaker's intuitions, or what those intuitions put him consciously in touch with, largely concern which arrangements of words are permissible as determined in accordance the underlying states of the speaker's linguistic competence. What he has knowledge of is which strings of expressions are permissible and which interpretations of those strings are possible: facts determined in part by his internal grammar. As a matter of phenomenology he may think he is trying to adhere to facts about a public language and what it permits, though, in actual fact, he is responding to something internal but inaccessible to him: the dictates of his language faculty. The picture is similar to Hume's account of moral judgements, where subjects take themselves to be reacting to the cruelty or viciousness of an act observed, though in fact they responding to something in their own breast: the feeling that observing the act gives rise to in them. Similarly, when performance factors don't intervene, what the speaker has knowledge of are linguistic facts fixed by his internal language faculty and his lexicon, and not knowledge of anything external.

The inward direction of fit for our linguistic judgements reveals a blunt stopping place. The states of the language faculty that such judgements are answerable to are not themselves answerable to anything: they are just facts about human cognition. The language faculty doesn't get things right or wrong; there are no correctness conditions for its states to meet. It simply permits (1) and proscribes (2). As a device, mechanism, or set of templates it is just configured this way. There is nothing suspect about this answer: this is just how we are built, with all human language users bounded by the same universal grammar principles, modulo the setting of values for variable elements within a fixed parametric range. Philosophers who hanker after the notion of unconscious or tacit knowledge *in* the faculty will be disappointed. How-

ever, the lack of answerability of the states of the faculty to anything beyond it does not prevent the speaker having personal level states of knowledge: knowledge of the correct linguistic arrangement of expressions of her language. Intuitions about these arrangements count as knowledge when they track states representing word-like items with particular features as standing in certain relations of grammatical dependency to one another. Tracking occurs when acceptability judgements are aligned with what is grammatical according to the output of the faculty, and *typically* what is acceptable is what is grammatical. It is this answerability of a speaker's intuitions to something more than whatever she comes up with that gives these judgements some objectivity and since they are largely right the speaker counts as authoritative about her language. Collins says, 'the knowledge at issue raises no epistemological concerns of objectivity or answerability at all.' (ibid. 20) But this is just wrong.

What seems to be bothering Collins is that, 'the evidential status of intuitions does not rest upon their tracking the states of the faculty: parsing failures occur precisely where such 'tracking' fails.' (ibid. 21) Agreed, but the status of intuitions as items of knowledge does depend on their tracking the requirements of the language faculty. None of this makes what we intuitively judge come to be *about* the inner linguistic system.

Collins is right to point out that 'the structural object that matches our first-person object will be the effect of the interaction of narrow syntax and the interpretive components, for the output of the former will be a sequence of phases. In other words, there is no LF structure that might be taken to embody our sentence sized syntactic intuitions independent of the external systems.' This is why I speak in terms of 'tracking' or being 'determined in accordance with the requirements' rather than in terms of outright correspondence or matching.

Representations of Language versus Knowledge of Language

Collins and I agree that the language faculty doesn't embody knowledge but it can give rise to it. As Collins says: 'the faculty must ultimately bear on the first-person states'. But what is denied is 'that the faculty is itself contentful' (20). He acknowledges:

a non-intentional conception of the language faculty does *not* involve a denial of **our** knowledge of language, if such knowledge amounts to intuitive judgement of the kind that serves as data for the linguist. What is denied is that such 'knowledge' requires some grounding or justification (ibid. 20)

That's right, and we've seen how authoritative first-person knowledge lacks grounding or justification. But it does require an object and Collins has yet to tell us what that is. Pace Collins, I think that states of the language faculty represent syntactic properties and need to do so in

order to give rise to knowledge. I will argue that states of the faculty, or those intermediate working states of the parser generated in accordance with the requirements of the faculty, have to represent syntactic structure; which is not to say they have to be representations of anything external to the speaker. They just have to represent relations of grammatical dependence and other features among items tokened. We need representational contents in order to say what these representations represent, and what is more, I think Collins himself cannot avoid talking about the faculty in terms of content (though Robert Matthews does). Let us begin with the charge of unavoidability.

Collins declares that current linguistic theory is attempting to specify a function in intension, a function that ‘encodes the structure of pairs of phonological-semantic representations’(p.3); namely, those that are explanatory of speaker/hearer’s judgements. The function ‘is a way of recursively defining a set of structures whose character is explanatory over the evidence of speaker/hearers’ understanding of the linguistic material’ (ibid.). It might seem as if the function was just the theorist’s way of specifying information about linguistic structure, but, as Collins says, the speaker/hearer’s ‘competence *is* the function’ (ibid.). This something *in* the speaker/hearer plays the role of the theoretically characterised function that appears in the linguist’s grammar, although what goes on in the speaker will not be represented in anything like the format of the theory. In particular, it will not be a set of propositions or a theory the speaker tacitly or unconsciously knows.¹³ However, some encoding of the information represented by the theory has to be present in the speaker/hearer if the specified structures are to *explain* aspects of his or her psychology? Surely, such structures can explain the patterns in a speaker/hearer’s linguistic judgements only if those structures are psychologically real; they must figure somewhere in the speaker/hearer’s psychology. But how do they figure? They cannot be characterized as neural structures, nor are they the structures our cognitive representations have, whatever structure that is. And why think the mind’s way of encoding structure would take on the shape of the structures encoded? Isn’t this a straightforward vehicle/content conflation, and empirically implausible? The representation of structure is one thing, the structure of representations another.

We have cognitive representations of all sorts of things: the visual layout of the room, the positions of our limbs, the colour of the carpet. None of these representations have to take on the *form* of what they represent. Our representations of the carpet’s being green do not themselves have to be green. Nor do our representations of the spatial layout of the room have to be in any similar way spatially arrayed. Whatever

¹³ To conceive a speaker’s being sensitive to the type of information encoded in a correct theory of his language as his having conscious or unconscious knowledge of the theory is a particularly crude way to way to construe the relation between speakers’ knowledge and their grammars, and as such invites pretty feeble responses. (See Hornsby 2005, 109 as quoted in Matthews 2006, 458.)

form they take, they have to *encode* information about the relative positions of objects to one another. Similarly, whatever form our linguistic representations have, their job is to encode information about the structural relations among constituents. That means there are cognitive states in speaker/hearers that represent syntactic structures. But in saying this, I am *not* claiming that these cognitive states represent anything external to the speaker/hearers.¹⁴

The encoding of structure will neither take the same form as the structure it represents, nor will it take the form of the theorists' specification of structure. However to say that the linguistic system does not contain propositional representations does not mean that it is 'non-intentional'. There are plenty of non-propositional though still intentional representations in the visual system, although these representations, unlike linguistic representations are also about something in the subject's environment. However, this need not be the case. There are percepts of motion created by the visual system, as when light seems to travel between two stationary light sources, switched on and off in sequence. At the right frequency the brain has no option but to interpret the signals as an instance of movement. A movement of light between the two light sources is what we 'see' in the resulting percept, even though the sources of light are stationary. In a similar way, hearing a sentence as structured, when what we are actually presented with is a continuous speech signal, is a matter of consciously experiencing a linguistic percept whose form and character are owed to the way the mind/brain represents constituent structure from the string of words we perceive in the sound signal.

There would appear to be no real alternative to a representational story about the mind/brain's handling of syntactic structure. For if the mind/brain satisfies 'the specification of a function that maps lexical items (clusters of features) onto pairs of structures that interface with external systems of articulation and conception-intention' ('Redux', 4) then it must represent or encode clusters of features and pairs of structures. Something in the mind/brain maps information of one kind onto information of another. It is this information that must be represented internally, with some of it—the lexical information—being 'stored'. (4). What is more, as an upshot of the mind/brain's encoding or storing of such information and deploying it when strings are produced or encountered the speaker/hearer will have conscious experiences of *hearing each string of sounds as structured*. The result is a percept depending on external prompting by the acoustic signal and on what is internally stored. How we hear a string as structured will depend on how the

¹⁴ In a slightly dangerous comparison one could claim that the marks of a canvas may succeed in depicting a scene even though the scene depicted does not exist in the world beyond the canvas. Of course, the scene might exist and the depiction could be better or worse. However, in the case of representational paintings, we tend to think of the depicted scene as being as if about a real scene. A case where no such match is envisaged could be an arrow sign on the road.

string of stored lexical items prompted by the acoustic input signal are parsed, and this depends in part on the assignment of linguistic structure the mind/brain makes to that string.

Collins says: 'Part of what it means to have a non-intentional conception of the language faculty is that one does not take its states to have correctness conditions; they are not answerable to an independent realm about which one might go right or wrong.' (21) But because states of the language faculty are not answerable to 'an independent realm' that doesn't mean they are non-intentional. They represent certain dependencies among selected lexical items. And how they represent things is how they are. There are no further correctness conditions to be met. The content of the representations, what is represented, namely, the syntactic structure of expressions, does not need, in addition, to match some independently existing structure of an external language. States of the language faculty have contents but not contents that represent any other linguistic reality. On the contrary, linguistic reality is (in part) constituted by how these states represent things to be. The sub-personal organisation of the language faculty constrains the representations that can be constructed when certain lexical items, or clusters of features, in whatever format, are tokened together. These representations of the relations among the features don't get things right or wrong, they simply represent certain syntactic relations and dependencies to hold among the constituents. Without getting something right or wrong we may wonder how they can have representational content. But this question targets the wrong states. We need to ask this question about states for which the question of knowledge arises, and we need not treat all representational states as knowledge states. There are representations in fiction that cannot be said to have correctness conditions or give us knowledge of anything external to the way they represent things.

Maybe it could be said that if there were people or places like those represented in the fiction these representations would portray them correctly or incorrectly. However, it is only because novels and plays portray scenes or events for which we can entertain the idea that these things could have happened like that that we think of what they represent as if it also corresponded to real states of affairs. But this idea of things featuring in the contents of representations also having to match something external to the representational system is an additional requirement and need not be at the heart of the notion of representations. There need not be as if aboutness of anything external to the fictional discourse, painting or mental state, for there to be a content represented by the state, painting or discourse. This additional element overshoots the mark. Representational states must have aboutness, and the states we are interested in are about grammatical properties and dependencies. This is genuine aboutness. Though, in order to have this we do not need to have the represented structures be, in addition, as if about the structure of a Platonic, or otherwise external language. Rather, the

internal states store information about lexical items and construct representations of the intricate syntactic relations among them when they are combined to form complex expressions. The encoded information is represented but Collins is right ‘the internal states of the faculty do not function as if to be about anything’ (10). Where he goes wrong is thinking that such a strong condition of being ‘as if’ about an independent realm is required for the internal states of the faculty to represent things in a certain way in the language faculty.¹⁵

Representing in the faculty doesn’t need to assume that what is represented goes proxy for something existing beyond it. Many things get represented that aren’t about anything else: the body’s temperature can be represented as high. This level is not about anything else—even though it can be correlated with sweltering heat outside or fever within. We do not need to assess representational states as if they were about a fictional or another possible world. This is not the only model of representing. Nevertheless, it is the content of representations that makes them representations of one thing and not another—representations of grammatical and not spatial relations, say.

So perhaps Collins and I have a merely verbal disagreement. If he could agree that there is a difference between a representation and what it represents (but see ‘Redux’, 20, footnote 26); something that makes it a representation of X and not a representation of Y, without that having to require additionally that we suppose X correspond to anything external to the system, or that the representations must be as if representations of some independent realm, then I think we can accept much of what he says and simply drop the paradoxical sounding non-relational talk of representations as not being representations of anything. If states of the language faculty take ‘Bill’ to be the agent of ‘want’ and not ‘seem’ and posit PRO controlled by ‘Bill’ and not ‘Peter’ in:

3. Bill seemed to Peter to want to leave

that is because those states represent the syntactic relations between these items in this way. No external, public or Platonic language need be invoked.

Chomsky’s reasons for anti-representationalism may be that since there is no independent language for the language faculty to provide knowledge of, the states of the faculty cannot be representations of anything either. The worry might be that if we say our states of the language faculty are representations of something we would be tempted to say that the states of the faculty gave us knowledge of that

¹⁵ Here, I am firmly in agreement with Georges Rey (2006) about such states having intentional contents and representing properties of linguistic expressions. Where I depart from him is in his terminology of ‘intentional inexistent’ and his treatment of these representations as like illusions. Like so many others, Rey seems to accept the additional ‘as if’ criterion for genuine representational content. By my lights, something like illusions, or better percepts enter at the personal level, not at the sub-personal level of the language faculty.

same thing. But the worry behind this thought is unfounded. Firstly, the representations are not of anything external to the mind of the speaker, so it need not license a view of language at odds with Chomsky's internalist conception of language. Secondly, genuinely representational states need not amount to knowledge states. And thirdly, we have seen that for personal level linguistic judgements to be answerable to how things are portrayed by representational states of our language faculties there must be a way those states represent things to be. The representational states of the internal linguistic system represent syntactic structure—structural relations between items drawn from the lexicon—and are answerable to nothing beyond themselves. How those representations represent things syntactically is how they are. And it is the content of those representations that makes them representations of syntax. On the other hand, as speaker/hearers, we have experiences of uttering or hearing sentences, and the intuitive judgements we form when attending to sounds we experience linguistically are answerable for their correctness to whether the string we judge acceptable or not really is a permissible grammatical arrangement of items as determined by our I-language.

Finally, Collins tells us his account of the underlying linguistic system does not 'impugn our epistemic phenomenology. The moral is only that we should not take the constitutive rationality of our manifest image to constrain our empirical theorising.' ('Redux', 21) I agree, it does not constrain our empirical theorising but it can be evidence. The ways we can and cannot hear strings can tell us something about the constraints the grammar imposes. Collins points out that:

Just how the faculty contributes to our personal-level states as knowers of language is as complex a problem as is possible to imagine...[and] surely we should expect our high-level judgements to be an ensemble effect. (21–2)

Certainly we should, but we should also expect the language faculty as part of the whole ensemble of systems to regularly deliver judgements about permissible ways to combine word-like items in a string, and not mislead us. Were this not the case, a corpus of a speaker's use of language would provide better evidence than the speaker's linguistic intuitions: something neither of us believe.¹⁶

Bibliography

- Chomsky, N. (1959). 'Review of B. F. Skinner's *Verbal Behaviour*.' *Language*, 35, 26–58
- Chomsky, N. (1965). *Aspects of the Theory of Syntax* (Cambridge, MA: MIT Press).

¹⁶ I would like to thank members of the Philosophy of Linguistics group in the University of London for very helpful discussions of these issues, especially Gareth Fitzgerald and Nick Jones. I am also very grateful for discussions both at and after the Dubrovnik Philosophy of Linguistics conference in 2007. Special thanks to John Collins, Ophelia Deroy, Michael Devitt, Dunja Jutrović, Peter Ludlow, Robert Matthews, and Georges Rey.

- Chomsky, N. (1968). *Language and Mind* (New York: Harcourt, Brace, Jovanovich).
- Chomsky, N. (1980). *Rules and Representations* (New York: Columbia University Press).
- Chomsky, N. (1981). *Lectures on Government and Binding* (Dordrecht: Foris).
- Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin and Use* (Westport: Praeger).
- Chomsky, N. (2000). *New Horizons in the Study of Language and Mind* (Cambridge: Cambridge University Press).
- Collins, J. (2004). 'Faculty disputes.' *Mind and Language*, 19, 503–533.
- Collins, J. (2006). 'Between a rock and a hard place: a dialogue on the philosophy and methodology of generative linguistics.' *Croatian Journal of Philosophy*, 6, 471–505.
- Collins, J. (2007). 'A review of Michael Devitt's *Ignorance of Language*.' *Mind*, 116, 416–423.
- Devitt, M. (2006a). *Ignorance of Language* (Oxford: Oxford University Press).
- Devitt, M. (2006b). 'Defending Ignorance of Language: responses to the Dubrovnik Papers.' *Croatian Journal of Philosophy*, 6, 571–606.
- Dummett, M. (1976). 'What is a Theory of Meaning II?' In G. Evans and J. McDowell, *Truth and Meaning* (Oxford: Oxford University Press).
- Dummett, M. (1978). 'What Do I Know When I Know a Language?' Reprinted in M. Dummett, *The Seas of Language* (Oxford: Oxford University Press, 1993).
- Johnson, K. (2007). 'Tacit and Accessible Understanding of Language.' *Synthese*
- Matthews, R. (2006). 'Could competent speakers really be ignorant of their language?' *Croatian Journal of Philosophy*, 6, 457–468.
- Rey, G. (2006). 'The intentional inexistence of language—but not cars'. In R. Stainton (ed.), *Contemporary Debates in Cognitive Science* (Oxford: Blackwell), 237–255.
- Smith, B. (1998). 'On Knowing Our Own Language'. In C. Wright, B. Smith and C. Macdonald (eds.), *Knowing Our Own Minds* (Oxford: Oxford University Press).
- Smith, B. (2006a). 'What I know when I know a language'. In E. Lepore and B. Smith (eds.), *Oxford Handbook to the Philosophy of Language* (Oxford: Oxford University Press), 941–982.
- Smith, B. (2006b). 'Why we still need knowledge of language.' *Croatian Journal of Philosophy*, 6, 431–456.