



Epistemic Heresies: Reply to John Collins' Redux

«Epistemic Heresies: Reply to John Collins' Redux»

by Robert J. Matthews

Source:

Croatian Journal of Philosophy (Croatian Journal of Philosophy), issue: 22 / 2008, pages: 45-55, on www.cceol.com.

*Epistemic Heresies: Reply to John Collins' Redux*¹

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

Elaborating on views I have expressed elsewhere, I argue that the common-sense notion of linguistic competence as a kind of knowledge is both required by common-sense explanatory and justificatory practice and furthermore fully compatible with the non-intentional characterization of linguistic competence provided by current linguistic theory, which is itself non-intentional.

Key words: Linguistic competence, knowledge of language, linguistic theory as non-intentional, common-sense conception of competence as knowledge.

Common English parlance, John Collins tells us, has it that people know languages, that ships sail rather than swim, that smokers inhale rather than swallow, but “none but the deluded would explicitly take common parlance to be a guide for theoretical reflection” (Collins 2008, 3).² Yet some of us, Collins claims, persist in thinking that linguistic competence is a matter of knowing something about the language in question, thereby supposedly taking common parlance “to constrain theoretical inquiry or, indeed, to mark out the proper domains of such inquiry” (ibid.). Collins is surely right to caution us not to put much stock in common parlance as a guide to or constraint on theoretical inquiry, especially as regards language. For it might be mere accident that speakers of English speak of ‘knowing’ a language, whereas speakers of French, Italian, German, and Spanish speak simply of ‘speaking’ a language. And yet for all that, linguistic competence *might* have something to do with knowledge of language. Common parlance, after

¹ This response has benefited greatly from the comments and criticisms of both John Collins and Frances Egan.

² Unless otherwise stated, all subsequent references to Collins are to this 2008 paper.

all, is not always wrong. Smokers (other than Bill Clinton, who claimed never to have inhaled) do inhale rather than swallow. And ships do sail rather than swim, both because lacking the requisite appendages ships are incapable of swimming and because until relatively recently all ships were sail-powered. So maybe, contrary to what Collins imagines, the common English parlance of knowing a language is onto something about linguistic competence, and exploration of the idea that linguistic competence involves knowledge of some sort might prove fruitful. At very least, the idea of knowing a language has this much going for it: the notion of competence as knowledge is suitably divorced from actual speech as to permit one to be linguistically competent, in the sense relevant to linguistic theorizing, yet unable to speak.

Chomsky himself seems at one time to have thought that linguistic competence had something to do with knowledge of language, even if he no longer does. Collins denies this, insisting that Chomsky has never thought seriously of linguistic competence in such terms, that he has not, as I claimed (in Matthews 2006), “changed his mind” on this matter. Perhaps Collins is right about Chomsky’s never having changed his mind on this matter, but this much is clear: at one point Chomsky talked a lot about knowledge of language, how it was acquired, and how it was used—there was even a book entitled *Knowledge of Language*—and such talk was not at all disparaging of the idea that linguistic competence was a matter of knowing certain things, namely a grammar. So Chomsky must have at one point thought the idea of linguistic competence as knowledge worth exploring, or at least worth talking about.

Now, to say that linguistic competence is a matter of having knowledge of language is not yet to say anything about its kind or in what it consists. Certainly nothing requires that we assume, as most philosophers do, that there are two distinct kinds of knowledge, propositional knowledge and practical knowledge, only one of which is the kind of knowledge constitutive of linguistic competence. Nor need we assume that this knowledge is explicitly represented in (or by) its possessor, much less that it is so represented by certain quasi-linguistic mental representations that express the precise contents of what is known. And certainly nothing requires that linguistic competence turn out to be exhaustively characterized in terms of knowledge. The knowledge of language that we speak of as constitutive of linguistic competence might turn out to be only the consciously accessible aspect of an otherwise largely inaccessible linguistic competence, with much of the explanatory richness of any adequate characterization of linguistic competence being decidedly non-epistemic. These, I am certain Collins would agree, are all matters for empirical investigation, not bald pronouncement based on speculation either about the aptness of English colloquial expressions or about how in general knowledge might be physically embodied in a possessor.

Collins describes his analogy of the locution ‘knowledge of language’ with talk of the ‘sailing’ of ships or the ‘inhaling’ of smoke by smok-

ers as intended to underscore his claim that “theoretical inquiry into language should not be beholden to any substantiation of the apparent epistemic relation by which we commonly designate linguistic competence and that current inquiry does not reflect any such substantiation” (3–4). He concedes that Barry Smith and I do not think that it does, but he nonetheless discerns in the two of us a slight whiff of heresy, despite our being, as he puts it, “on the side of the angels”: we claim, he says, that “our conception of speaker/hearers as the object of linguistic theory requires a certain intentional optic” (4). Now, I’m not quite sure how to understand this nice turn of phrase, but speaking for myself, I am certain that I claim, and furthermore believe, nothing of the sort, though I freely concede that in matters ecclesiastical, the heretic is often blind to his heresy. More than a few have been surprised to find themselves hauled before an inquisitor. So let me state clearly: In my view, linguistic theory is non-intentional: in characterizing and explaining the linguistic competence of speakers, linguistic theory does *not* attribute intentional (contentful) states to those speakers. But, and this may be the point that Collins misunderstands about my view, linguistic competence, like other competences, is nonetheless often described informally in intentional terms, specifically in terms of knowledge. There are several reasons why common sense should opt for such descriptions, one of which is to allow the fact of a speaker’s linguistic competence to figure in epistemic justifications of one sort or another. Similarly, linguistic theory is often given an intentional gloss, typically in order to make clear the explanatory character of the theory, but also as a means of informal presentation. That linguistic competence and linguistic theory is often characterized in intentional terms does not challenge the current orthodoxy that Collins wants to defend. So with these abjurations made, let me respond to some of Collins’ specific criticisms of what he takes to be my view as well as comment on some of his own claims. As will be clear from what follows, our disagreements are minor relative to our points of agreement. I begin with his discussion of what makes for the non-intentional character of contemporary linguistic theory.

Collins notes that he and I are agreed that current linguistic theory is non-intentional. But it is not at all obvious that we agree about what linguistic theory being non-intentional comes to or what it is about linguistic theory that makes it so. Collins says this:

I take [the claim that current linguistic theory is non-intentional] to amount to the claim that linguistic theory is the attempt to specify an intensional functional that underlying neuronal systems of production, consumption, and understanding respect in the sense that the function encodes the structure of pairs of phonological-semantic representations that are explanatory over speaker/hearers linguistic judgements [sic?] and, in part, parsing performance, inter alia. This conception stands in contract to a ‘processing’ model insofar as specifying the function is not to specify *how* speaker/hearers process linguistic material. The function, rather, 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. (Collins 2008, 4–5)

There is no disagreement between us about the formal character of current linguistic theory. It specifies a speaker's linguistic competence by specifying intensionally (and no doubt under idealization) the pairing of sounds and meanings that the language faculty effects in the course of language processing, though without saying at all how this pairing is effected. But what makes linguistic theory non-intentional is *not* the fact that the specification of linguistic competence takes this particular form. It is compatible with a theory specifying linguistic competence in such terms that speakers are linguistically competent in virtue of being in certain intentional states such that being in these states explains why this intensional specification is true of them. That linguistic theory specifies linguistic competence in the way that it does entails nothing one way or the other regarding the intentional character of the states that are the truth-makers of this specification. What makes linguistic theory non-intentional has to do with the character of the internal states that are attributed to the language faculty by the theory: there is nothing essentially intentional about these states; these states are not in any essential way about anything else, e.g., about some external language ('E-language', as Chomskyans would have it); nor does construing the theory as explanatory of linguistic competence presume that they are in any way intensional.

The non-intentional character of current linguistic theory is not a peculiarity of linguistic theory. The same sort of reasons that dictate that theories of linguistic competence are non-intentional would entail that competence theories of other cognitive competences (vision, face recognition, deductive reasoning, motor control, etc.) are similarly non-intentional. Performance theories, including theories of linguistic performance, are similarly non-intentional. All cognitive computational theories worthy of the name, I would argue, are non-intentional.

Immediately following the previous quotation, Collins goes on to say this:

This point, of course, goes back to chapter 1 of *Aspects* (1965), where Chomsky makes clear that linguistic inquiry should partition the linguistic mind into at least two parts: a *competence* system, our intensional function, and *performance* systems, including at least a parser. (Collins 2008, 5)

Here, too, we disagree. Competence and performance, even as Chomsky construed them in *Aspects*, are *not* two parts of the 'linguistic mind'. There is but one language faculty, viz., the cognitive faculty responsible for language production and understanding. But that one faculty can be studied scientifically from different perspectives, with different explanatory goals in mind: most crucially here, it can be studied from the perspective of what the faculty is a capacity for, its *competence*, or alternatively from the perspective of how the faculty manages to possess this competence and how this competence gets deployed in the

course of language production and understanding, what has come to be called, perhaps misleadingly, its *performance*. Current linguistic theory characterizes competence intensionally by specifying a function from lexical items into sound-meaning pairs, but what it characterizes is the language faculty itself, and *not* a component or system of that faculty. It characterizes the language faculty, but in a way that abstracts away from all detail as to how this faculty manages to possess the competence that it does, by specifying (no doubt under considerable idealization) the pairing of sound and meaning effected by the language faculty.

Turning now to another matter, Collins bristles at my claim that “Collins (2004) argues, unconvincingly to my mind, that there has in fact been no change in Chomsky’s conception of linguistic competence over the years...” (Matthews 2006, 219, fn. 11). Collins insists:

I did not and do not argue that Chomsky has had a constant view of linguistic competence; my claim, rather, is that recent developments, in fact dating back to the late 1970s, allow one to see that the apparent epistemic relation of knowledge never played a serious role in theory construction and simply marked a puzzle or problem for acquisition. In other words, there were never any serious epistemic or intentional commitments, but there were serious disputes about knowledge of language, for that notion appeared to be required to keep the topic of inquiry in view, at least if the only apparent alternatives were species of behaviourism or moronic philosophical quietism. (Collins 2008, 5–6)

Collins’ response here confirms the claim that I attributed to him, for as the unquoted remainder of my footnote makes clear, I was speaking precisely about Chomsky’s often expressed endorsement of an intentional and furthermore epistemic conception of linguistic competence. There cannot therefore really be any serious dispute about what Collins argued in his 2004 paper; after all, defending the view that I attribute to him, viz., that “there were never any serious epistemic or intentional commitments” (ibid.) on Chomsky’s part, is much the burden of that paper. The only dispute here is who’s right and who’s wrong about Chomsky’s earlier views. Collins is much more willing than I am to discover in Chomsky’s earlier writings, which certainly *appear* to endorse an intentional and furthermore epistemic conception of linguistic competence, a view entirely consistent with his more recent non-intentional conception. I don’t think it much matters who is right here. But without getting into a pointless exegetical exercise,³ let me just note that Chomsky was, as Rey (2003, 160) points out, quite firm in his rejection of my suggestion (in Matthews 1980) that language acquisition, and linguistic competence more generally, might better be construed in non-intentional terms (the words within quotation marks were mine):

We agree that, *at some level*, much of what is called ‘learning’ ... should be characterized in an ‘non-intentional physiological vocabulary....’ But I do not see that this amounts to abandoning a ‘rationalist’ account of language

³ Rey 2003 makes a pretty good case, in my view, for the claim that Chomsky *has* changed his mind on these matters.

acquisition in which 'the various processes ... are defined over ... contents [of a state]', and innate structure 'is characterized intentionally in terms of both the *content* of a state and the learner's *relation* to that content' (say, cognizing). (Chomsky 1980, 47)

Turning to a third and perhaps more important matter, based on a misreading of my 2006 paper, Collins mistakenly concludes that while I am willing to dispense with an attitudinal construal of linguistic theory, I remain wedded to a intentional (content) construal of linguistic theory, whereby the contents of linguistic theory are the intentional contents of mind/brain states. Were this so, my claim that linguistic theory is non-intentional would of course be so much word play. I believe no such thing. Let me be clear: the *theories* that linguists construct (like the theories that physicists construct) have contents, but it does not follow from this that linguistic competence, or the states constitutive of linguistic competence, have content. Collins is apparently misled by the following passage, which he quotes:

The crucial point of the gloss is to get the reader to see *that the content alone* of the non-intentional theory is sufficient for the explanatory task. It does this by getting him or her to see that if we were to think of the theory as the content of the appropriate epistemic attitude, then we would accept the theory as the theory of competence that it claims to be. (Matthews 2006, 213)

The confusion here is a simple one: Collins is confusing the trivial, uncontentious claim that linguistic theory, like any scientific theory, has content, i.e., is about whatever it's about, with the non-trivial, contentious claim that what the theory is about, viz., linguistic competence, or the states of the language faculty, has content. The point being made in this passage, a point elaborated in the context in which this quotation appears, and defended in detail by Egan (1995), is this: It may often not be apparent to an outsider that a proposed theory of some cognitive competence is capable of explaining what it is claimed to explain. This is especially true in the case of computational or formal theories. (Consider, e.g., the claim that linguistic theory explains a speaker's linguistic competence by specifying intensionally a function from lexical items to sound-meaning pairs, or the claim that a theory of vision explains our ability, as Marr put it, "to know what's where in our visual environment" by specifying the computational processes that map 2-dimensional arrays into 3-dimensional, object-centered representations.) The problem in part is that the explananda that the theory takes itself to explain are typically couched in intentional terms, e.g., Marr's "knowing what's where in our environment," such that given our common-sense practice of explaining intentional phenomena in terms of other intentional states, we are prone to expect that the explanantia themselves must accordingly be intentional.⁴ We can address this problem of the theory's not being recognized as explaining what it claims to explain, by giving the theory an intentional cast, embedding it within the context of

⁴ For a standard defense of the view that intentional explananda demand intentional explanantia, a view that I reject, see Graves *et al.* 1973.

one or more propositional attitude verbs. This gives the theory the sort of form that common sense seems to demand. It also gives us license to then gloss the theory so cast in ways that make it more intuitively accessible, allowing us to say such things as “the speaker knows in effect that movement is bounded by certain locality constraints,” “the visual system assumes that objects are rigid in translation,” and so on. This, I (and Egan) claim, is precisely how computational theories in cognitive psychology are typically presented, again not because the theories are genuinely intentional in the sense of attributing certain intentional states constitutive of the competence, but for practical purposes of, as Chomsky (2000, 161) puts it, “informal presentation, intended for general motivation.” Now in the quotation which Collins cites, I was simply pointing out that what’s doing the presentational and motivational work is not the embedding attitude (knowing, cognizing, etc.), but the content of what’s known (what’s cognized, etc.), which explains why Chomsky could always be so nonchalant about the relevant attitude type that a native speaker was said to bear to his grammar. But the crucial point here, again, is that no claim is being made to the effect that intentional (contentful) states are constitutive of linguistic competence.

The pragmatic strategy by which linguists present and motivate a competence theory first by casting it in intentional terms and then by glossing as necessary the contents of the theory is dictated and facilitated not only by the fact that the explanatory goals of the theory are typically formulated in intentional terms but also by our common-sense practice of characterizing competences, both physical and intellectual, in terms of knowledge (for details, see Matthews 2006). Thus, we describe our competence in swimming, chess, and so on both in terms of knowing how to swim, knowing how to play chess, etc. and in terms of knowing various things, e.g., that certain swimming strokes are more efficient than others, that certain chess moves in certain circumstances will lead to disaster. Now, philosophers, beginning perhaps with Ryle, have made much of the putative distinction between knowledge-how and knowledge-that, arguing at great length as to the kind of knowledge constitutive of linguistic competence. It is an open question in my mind whether there is any such principled distinction, and more importantly whether it will bear the epistemological weight or do the epistemological work that many suppose. But this much is clear: within common-sense practice, our characterization of competences seems not to worry too much about this putative distinction. Given the way we commonly conceive of and talk about competences, competence would seem to be a matter of know-how informed by relevant knowledge-that, though when we attempt to say in some detail just what the know-how constitutive of some competence comes to, we find ourselves quickly forced to the knowledge-that locution. It is as if the relation of knowledge-that to knowledge-how is more intimate than simply the one informing the other, which should give us pause when we imagine that there are two distinct sorts of knowledge, practical knowledge and propositional

knowledge, that in some mysterious fashion somehow come together to constitute the competence.

This intuitive, commonsensical way of thinking about competence in terms of knowledge underpins the idea that linguistic competence is a matter of knowing a language, an idea that in turn underpins our common-sense practice of explaining and justifying in epistemic terms action that manifests linguistic competence. That practice, I argued (in Matthews 2006), needs this notion of competence as knowledge, even if knowledge plays only a pragmatic (informal presentational) role in explicating linguistic theory. Collins has little patience with the idea that knowledge of language might play some role in explaining and justifying action. In my discussion of these matters, I said this:

As competent speakers we are in most causes *authoritative* both about what we ourselves and about what others with whom we converse say. And it is in virtue of our being authoritative in this regard that it is rational for us to rely on what we take ourselves and others to say. Such authority is grounded in our linguistic competence. (Matthews 2006, 215)

I go on to say that when it comes to justifying our reliance on language as a reliable means of communication, “any explanation of the role of linguistic competence as the ground of this authority is going to have to characterize this linguistic competence in epistemic terms” (ibid.). Collins concedes that we do take ourselves to have authority over what we say, but he complains that “the distance between this fact and the putative evidence for it is no distance at all” (14). Apparently Collins imagines that we cannot sometimes justify our claim to know what we or another said, simply by appealing to the fact that we are competent speakers of the language in question. But surely an English-speaking witness testifying at a trial might, if challenged, justify his claim to know what one Pashto-speaking defendant said to another by pointing out that he is also a competent speaker of Pashto, just as in response to a different sort of challenge this witness could point out that he knows what one defendant said to the other because he was standing within easy earshot of them. No doubt an epistemologist might discover all sorts of defects in the sort of justifications I am describing here, but it seems undeniable that our linguistic competence does ground and justify our reliance on language as a reliable means of communication. How else are we to explain my confidence in acting on the words of another in a language I know (i.e., in a language in which I am competent)? How else are we to explain my reluctance to act on the words of another, when those words are in a language I don't know or don't know well? It's not for nothing that I would be very reluctant to rely on directions given to me in German, somewhat cautious and tentative in relying on directions given to me in French.

Collins concedes that we *know* the pairings of sounds and meanings in our language, thus conceding my central point, namely, that we do have the sort of knowledge of language that I claim we have. But he goes on to say:

But no epistemic appeal is required to what effects the pairing in order for us to have authority over our employment of such pairings. The knowledge, along with whatever ground for it we like, simply arise from our subpersonal systems being in working order, as it were. (Collins 2008, 14)

But the epistemic appeal that Collins is apparently attributing to me here is not one that I made. My claim was that we are justified in relying on language as a reliable means of communication because we know the language being used, where the knowledge in question is knowledge of the pairing of sounds and meaning effected by the language. No appeal whatever is being made to whatever effects this pairing, for the simple reason that, as I emphasized, for the purposes at hand, being competent in the language *just is* a matter of knowing this pairing. So why the dust-up here? It is presumably because Collins is concerned that communicative authority, which he takes to be a linguistic phenomenon that is “essentially epistemic insofar as it can be in view as a phenomenon,” not be allowed to justify the resuscitation of knowledge of language as an explanatory notion “just because linguistic authority might submit to some explanation some day” (pp. 14–15). If this is Collins’ worry, then it is unwarranted. Even if linguistic authority, understood here as a fact about our relation to the language we speak, were ever to submit to a naturalistic explanation, the explanation would presumably be every bit as non-intentional as other scientific explanations of linguistic phenomena. To think otherwise is to think that talk of knowledge of language in these contexts would carry some implication about the intentional (contentful) character of the states and processes constitutive of linguistic competence not carried by talk of knowledge of language in other contexts.

Finally, turning to informant judgments, which provide the data for linguistic theorizing, Collins takes issue with my claim (Matthews 2006, 215–6) that linguists need some justification for crediting these judgments as an accurate expression of competence in the language under study. I argued that there would seem to be no reason to credit these judgments as being such unless these informants actually *know* certain things about their language, and furthermore their judgments actually express this linguistic knowledge. Collins doesn’t like this, my second way of establishing the claim that to be linguistically competent is to have knowledge of one’s language. Collins opines, “I am afraid that I find this too philosophical by half. Generally, in science no issue arises about *justification of data*” (15). In general, yes. But we accept the informant’s judgments as evidence about the language under study only on the assumption that the informant knows the language and furthermore his judgments express or reflect that knowledge. For if that assumption were not operative, why should linguists bother eliciting the linguistic intuitions and judgments of a native speaker? Anyone would do as well. To be sure, issues of justification do not regularly arise in the course of linguistic inquiry, but that is because linguistic methodology constrains the collection of data in such fashion as to insure that we elicit data only

from individuals for whom this justificatory worry does not arise. If I want to study the phonology of Catalan, I seek out a native speaker of Catalan, someone who knows Catalan. We don't put much stock in the linguistic judgments of those for whom we question their knowledge of the language, for whom, as we could also put it, we question their competence in the language in which we are interested. Collins concedes as much when he acknowledges (15) that we want our informants to be native speakers of the language at issue. But if Collins is prepared to concede this, then what is it that vexes him so terrifically about my claim that competent speakers of a language know their language? Apparently it is the thought that knowledge of language requires an object to which the possessor of this knowledge is epistemically related. But what could that object be? Language, you say, but understood how? There seem to be no good options: Current orthodoxy precludes identifying language with the products of the language faculty, but it seems forced to think that the epistemic relation is one that holds between the possessor of linguistic knowledge and the possessor's language faculty. But the latter is precisely the relevant relation, as Collins himself admits, provided we understand the language faculty under this epistemic conception in the right way: "a non-intentional conception of the language faculty does *not* involve a denial of our knowledge of language, if such knowledge amounts to intuitive judgment of the kind that serves as data for the linguist" (18–19). This is exactly the sort of linguistic knowledge that will underpin our justification in relying on language as a reliable means of communication, exactly the sort of knowledge that an informant must have if we are to rely on this informant as a reliable source of data for linguistic theorizing. Put another way, this is the only knowledge of language that we expect our native informant to possess. Acknowledging this sort of knowledge of language in no way impugns a non-intentional conception of the language faculty; nor, as Collins imagines (p. 19), does it entail that whatever relation holds between the language faculty and this knowledge be epistemic. The latter is simply an expression, a product of the former.

Collins has suggested to me that my notion of knowledge of language is so "deflationary" as to render my claim that linguistic competence is a matter of knowing a language little different from his epistemically "nihilistic" conception of linguistic competence. (His conception might better be termed 'eliminativist'.) I think Collins is right to underscore our fundamental agreement on the non-intentional character of linguistic theory. But it nonetheless seems to me that we miss something important if we fail to appreciate that our common-sense conception of competence does take competence to be a matter of knowledge and furthermore that this conception plays a role in common-sense explanatory and justificatory practice. That this common-sense conception of competence deploys a notion of knowledge that is deflationary when measured against certain more extravagant conceptions of knowledge popular among some philosophers simply underscores theoretical innocuousness of common-sense talk of knowing a language.

References

- Chomsky, N. (1980) "Précis of *Rules and Representations* with Replies to commentators.", *Behavioral and Brain Sciences*, 3, 1–61.
- Chomsky, N. (2000). *New Horizons in the Study of Language* (Cambridge: Cambridge University Press).
- Collins, J. (2004). "Faculty Disputes." *Mind and Language*, 19, 503–33.
- Collins, J. (2008). "Knowledge of Language Redux" (this volume).
- Egan, F. (1995). "Computation and Content." *The Philosophical Review*, 104, 181–203.
- Graves, C., Katz, J., Nishiyama, Y., Soames, S., Stecker, R., and Tovey, P. (1973). "Tacit Knowledge." *Journal of Philosophy*, 70, 318–30.
- Matthews, R. (1980). "Language Learning *versus* Grammar Growth." *Behavioral and Brain Sciences*, 3, 25–6.
- Matthews, R. (2006). "Knowledge of Language and Linguistic Competence." *Philosophical Issues*, 16, *Philosophy of Language*, 200–20.
- Rey, G. (2003). "Intentional Content and Chomskyan Linguistics." In A. Barber (ed.), *Epistemology of Language* (Oxford: Oxford University Press), 140–86.