

An early version of this paper was published in *Language Sciences* 24 (2002) 97–106.
The present, revised, version appears in Savas L. Tsohatzidis, *Truth, Force, and Knowledge in Language: Essays on Semantic and Pragmatic Topics* (Berlin & Boston: De Gruyter 2020).

Grammars as objects of knowledge: The availability of dispositionalism

Savas L. Tsohatzidis

1. Introduction

As is well known, Chomsky and several linguists following him believe that the rules and principles comprising the grammar of any given natural language are things that are *known* by speakers of that language—and that speakers possess that knowledge independently of any linguist’s succeeding, or even trying, to formulate the rules and principles in question. Chomsky and his followers, however, do not believe that all reinterpretations of their professed belief are correct. In particular, they are strongly opposed to a reinterpretation of their professed belief that would appear to radically impoverish its content by adopting the—on at least one construal, characteristically Wittgensteinian—view that ascriptions of knowledge in a given domain are logically indistinguishable from ascriptions of dispositions to behave in certain ways in that domain (cf. Wittgenstein 1953, 1980). For, given that view, the claim that the rules and principles comprising the grammar of a natural language are things known by speakers of that language, can easily be taken to amount to nothing more than the claim that these speakers are disposed to exhibit certain patterns of verbal behaviour that happen to be correctly predictable (under appropriate idealisations) on the basis of those rules and principles, *without* those rules and principles being in any way represented in the minds or brains of the speakers. And this certainly appears to be an impoverishment of the idea—recurring with remarkable stability in Chomsky’s works—that the sorts of rules and principles that he and his followers are proposing do not merely constitute appropriate bases for correct predictions about linguistic behaviour, but are also *causally* involved in the production of that behaviour (an involvement that would, of course, be impossible unless the rules and principles in question were somehow represented in the minds or brains of the subjects whose behaviour they are alleged to causally influence).

Chomsky’s reason for rejecting the suggestion that a system of grammatical rules and principles, if accepted at all, might be construed merely as predictively reliable and not as causally efficacious is not an empirical one. For, as Chomsky is aware, no one has yet succeeded in locating the neural embodiments of any specific grammatical rules and principles, and no one is, accordingly, in a position to empirically defend the claim that a particular system of rules and principles, as opposed to innumerable others with the same predictive power, is causal-ly involved in the brain’s control of linguistic behaviour. What is more, the very idea that *some* such system, whether or not it is at present neurologically identifiable, *must* be supposed to be involved in the brain’s control of linguistic behaviour appears to progressively lose its original aura of inevitability with the advent and expansion of connectionist approaches to brain modelling (which, incidentally, may be

seen as offering perspectives complementary to the Wittgensteinian ones; cf. Mills 1993); for, in many cases, these approaches have been successful in showing that the assumption that the brain controls a pattern of behaviour which is describable by means of a given set of rules or principles is consistent with the assumption that the brain does *not* refer to representations of *any* rules or principles in order to exercise its control; if, therefore, these approaches are correct even in some cases, they certainly block any *automatic* inference from the describability of a behavioural pattern by means of a rule system to the conclusion that the organisms exhibiting the pattern possess and consult neural representations of that rule system.

Chomsky, however, believes that, although he is not, given the present state of the neurosciences, in a position to demonstrate that a proposed system of grammatical rules and principles is indeed causally efficacious, he *is* in a position to justifiably assert, in advance of any empirical inquiry, that such a system's predictive utility does *not* exhaust its content. For, the claim that such a system's predictive utility does exhaust its content—a claim, it should be noted, that would also be embraced by many philosophers, like Quine (1970), who could hardly be described as typical Wittgensteinians—crucially depends, in Chomsky's opinion, on the Wittgensteinian idea that ascriptions of knowledge in a given domain are logically indistinguishable from ascriptions of dispositions to behave in certain ways in that domain. And, according to Chomsky, there is a decisive *a priori* argument against that idea, at least as far as the domain in question is the domain of language—there is, in other words, a decisive *a priori* argument to the effect that the concept of linguistic knowledge cannot be identified with the concept of a disposition to exhibit linguistic behaviour.

After describing how this *a priori* argument was formulated in Chomsky (1980), criticised in Kenny (1984), and defended against Kenny's critique in Chomsky (1988), I shall inquire in what follows whether Chomsky's defence is successful, reaching the result that it is not. My first conclusion will therefore be that, despite appearances, the dispositionalist position is *not* vulnerable to Chomsky's critique. I will then examine whether the dispositionalist has the conceptual resources not only for countering Chomsky's attack but also for mounting an original attack of his own against the anti-dispositionalist position that Chomsky represents; and, finding that he does indeed have such resources, I shall derive my second, complementary conclusion, that it is Chomsky's position, rather than the dispositionalist's, that emerges as the really weak one in this particular debate.

2. The terms of the dispute

Chomsky's claim that a person's linguistic knowledge cannot be identified with that person's disposition to exhibit linguistic behaviour comes as a result of the following thought experiment:

Imagine a person who knows English and suffers cerebral damage that does not affect the language centers at all but prevents their use in speech, comprehension, or let us suppose, even in thought. Suppose that the effects of the injury recede and with

no further experience or exposure the person recovers the original capacity to use the language. In the intervening period, he had no capacity to speak or understand English, even in thought, though the mental (ultimately physical) structures that underlie that capacity were undamaged. Did the person know English during the intervening period? (Chomsky 1980: 51)

Chomsky's answer is that the protagonist of the above thought experiment did know English during the intervening period, and that, therefore, linguistic knowledge cannot be identified with the capacity to exhibit linguistic behaviour.

Kenny's reply is that Chomsky's argument cannot impugn the propriety of that identification, since one can legitimately maintain that the identification is correct while plausibly redescribing the imagined situation by invoking a distinction—whose independent desirability was pointed out by Kenny (1975) at an earlier time, and which was further elaborated in Kenny (1989)—between the existence or non-existence of a behavioural capacity and the existence or non-existence of the capacity *to exercise* that behavioural capacity. Considering Chomsky's question, "Did the person know English during the intervening period?", Kenny responds:

Well, we can say what we like as long as we know what we are doing: it is up to us to decide whether what is left is sufficient for us to call it "knowledge of English". Perhaps Chomsky is right that the more natural decision is to say that it is sufficient. Fine, then, let us say that the person knows English. But why should we not also say that the person retains the capacity to speak English? For extraneous reasons, he cannot use or exercise that capacity at the moment; but since, *ex hypothesi*, he is going to use it in future without any of the normal acquisition processes, is it not natural to say that he still holds on to it in the meantime? The concept of *capacity to use English* has exactly the same fuzzy edges as the concept of *knowledge of English* and Chomsky's example does nothing to separate the two concepts. (Kenny 1984: 141)

Chomsky's response to Kenny's reply is that it is ineffectual, since the distinction between the existence of a capacity and the existence of the capacity to exercise that capacity does not correspond to anything present in our ordinary concept of a capacity, and merely represents an arbitrary move designed to save the dispositional account. After repeating what he regards as the proper conclusions to draw from his thought experiment, Chomsky notes:

To avoid these conclusions, philosophers committed to the identification of knowledge and ability have been forced to conclude that [the person] who lost the ability to speak and understand [English] after brain injury in fact retained this ability, though he lost the ability to exercise it. We now have two concepts of ability, one referring to the ability that was retained and the other to the ability that was lost. The two concepts, however, are quite different. It is the second that corresponds to ability in the sense of normal usage; the first is just a new invented concept, designed to have all the properties of knowledge. Not surprisingly, we can now conclude that knowledge is ability, in this new invented sense of "ability" that is quite unrelated to

its normal sense. Plainly nothing is achieved by these verbal maneuvers. We must conclude, rather, that the attempt to account for knowledge in terms of ability (disposition, skill, etc.) is misconceived from the start. (Chomsky 1988: 11–12)

Indeed, not only is Chomsky convinced that the Wittgensteinian account of knowledge in terms of ability is vitiated by being forced to adopt a conception of “ability” that diverges, according to him, from “normal usage”, but he also contends that, by being forced to adopt such a divergent conception, it contradicts Wittgenstein’s own much more general thesis that departures from “normal usage” should be always avoided because they are among the primary sources of conceptual confusion. As Chomsky puts it in a later presentation of his response to Kenny’s reply, “the Wittgensteinian construal of knowledge as a species of ability seems to be a paradigmatic example of the practice that Wittgenstein held to be a fundamental source of philosophical error” (Chomsky 1992: 104). Apparently, then, Wittgensteinians would have to concede not merely a local but a global defeat if Chomsky’s suggestions could withstand scrutiny—which is one more reason for finding out whether in fact they could.

3. The structure of the dispute

Let Q stand for the person involved in Chomsky’s thought experiment and *t* for the period between that person’s injury and recovery. Chomsky’s original argument can then be summarised by saying that, since statement (1a) below is true of Q at *t*, and since statement (1b) below should also be true of Q at *t* *if* the dispositional account of knowledge was correct, proponents of that account would be constrained by the canons of deductive reasoning to accept as also true of Q at *t* the statement in (1c), even though that statement is, in Chomsky’s view, *not* true of Q at *t*.

- (1) a. Q does not have the ability to use English
- b. Q knows English if and only if Q has the ability to use English.
- c. Therefore, Q does not know English.

In his reply, Kenny grants to Chomsky the right to assert that Q does know English at *t*, but insists that saying of Q that he cannot *exercise* the ability to use English at *t* (rather than that he does not *have* the ability to use English at *t*) is all that is required to truthfully represent Q’s predicament. Once this is admitted, however, there is no way—Kenny suggests—in which a dispositionalist could be threatened by the inference that Chomsky is planning to impute to him. For, the inference in question will then have to take either the form in (2), which, though faithfully representing the dispositionalist position in premise (2b), is *not* deductively valid, or the form in (3), which, though deductively valid, does *not* faithfully represent the dispositionalist position in premise (3b):

- (2) a. Q cannot exercise the ability to use English.
- b. Q knows English if and only if Q has the ability to use English.
- c. Therefore, Q does not know English.

- (3) a. Q cannot exercise the ability to use English.
- b. Q knows English if and only if Q can exercise the ability to use English.
- c. Therefore, Q does not know English.

In his rejoinder, Chomsky correctly perceives that the only way to counter Kenny's objection is to deny that, as far as our ordinary concepts are concerned, there is a real distinction to be made between the idea of a person's *being able or unable to exercise an ability* and the idea of that person's *having or not having that ability*. He therefore thinks he can dispose of Kenny's objection by swiftly proceeding to that denial: there can be no possible difference, Chomsky contends, between the existence or non-existence of an ability and the existence or non-existence of the ability to exercise that ability; and since no such difference can possibly exist, Kenny is not entitled, Chomsky argues, to invoke it in order to defuse his objection to dispositionalist accounts of knowledge ascriptions. The question before us is, then, whether Chomsky is justified in his claim that statements of the form (4) are equivalent to corresponding statements of the form (5)—in other words, whether he is justified in his claim that every statement of the form (6) is necessarily true:

- (4) x has the ability to f
- (5) x can exercise the ability to f
- (6) x has the ability to f if and only if x can exercise the ability to f

4. A dispositionalist defence

In order to show that Chomsky's equivalence thesis does not in fact hold, it would be sufficient to show that there are circumstances where a statement of the form (4) would be true whereas a corresponding statement of the form (5) would be false. Showing this, however, is not at all difficult, since circumstances where a subject's inability to exercise a certain ability is consistent with that subject's continued possession of that ability are not at all rare.

The abilities (or skills) to play football, to conduct an orchestra, to swim and to dance are certainly abilities that can truthfully be ascribed to persons. Now, each person to whom these or any other abilities are truthfully ascribed is, like any other person, an organism whose continued existence requires the regular immersion in states of sleep. During sleep, none of the abilities just mentioned can, of course, be exercised (and, during deep sleep, a host of many other, even more common, abilities cannot be exercised either). From the fact, however, that, during sleep, an ability cannot be exercised, it hardly follows that it is *lost*. We would certainly not say that a football player, an orchestral conductor, a swimmer and a dancer lose their playing, conducting, swimming and dancing skills every time they go to sleep, and newly acquire them every time their sleeping sessions are over. Indeed, if these abilities were literally *lost* during each sleeping session, it would be a miracle that their bearers instantly reacquire them after each sleeping session, since these abilities cannot be acquired *at all* without

considerable training (which is why they are commonly described as skills). The solution, of course, is to say that what the subjects in question lose when they fall asleep are *not* the above mentioned abilities but rather the ability *to exercise* those abilities—and, correlatively, that what they newly acquire when they wake up are *not* the above mentioned abilities but rather the ability *to exercise* those abilities. It is in order to conceptually accommodate perfectly ordinary situations like these that the distinction invoked by Kenny has been designed. And since it does succeed in accommodating them, its legitimacy cannot credibly be denied.

On the contrary, it is Chomsky's refusal to accept that distinction that cannot be legitimised, since it leads to palpably absurd redescriptions of the perfectly ordinary situations we have been considering. For, given that, according to Chomsky, an ability exists if and only if it can be exercised, Chomsky would have to say that an experienced football player and an experienced orchestral conductor completely *lose* their playing and conducting skills *every* time they go to sleep (since, during sleep, these skills cannot be exercised). And since a completely lost skill cannot reappear unless it is newly acquired, Chomsky would have to explain the disturbing fact that, once awake, the player and the conductor are immediately capable of manifesting their playing and conducting abilities, by saying that, though completely lost during sleep, these complex abilities are *instantly* re-acquired by their bearers after *every* sleeping session (which, of course, amounts to saying that the ability to play football or to conduct an orchestra are, by turns, totally lost and totally regained during *each day* in the lives of their bearers). Now, "normal usage", which Chomsky appears willing to invoke on his behalf, may be notoriously tolerant, but even "normal usage" would not tolerate that ordinary lives are, as Chomsky's proposal implies, constant successions of miracles. Indeed, by providing the two *distinct* expression types 'x has the ability to *f*' and 'x can exercise the ability to *f* without in any way *imposing* on us to treat them, in Chomsky's way, as necessarily co-extensive, "normal usage" performs two excellent services: first, it provides us with the opportunity of *inquiring* whether they are in fact co-extensive; and secondly, it provides us with the means, once we realise that they are *not* in fact co-extensive (once we realise, for example, that the abilities we cannot exercise during sleep are not lost during each sleeping session and instantly reacquired after each sleeping session), of explaining why ordinary lives are *not* constant successions of miracles.

Our conclusion must be, then, that, since the distinction between a subject's possessing an ability and a subject's being able to exercise that ability is both real and vital in preventing the generation of manifest absurdities, Kenny had every right to invoke it in objecting to the anti-dispositionalist argument that Chomsky had built on the basis of his thought experiment; and that, consequently, Chomsky's attempt to save that anti-dispositionalist argument by denying the distinction's viability is unsuccessful. As a result, Kenny's original argument retains its full force, as well as the full range of its implications: Chomsky has *not* yet shown, contrary to what he and many of his followers may believe, that there is anything wrong with the proposal of identifying knowledge in the linguistic domain with the ability to exhibit certain behavioural patterns in that domain; and he has not, therefore, provided a decisive *a priori* reason for denying that a proposed system of grammatical rules and principles, if accepted at all, can be construed

merely as a basis for accurate predictions (under appropriate idealisations) of certain behavioural regularities and not as a causal force responsible for these regularities.

5. A dispositionalist attack

Of course, the fact that Chomsky has failed to show that the dispositionalist position must be ruled out does not in itself determine whether his own anti-dispositionalist position is or is not ultimately defensible. It would be interesting to examine, therefore, whether the dispositionalists would be more successful in arguing, in their turn, that it is the anti-dispositionalist position that ought to be ruled out. It seems, in fact, that they could do precisely that, by exploiting a subtle, and no doubt unintended, concession to dispositionalism that is implicit in the way in which Chomsky's thought experiment is set up.

Recall that the subject of Chomsky's thought experiment is understood as a subject who *did behaviourally manifest* its knowledge of English *before* its injury and also as one who *did behaviourally manifest* its knowledge of English *after* its recovery. It is only after describing the subject in such terms that Chomsky goes on to raise the question whether it knew English "during the intervening period"—during the period, that is, in which it did not behaviourally manifest its knowledge—and argues that the affirmative reply to that question, which he takes to be correct, creates a problem for the dispositionalist. Now, as we have seen, the dispositionalist is not in fact threatened by the affirmative reply, even if he accepts it as correct: provided that the subject's linguistic ability *was* exercised before and after the injury, the dispositionalist is under no pressure to deny that it existed, but could not be exercised, *during* the injury. It seems, however, that the dispositionalist can now create a problem for the position Chomsky himself represents—a problem, that is, for the position that linguistic knowledge *can* be characterised at all without *any* reference to its behavioural manifestations.

The dispositionalist would first observe that if, as Chomsky believes, linguistic knowledge could be characterised independently of any reference to its behavioural manifestations, then, in the description of Chomsky's thought experiment, the reference to the subject's pre-injury and post-recovery behavioural manifestations of its knowledge of English should be *inessential*: everything we would say about the subject's knowledge of English "during the intervening period" we should also be able to say *without* taking into account the fact that it did behaviourally manifest that knowledge before the injury and did behaviourally manifest the same knowledge after the recovery. But this radical separation of attributions of knowledge from references to behavioural manifestations—the dispositionalist would object—does not appear to be in fact possible, as we can appreciate by considering a different thought experiment, which abstracts away from precisely those factors that ought to be inessential on Chomsky's account:

A child is born to a couple of English speaking parents. A few weeks after its birth—and so, before the child is in a position to use English or any other natural language—it suffers cerebral damage which, though in all other respects does not prevent the natural growth of his brain, completely prevents, for ten consecutive

years, the use of English or of any other natural language “in speech, comprehension, or let us suppose even in thought”. On its tenth birthday—and so, while the effects of the injury are still fully in place—the child suddenly dies. Did the child know English or any other natural language at any point in its short life?

The answer to *this* question is, obviously, negative. A child who, at no point in its entire life, *used* a language—has never spoken it, has never understood it, has never even thought in it—cannot be said to *know* that language, no matter what the architecture of its brain is: even if a post-mortem examination of the child in question was to reveal that the so-called “language centres” in its brain (that is, the brain structures that are supposed to be linguistically relevant but strictly distinct from whatever brain structures are dedicated to the management of language *use*) were as well developed as anyone else’s in its community, the fact that the reputed “centres” were, in view of the child’s total incapacity for linguistic behaviour, *never* in its life connected to the production or understanding of any element of English or of any other natural language would surely suffice for concluding that the child never knew English or some other natural language. But this obviously true negative answer—the dispositionalist would observe—is clearly not one that Chomsky’s position would allow him to acknowledge. For the whole point of Chomsky’s anti-dispositionalist campaign is precisely to assert that a person’s *knowledge* of a language simply consists in the existence of certain specialised structures in its brain that are distinct from whatever brain structures may be dedicated to the management of language use, and can be fully characterised without reference to any aspect of the person’s behaviour (for example, producing sentences, understanding sentences, etc.) that constitutes language use. Assuming, therefore, that Chomsky would not want to concede that direct or indirect reference to language use is necessary for ascriptions of linguistic knowledge (a concession that would make his position, in all relevant respects, indistinguishable from the dispositionalist’s), the only option available to him when confronted with the above thought experiment would be to claim that, if the “language centres” in this child’s brain were found, after a post-mortem examination, to be as well developed as anyone else’s in its community, the child should indeed be credited with having possessed a great amount of linguistic *knowledge*, even though it was, throughout its entire life, fully unable to either produce or understand anything occurring in any natural language whatsoever. But this only shows—the dispositionalist would conclude—that if Chomsky was to maintain his position, he would have to use a concept of “knowledge” that is so idiosyncratic and obscure that any attempt at further communication with him on these matters would be bound to fail (as, in fact, many philosophers have, on independent grounds, long ago suspected that it would; cf., among others, Nagel 1969, Stich 1971, Cooper 1975).

It seems to me that this argument is very difficult to counter, and certainly impossible to ignore. If so, the conclusion we are finally entitled to draw is stronger than the one previously derived: not only has Chomsky failed to produce a valid argument against the dispositionalist construal of linguistic knowledge, but it is also the case that the dispositionalist can produce a very effective argument of his own against Chomsky’s anti-dispositionalist construal. To put it concisely, the relative strengths of the dispositionalist

and the anti-dispositionalist conceptions of linguistic knowledge appear to be exactly the opposite of what Chomsky thought they were.

References

- Chomsky, Noam. 1980. *Rules and representations*. New York, NY: Columbia University Press.
- Chomsky, Noam. 1988. *Language and problems of knowledge: The Managua lectures*. Cambridge, Mass.: The MIT Press.
- Chomsky, Noam. 1992. Language and interpretation: philosophical reflections and empirical inquiry. In John Earman (ed.), *Inference, explanation, and other frustrations: Essays in the philosophy of science*, 99–128. Berkeley, CA: University of California Press.
- Cooper, David E. 1975. *Knowledge of language*. London: Prism Press.
- Quine, W.V. O. 1970. Methodological reflections on current linguistic theory. *Synthese* 21. 386–398.
- Kenny, Anthony. 1975. *Will, freedom, and power*. Oxford: Blackwell.
- Kenny, Anthony. 1984. *The legacy of Wittgenstein*. Oxford: Blackwell.
- Kenny, Anthony. 1989. *The metaphysics of mind*. Oxford: Oxford University Press.
- Mills, Stephen. 1993. Wittgenstein and connectionism: A significant complementarity? In Christopher Hookway & Donald M. Peterson (eds.), *Philosophy and cognitive science*, 137–157. Cambridge: Cambridge University Press.
- Nagel, Thomas. 1969. Linguistics and epistemology. In Sidney Hook (ed.), *Linguistics and philosophy*, 171–182. New York: New York University Press.
- Stich, Stephen P. 1971. What every speaker knows. *The Philosophical Review* 80. 476–496.
- Wittgenstein, Ludwig. 1953. *Philosophical investigations*. Oxford: Blackwell.
- Wittgenstein, Ludwig. 1980. *Remarks on the philosophy of psychology*, Vols. 1 and 2. Oxford: Blackwell.