
Language and the Mind

A. J. P. Kenny

Phil. Trans. R. Soc. Lond. B 1981 **295**, 245-251

doi: 10.1098/rstb.1981.0137

Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click [here](#)

Language and the mind

BY A. J. P. KENNY, F.B.A.

Balliol College, Oxford OX1 3BJ, U.K.

‘We may usefully think of the language faculty, the number faculty, and others, as “mental organs”, analogous to the heart or the visual system or the system of motor coordination and planning. There appears to be no clear demarcation line between physical organs, perceptual and motor systems, and cognitive faculties in the respects in question.’ (N. Chomsky, *Rules and representations*, p. 39.) I shall argue that the view thus expressed is open to philosophical criticism that is not adequately rebutted by the lengthy philosophical discussion in *Rules and representations*. The notion of a mental structure appears to involve a philosophical confusion, one that is sometimes nicknamed the ‘hardware–software fallacy’; but this nickname is unfortunate since the postulated structures appear to be intended to be too ghostly to be hardware and too concrete to be software. Philosophical criticism of the notion of a mental structure does not imply that there cannot be illuminating structural descriptions assignable by grammar to linguistic expressions. But it concerns the characterization of the appropriate vehicle to embody the capacities that are expressed in linguistic performance, and thus the relation between mind and body. If the role of linguistics, like the role of psychology, is described in a way that is free of philosophically inappropriate analogy, then the discipline of linguistics can indeed be regarded as a branch of psychology; for language is the paradigm expression of mind.

To view the language faculty as an organ like the heart involves a deep philosophical confusion. Chomsky’s description of the mental structures that he investigates introduces an irrelevant metaphysical element at the interface between physiology and psychology. I intend to justify this complaint by a detailed examination of some crucial passages in Chomsky’s latest book, *Rules and representations* (1980). But before doing so let me, in order to avert misunderstanding, list a number of points on which philosophers have picked quarrels with Chomsky and on which I think it is he, and not his philosophical critics, who is in the right.

I have no quarrel with the idea that there are faculties of the mind, and that the mind in that sense has a modular structure. I have no quarrel with the notion of deep structures, or mental representations in the only sense in which these are really relevant to the exciting empirical inquiries that Chomsky and his associates are engaged in. I have no quarrel with the idea that in using language we display tacit knowledge, operating rules and principles that cannot in the normal way be brought to conscious formulation. Finally, I have no objection to innate mental structures on the grounds of their innateness. Obviously, human beings are born with certain abilities, including abilities to mature as well as abilities to learn. Whether the ability to acquire grammars of a certain kind is an ability to learn or an ability to mature under certain conditions seems to me a philosophically open question, capable in principle of being settled by empirical inquiry.

Despite this range of agreement, I think that Chomsky employs in his writing a confused notion of the mental. I should perhaps begin by explaining what I think a non-confused notion of the mental looks like.

The mind is the capacity to acquire intellectual skills. The chief and most important intellectual skill is the mastery of language. Others, such as knowledge of mathematics, are acquired by human beings through the languages that they have mastered. So the study of the acquisition and exercise of language is the way *par excellence* to study the nature of the human mind.

Someone who has acquired a language *knows* that language. Knowledge of a language is an ability: an ability that can be exercised in many different ways, for instance, by speaking the language, by understanding what is said to one in the language, by reading the language, by talking to oneself in one's head in the language. To know a language just is to have the ability to do these and similar things. It is a conceptual truth that to study an ability you have to study the exercise of that ability: to investigate what the ability to ϕ is you have to investigate what ϕ ing is. So to study knowledge of language you have to consider and examine what the exercise of linguistic knowledge is. The exercise of linguistic knowledge can be called, if you like, linguistic behaviour. But 'behaviour' must be understood in a broad sense, so that, for instance, reciting a poem to myself in my head imperceptibly to others will count as an instance of my linguistic behaviour.

We must distinguish between abilities and their possessors on the one hand and their vehicles on the other. The possessor of an ability is what *has* the ability. I am the possessor of my linguistic ability; it is I (and not my mind, or my brain) who know English and am exercising this capacity in giving this lecture. Similarly, my car has the capacity to decelerate: it can go slower in answer to my touch on the foot-brake. The *vehicle* of the car's ability to decelerate is the brake mechanism; similarly, my eye is a part of the vehicle of my ability to see. The vehicle of an ability is that part of its possessor in virtue of which it is able to exercise the ability. A vehicle is something concrete and more or less tangible; an ability, on the other hand, has neither length nor breadth nor location. It is, if you like, an abstraction from behaviour.

An important instance of the distinction between possessor, ability, and vehicle is the distinction between people, their minds, and their brains. Human beings are living bodies of a certain kind, who have various abilities. The human mind is the capacity that human beings have to acquire intellectual abilities: a capacity is itself an ability, but a second-order ability, the ability to acquire abilities. The vehicle of the human mind is, very likely, the human brain. Human beings and their brains are physical objects; their minds are not, because they are capacities. This does not mean that they are spirits. A round peg's ability to fit into a round hole is not a physical object like the round peg itself, but no one will suggest it is a spirit. It is not any adherence to spiritualism, but simply concern for conceptual clarity, that makes us insist that a mind is not a physical object and does not have a length and breadth.

If a mind is not a physical object, can it have a structure at all? Yes, it can. The set of abilities through which the mental capacity is exercised have relationships to each other – there are relationships, for instance, between the ability to multiply and the ability to take square roots – and these relationships between abilities form the structure of the mind. Not only human beings have abilities that are structured in this way: we can discover the structure latent in the operations of a pocket calculator by identifying the algorithms that it uses. To discover the algorithm that a calculator uses, say, for the extraction of square roots calls for mathematical, rather than electronic, inquiry. When considering the human mind, the physiologist is in the position analogous to the electronic engineer, and the psychologist is in the position of a mathematician who is trying to deduce from the form the calculator's output takes (what kind of rounding errors it commits, etc.) what is the algorithm that it is using.

Chomsky makes a distinction between capacities and their vehicles, as I have done. He describes the object of his study as ‘human cognitive capacities and the mental structures that serve as the vehicles for the exercise of these capacities’. But in terms of the distinctions that I have drawn the mental structures that Chomsky is interested in are capacities, not the vehicles of capacities; it is the physiological hardware characteristic of the exercise of the relevant mental capacities that is the vehicle.

In fact Chomsky’s mental structures seem to belong at times to the world of software, at times to the world of hardware. The great majority of what he says, as a linguist, about the knowledge and use of grammar by language users, is perfectly intelligible in terms of mental structures’ being particular abilities and their exercises: the exercise of the ability to operate an algorithm, to discover the value of a particular grammatical function for a given grammatical argument. But from time to time Chomsky the philosopher intervenes to tell us that what he is talking about is not to be understood as a capacity or ability at all. It is something quite different, which underlies the ability and its exercise, and which yet does not underlie it in the way that the physiological structures and processes of the brain do.

To show that it is possible to know a language without having the capacity to use it, Chomsky offers (p. 51) the following argument:

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?

The answer, Chomsky says, is ‘yes’: that is shown by the fact of recovery.

I do not wish to contest the answer that Chomsky gives to his question: what does seem to me surprising is his clear assumption that there is a fact of the matter here, to be settled by considering empirical evidence. If it really were a factual matter whether a person in the condition described by Chomsky knew English or not, then a thought experiment would be a most inappropriate way to settle the question. A thought experiment is not an experiment and does not provide empirical evidence: its function in philosophy is rather to draw attention to the shape and structure of our concepts. This Chomsky’s illustration effectively does: it shows the fuzzy edges of the concept *knowing English*. In the normal case, a large number of criteria for the application of the concept are present: the person in question can readily speak, understand, and think in English. We imagine the criteria whittled away, so that all that is left is that the person is going later to use English normally. Shall we say he knows English in the interval? 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 the 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.

Of course, Chomsky denies the presence of the capacity during the intervening period. He says:

In the intervening period, he had no capacity to speak or understand English, even in thought.

But why should he say that? It is possible that he thought that it followed from something else that is, on the hypothesis, genuinely true:

He had no capacity to speak or understand English, even in thought, in the intervening period.

But of course the first proposition does not follow from the second, and if Chomsky thinks that it does he is mistaken.

Chomsky goes on to consider a second aphasic who is like the first but who never recovers speech. To deny that his person lacks knowledge of English would be perverse, he argues: we have agreed that the first aphasic knew English, and this second one is in exactly the same mental (ultimately physical) state, as might be shown on autopsy. Well, once again this would be a matter for decision, not for discovery; but surely this time the natural decision would go the other way. The one remaining prop to support the application of the concept *knowing English* – that the unfortunate was later going to use the language again without relearning – has been pulled away. So what grounds are left for saying that he knows English? ‘But he is in the same brain state as somebody who does know English’. But that is to beg precisely the question at issue: that it is the brain state, and not the presence or absence of the capacity, that settles whether someone knows English or not.

The question-begging nature of Chomsky’s procedure is masked by his use of expressions such as ‘mental (ultimately physical) structures’ and ‘mental (ultimately physical) states’. Chomsky uses such expressions to indicate that his mentalism does not involve any sort of immaterialism: mental structures are simply physical structures described at a certain level of abstraction. But the expressions are ill-chosen, whatever one may think of immaterialism, because they conceal the fact that the criteria of identity for a mental state are not the same as those for a physical state. Two people can be in the same mental state while being in a different physical state, and can be in the same physical state while being in a different mental state. To say this does not beg any question about materialism, since it is equally true of computers that there is no one–one correlation between software structures and hardware structures.

An analogy may help here. A monarch is a legal person, in the sense that what it is to be a monarch is defined by a set of legal relationships. But all actual monarchs, from Hammurabi to Elizabeth II, are physical persons, human beings of flesh and blood. We might therefore, in the style of Chomsky, call monarchs ‘legal (ultimately physical) persons’. But if we do so, we invite confusion. Suppose that I met the King of England in 1937 and in 1940 I say ‘I have just met again the same legal (ultimately physical) person as I met in 1937’. Whom did I meet in 1940: the Duke of Windsor or King George VI? The confusion in talking of the same mental (ultimately physical) structure is less obvious but no less serious.

Indeed, the confusion is more serious. For monarchs are indeed persons; whereas states of the mind and states of the brain are not states of the same kind of thing. I agree with Chomsky that to describe a state of mind is to describe, at a certain degree of abstraction, a physical object; but the physical object so described is a human being and not a brain. The brain states

characteristic of speakers of English – if we go along with Chomsky in assuming that there are such – may, for all we know, be reproducible in a brain *in vitro*. However successfully they were reproduced they would not constitute knowledge of English; for it is people, not brains, to whom it makes sense to attribute such knowledge.

Let me emphasize again that I am not arguing for immaterialism or spiritualism. The conceptual points that I have been making can be made about pocket calculators no less than about human beings. My calculator works out the square root of 123456789. In a flash, there comes the answer 11111.11106. Between my pressing the square-root key and the display appearing, complicated events took place in its electronic innards. Those events, whatever they were, could have taken place in a different calculator doing a different job; and a different calculator doing the same job might well have taken an electronically totally different route. Moreover, the hardware might have been taken out of the case, separated from the input keys and the output display. Whatever electronic events then took place inside it would not have been the working out of the square root of 123456789. For in the sense in which calculators can work out square roots for us, it is only whole calculators, and not portions of their electronic anatomy, however sophisticated, that can do the working out.

Chomsky's final argument against the identification of knowledge of English with the capacity to use it goes like this (p. 52):

Were we to identify capacity and knowledge, we would presumably be led to say that the aphasic does not know English when the capacity is lacking, and hence would be committed to the belief that full knowledge of English can arise in a mind totally lacking this knowledge without any relevant experience whatsoever, as the case of recovery shows, something that is plainly not true of the child's mind and seems an exotic claim.

Such a claim would indeed be exotic; which makes it more surprising that only 40 pages later, on p. 93, Chomsky should commend Arthur Danto for pointing out the possibility 'that in principle there might be a "Spanish pill" with the property that by taking it we should have been caused (adventitiously) to be masters of Spanish without having learned the language'. Such a master of Spanish would undoubtedly know Spanish, Chomsky says, hence

We cannot rule out in principle the possibility that taking a pill might bring about the mental state that constitutes knowledge of Spanish.

In such a case, of course, knowledge of Spanish would arise in a mind totally lacking the knowledge without any relevant experience whatsoever: the claim that Chomsky dismissed as exotic in the earlier passage. Indeed it is a more exotic claim, for in the case of the recovered aphasic it might well be claimed that there *was* relevant experience: the experience at the time when the language was originally acquired in the normal way.

I think that Chomsky is right not to reject as inconceivable the notion that a pill might give us mastery of Spanish: the inconsistency between this concession and his earlier position merely brings out further the futility of his attempt to separate knowledge of English from the ability to use – the mastery of – the language. The idea is indeed exotic, but it is not literally inconceivable. What would be inconceivable would be the idea that a pill might give one knowledge of Spanish *without giving one the capacity to use the language*.

The root of Chomsky's confusion is his failure to distinguish between two different kinds of evidence that we may have for the obtaining of states of affairs: to distinguish between *criteria*

and *symptoms* (to use the terminology introduced by Wittgenstein). Where the connection between a certain kind of evidence and the conclusion drawn from it is a matter of empirical discovery, through theory and induction, the evidence may be called a *symptom* of the state of affairs; where the relation between evidence and conclusion is not something discovered by empirical investigation, but is something that must be grasped by anyone who possesses the concept of the state of affairs in question, then the evidence is not a mere symptom, but is a *criterion* of the event in question. A red sky at night may be a symptom of good weather the following morning; but the absence of clouds, the shining of the sun, etc., tomorrow are not just symptoms but criteria for the good weather. Similarly, the occurrence of certain electrical brain patterns may be, or may someday come to be, symptoms of the presence of knowledge of English in the person whose brain is in question. But his ready use of English is not just a symptom of, it is a criterion of, a knowledge of English.

To grasp the importance of the distinction between criteria in symptoms in connection with knowledge of language, consider the following case. Suppose that Professor Chomsky were now to die, and on opening his skull we discovered that there was nothing inside it except sawdust. This is indeed an exotic suggestion: if it happened it would be an astonishing miracle. But if it happened it would not cast the slightest doubt on what we all now know, namely, that Chomsky knows English extremely well. But on Chomsky's view we would have to say that in fact it showed he never knew English at all, since on his view knowledge of English is ultimately a certain physical structure. But the supposition that someone can use English as Chomsky does and yet not know English is not just a miraculous supposition: it is a piece of literal nonsense.

Chomsky's characterization of the mental structures that interest him makes them straddle uneasily the distinction between hardware and software. They seem to be too ghostly to be hardware (from time to time he reminds us that it is no part of his theory that they should be in the brain rather than in the liver); but they also seem to be too concrete to be software, otherwise they could not be characterized as ultimately physical structures. But though his characterization of mental structures is, I have argued, confused and incoherent, what he is studying when he is studying mental structures is something genuine, important and fascinating. It is, precisely, the relationship between different capacities and sets of capacities.

Chomsky would deny this. If we reject his conception of mental structures, he says (p. 50),

We are left with a descriptive study of behaviour, potential behavior, dispositions to behave and so on, a study that in my opinion cannot, in the end, be pursued in a coherent way . . .

What in fact we are left with, and what Chomsky in fact studies, are relations between different abilities: in particular the ability to render the value of certain functions (e.g. linguistic transformations) for given arguments. It is abilities of this kind that Chomsky is studying when he aims to uncover the structures underlying our use of language.

There are many intellectual tasks that we can perform in more than one way. When we multiply a number by ten, for instance, we can (*a*) simply add a 0 (*b*) recite the appropriate part of the ten times table (*c*) write the number down ten times and add up the result. Now, wherever we have a case of doing A by doing B there will be questions to be raised about the relations between the ability to do A and the ability to do B; we can only multiply by ten in the first way, for instance, if we know the decimal notation, and in the second way if we know the ten times table. And when we do A by doing B, it may well be that we know very well

that we are doing A, but do not know without reflection that we are doing B; when we return a serve at tennis there are many movements of hand and arm that we are not normally aware of by which we make the return. Similarly, in performing intellectual tasks – including the comparatively modest ones of pronouncing a word or constructing a sentence – there are many sub-tasks that we perform without conscious advertence. When we ask what rules or principles we employ in performing these tasks, we are asking what sub-abilities we are exercising when we exercise the ability to use language.

In what sense, then, does the performance of such tasks have ‘psychological reality’ when it is not conscious? The ability being exercised is a *psychological* reality in the sense that it is an ability that is being exercised in a task that is an intellectual one and not a merely physical one. It is a psychological *reality* in that it is open to testing which of the various possible algorithms that I might use in performing the task is the one that I am actually using. It is here that it is relevant to study the reaction time of subjects and similar phenomena to which Chomsky’s attitude has become increasingly cavalier.

I entirely agree with Chomsky in regarding the consciousness of a mental process as being quite inessential to the reality of the process. But this does not mean that psychological reality is unimportant. Once again, the point can be made in terms of simple non-human machines. It is a genuine empirical question which particular algorithm is being used by a computer or calculator to produce the solution to a problem, and the speed of computation and the nature of rounding errors and the like might provide empirical evidence for or against the use of a particular algorithm. It would in no way count against this that there was no representation of the algorithm in the output display of the calculator, or in the monitoring file provided by the computer. It is the monitoring file that is the analogue, in the computer context, of the conscious account that we can give of the way in which we perform intellectual tasks.

The philosophical confusions that I have claimed to detect in Chomsky’s presentation of his theories of knowledge of grammar are in no way organic to the theories themselves. The theories, if I have understood them, can be stated in such a way as to be quite free of the particular form of mentalism in which they have been wrapped up and which has provided irrelevant distraction to philosophical and psychological critics.

REFERENCE (Kenny)

Chomsky, N. 1980 *Rules and representations*. New York: Columbia University Press.