THE BRAIN AS HARDWARE, CULTURE AS SOFTWARE

1. Neo-Carnapians vs. Neo-Wittgensteinians: Chomsky vs. Davidson

For the last fifty years or so there has been a struggle going on between the heirs of Carnap's "unified science" movement and philosophers who take their inspiration from the later Wittgenstein. The differences between them are crystallized in their attitude toward cognitive science.

The Carnapians see the failure of the sort of behavioristic reductionism that Carnap himself inauspiciously attempted as leaving us with a problem about the place of intentionality in a world of physical particles. Many of them regard that problem as having now been solved, thanks to "the computational revolution". That revolution, Jerry Fodor tells us, was made possible because "computers showed us how connect semantics with causal properties for symbols" (Psychosemantics, p. 18). The analogy between brains and computers let us realize that symbols could, so to speak, become incarnate as neural states and thereby exert causal power. Fodor and the other heirs of Carnap see this revolution as a great intellectual breakthrough, giving us an insight into how the mind works that was not available previously. They think that collaboration between philosophers and cognitive scientists will fulfill Carnap's hopes for a unified science.

Philosophers of mind and language who are skeptical about such collaboration, such as Davidson, Brandom and Descombes, are not troubled by the irreducibilty of the intentional. They They do not think the irreducibility of one vocabulary to another presents any more problems than the unsubstitutability of one tool for another. Nor do they think of the mind as a mechanism whose workings cognitive scientists can study. They agree that the brain is such a mechanism, but they do not think it helpful to identify the mind with the brain.

They urge that we drop what Ryle called "Descartes' para-mechanical hypothesis", rather than reinterpreting it in physicalist terms by viewing brain states as representations. Davidson has said that we should not think of linguistic competence as possession of a "portable interpreting machine" that is "set to grind out the meaning of

an arbitrary utterance". ("Nice…", p. 445). He and the other heirs of Wittgenstein would like to get rid of the notion of internal representations, and thus of mechanisms that produce such representations. Whereas the neo-Carnapians regard mind and language as things that can be understood by taking them apart and seeing how their parts gear in with one another, the neo-Wittgensteinians see them as social skills. Thus Davidson says that, although a theory of meaning must describe what an interpreter can do, there is not reason to think that "some mechanism in the interpreter must correspond to the theory".

Chomsky's comment on this claim is that "for anyone approaching these problems from the standpoint of the natural sciences" Davidson's way of thinking is "utterly wrongheaded". (Chomsky, New Horizons, p. 56). He backs up this criticism by saying that a similar skepticism about underlying mechanisms would have inhibited the development of chemistry. Just as Dalton and Mendelejev enabled us to see mechanical interactions between atoms behind the ability of elements to combine into compounds, so cognitive science will enable us to see mechanical interactions behind the ability of humans to acquire social skills. For Chomsky, philosophers like Quine, Davidson and Dummett are Canute figures, trying to hold back a rising tide of established empirical results.

These results include, he claims, the discovery that "an initial state of the language faculty incorporates certain general principles of language structure, including phonetic and semantic principles". (NH, p. 60) Whereas Ryle thought that it was only "Descartes' myth" that the study of what he called "qualities of intellect and character" could be a study of mechanisms, Chomsky are as sure that Descartes' and Locke's analogies between physical atoms and mental atoms put us on the right track as the Wittgensteinians are that the notion of mental atoms is unhelpful and misleading. Ever since he published *Cartesian linguistics*, some twenty-five years ago, he has made clear his distaste for Wittgensteinian anti-Cartesianism. His later work has deepened and widened the division between those who think that the seventeenth-century "idea idea" has outlived its usefulness and those who do not.

Chomsky is caustic about Davidson's claim that Quine "saved philosophy of language as a serious subject" by getting rid of the

analytic-synthetic distinction. That distinction, Chomsky says, is taken for granted by "virtually everyone who works in descriptive semantics". As he writes:

> **One would be hard put to find studies of language that do not assign structures and describe the meaning of *kill*, *so*, etc., in such a way that there is a qualitative distinction—determined by the language itself—between the sentences "John killed Bill, so Bill is dead" and John killed Bill, so John is dead". ("New…", p.**

Chomsky says that we need the distinction between what is "determined by the language itself" and what is not in order to explain such phenomena as that "each child knows the relevant difference between 'who did John see Bill with?" and "who did John see Bill and?" Since, as he says, "children do not…produce 'who did John see Bill and?', then to be informed by their parents that this is not the way it is done", the only explanation available is the innate structure of the language faculty. ("New…", p. 56)

Chomsky's argument here depends on the assumption that any behavior that cannot be accounted for by interaction between the child and its parents must be accountable for by reference to an innate faculty, and on the further assumption that the absence of certain behavior is as good an explanandum as its presence. But this second assumption is, Wittgensteinians would rejoin, highly counter-intuitive. Suppose we asked for an explanation of why no child continues the sequence "2, 4, 6, 8", after reaching triple digits, with "104, 108, 112", and of why no correction or instruction by parents is necessary to insure that the child stays on track. The Chomskian explanation would presumably be that an innate mechanism is at work. To Wittgensteinians, this a "dormitive power", explanation of a non-event, one that nobody had previously thought to be in need of explanation.

The standoff between Davidson and Chomsky over these matters is clearest when Chomsky criticizes Davidson's attempt to ""erase the boundary between knowing a language and knowing our way around the world generally." This erasure would give empirical inquiry into language no foothold, for it would make a theory of language-learning what Chomsky calls "a theory of everything". "The proper conclusion", he continues,

> **Is not that we must abandon concepts of language that can be productively studied [such as Chomsky's own concept of "a internal representation of a generative procedure" (NH, p. 69)] but that the topic of successful communication in the actual world of experience is far too complex and obscure to merit attention in empirical inquiry (NH, 70)**

Chomsky thinks that the sort of common-sense explanation of how children learn language with which Davidson is content with will not do for scientific purposes, because "Reference to 'misuse of language', to 'norms,' to 'communities,' and so on…require[s] much more care than is taken. These concepts are obscure, and it is not clear that they are of any use for inquiry into language and human behavior"(NH, 72).

In passages such as this, Chomsky shows himself the true heir of Carnap. Carnap too would have found the notion of "norm" unclear, and unsuitable for purposes of scientific inquiry. But for Wittgensteinians to make a concept clear is simply to become familiar with the use of a linguistic expression. For them, the use of "norm" is less problematic than the use of "internal representation of a generative procedure".

From a Wittgensteinian perspective, the approach taken by Chomsky and his fellow cognitive scientists looks like that taken by the man who searches for his missing keys under the lamp-post, not because he droppedthem near there but because the light is better. What Chomsky calls "the standpoint of the natural sciences" is simply the habit of looking for micro-mechanisms behind macroscropic behavior. The idea that adopting this standpoint will always pay off looks to Wittgensteinians like Carnapian dogmatism.

Chomsky would presumably counter that linguists and other cognitive scientists have produced a body of knowledge that is much more than mere verbal manipulation. One item in this body of knowledge, for example, is that there is "a fixed biologically-determined function that maps evidence available into acquired knowledge, uniformly for all languages". (NH, p. 53) But it is only politeness, I think, that prevents Wittgensteinians like Davidson from rejoining that psychologists have been producing what they claimed to be bodies of

knowledge for a century or so without these putative contributions having had any significant impact either on debates among philosophers or on human beings' self-images.

The study of human cognition by psychologists in the period prior to the emergence of what Steven Pinker calls "the computational theory of mind"—roughly, the hundred years between 1860 and 1960—was accompanied by a lot of drum-beating and a lot of proclamations that a priori philosophizing was about to be superseded by a triumphant empirical discipline. Think of Watson, Piaget and of Skinner. But that period produced no new instruments of reliable prediction and no technological advances—the sort of things we think we expect scientists to produce.

Exasperation with the ratio of drum-beating to the production of useful instruments of prediction and control was one of the reasons why Wittgenstein ended *Philosophical investigations* and Ryle ended *The concept of mind* with doubts about psychology as a discipline. Wittgenstein said that the idea that "psychology treats of processes in the psychical sphere, as does physics in the physical" was a "misleading parallel", (PI, sec. 571) and that "the confusion and barrenness of psychology is not to be explained by calling it a 'young science'" (PI, p. 232) Contemporary Wittgensteinians have similar sentiments.

It is one thing to say that Chomskian linguistics, and the other academic specialities that bill themselves as parts of "cognitive science", are respectable disciplines—arenas in which very bright people engage in spirited debates with one another. It is another thing to say that these disciplines have contributed to our knowledge. Many such disciplines have flourished and decayed without leaving such contributions behind them. Fifteenth century Aristotelianism, seventeenth century hermeticism, and twentieth century logical empiricism are familiar examples.

The Wittgensteinians think that it is an open question whether cognitive science will go down in history as a successful attempt to bring the procedures of natural science to bear on the study of mind and language or as yet another attempt to set philosophy on the secure path of a science—one that eventually collapsed, like all the others, of its own weight. They have doubts about whether cognitive science will ever be

able to validate itself, and extricate itself from philosophy, in the way that chemistry did—by spinning off a new technology. Whereas the fans of cognitive science the Wittgensteinians as dogmatic behaviorists, the Wittgensteinians criticize the Chomskians in the same terms as Bacon criticized late scholasticism. They think of Chomsky and Fodor as he thought of Occam and Scotus: their beautiful theories and subtle arguments cannot be brought to bear on practice. They are building mechanisms in the air.

2. Compositionality: Fodor vs. Brandom

Fodor has attempted to break through the impasse about the utility of cognitive science by arguing that the Wittgensteinians are trying to fight the plain facts. If they are to do more than express quasi-behaviorist prejudices, he says, they must come up with some theory of their own about the way in which language works. All they can produce, he suspects, is "semantic holism", the doctrine that "the meaning of an expression is constituted by all of its inferential relations, hence by all of its role in a language". (MR, 153)

That doctrine is put forward implicitly by Davidson and explicitly by Brandom. Since the study of the roles of expressions in languages is the study of what Chomsky calls "successful communication in the actual world of experience", and since holists cannot easily distinguish between knowing a language and knowing one's way about in the world generally, this study cannot avoid becoming what he dismissively calls a "theory of everything". So semantics cannot be the sort of discipline that sustains an analogy with chemistry, and perhaps cannot be a discipline at all. That is why Davidson and Brandom offer no research programs for eager young cognitive scientists to carry out, and why they seem to Chomsky like obstructionist Luddites.

Fodor thinks that there is a decisive reply to the semantic holists, one that leaves the Chomskians in possession of the field. It is that language is compositional—that is, "the meaning of syntactically complex expressions is a function of their syntactic structure together with the meaning of their syntactic constituents" (MR, 146) . Since "meanings are compositional but inferential roles are not compositional, meanings can't be inferential roles". (MR, p. 147) Fodor spells out the point as follows:

> **The meaning of the phrase 'brown cow'…depends on the meanings of 'brown' and 'cow' together with its syntax….But now, prima facie, the inferential role of brown cow depends not only on the inferential role of 'brown' and the inferential role of 'cow' but also on *what you happen to believe about brown cows*. So, unlike meaning, inferential role is, in the general case, not compositional. (MR, 147-8)**

At this point in the argument, Wittgensteinians are inclined to say that if there were such things as meanings, languages would indeed be compositional, but that Quine and Davidson have given us good reason to think that there are not. There are just organisms that influence each others' behavior by making noises and marks rather than hitting each other with sticks. There is no fine structure—consisting of relations between mental or linguistic atoms called "meanings" or "concepts" of "representations"—to be discovered beneath their possession of these social skill. There are no mechanisms underlying their abilitites. To think that there is an feature called "sameness of meaning" to be detected within the conversational flux is (as Harman once put it, in the course of defending Quinean holism) to overestate the difference between dictionaries and encyclopedias. It is to insist on a distinction between language and fact that serves no purpose except to keep cognitive scientists employed.

Fodor's response is to describe this Wittgensteinian line of argument as

> **a well-greased, and well-traveled slippery slope: having arrived at the bottom, one finds oneself accepting such prima facie outlandish doctrines as that no two people ever share a belief, that there is no such relation as translation; that no two people ever mean the same thing by what they say; that no two time slices of the same person ever mean the same thing by what they say; that no one can ever change his mind; that no statements, or beliefs, can ever be contradicted (to say nothing of refuted), and so forth. It is a moot question how to get the goodness out of inferential role semantics without paying this extravagant price. (MR, 143)**

At this stage of the argument, Wittgensteinians typically make the Rylean point that two people can have the same thing—for example, the same build or the same outlook--without there being something that they both have that will repay study. When we say that a sentence in French and a sentence in English have the same meaning, or that a Frenchwoman and an Englishman share the same belief, we do not need criteria of sameness and difference of meaning or belief to back up our claims. Fodor, the Wittgensteinians think, is trying for more precision than the subject-matter admits, and more than anyone except cognitive scientists need.

Ryle's point was that there is an everyday sense in which many people have the same build and the same outlook, and another, equally useful, sense in which no two people do. Similarly, the utility of the everyday sense of "means the same" and "believes the same" should not make us reject the point that Descombes approvingly quotes from Sartre: "whenever I form a sentence its meaning escapes me, is stolen from me; meanings are changed for everyone by each speaker and each day; the meanings of the very words in my mouth is changed by others". (MP, p. 247) Sartre's point is the same as Brandom's when he says that "A word—'dog', 'stupid', 'Republican'—has a different significance in my mouth than it does in yours, because and insofar as what follows from its being applicable, its consequences of application, differ for me, in virtue of my different collateral beliefs". (ME, p. 587)

Fodor inisists that the Achilles heel of the Wittgensteinians is the productivity of natural languages, a property that can only be explained by compositionality. He defines productivity, roughly, as the ability to express an open-ended set of propositions—well-formed strings of potentially infinite length. But the obvious Wittgensteinian reply to this move is to say that if we think of languages as generative mechanisms—as brains equipped with the neurological equivalents of grammars and dictionaries—then they will indeed have this feature. But if we follow Davidson in denying the need for "portable interpreting machines set to grind out the meaning of arbitrary utterances" we shall not think of them in this way. If we think of knowing a language as having a social skill, the facts about infinite productivity that Fodor cites will seem irrelevant.

"Productivity", Fodor says, "is the property that a system has when it has an infinite number of syntactally and semantically distinct symbols." (CP, 1) But from a Wittgensteinian perspective that is a good reason to deny that what speakers of a natural language have is such a system. If one abstracts out an abstract entity called "English" from the conversational interchanges of various organisms, one can say, as Fodor does, that "English contains the open-ended sequence of nonsynonomous expressions: 'missile shield', anti-missile shield', anti-anti-missile-shield-shield'…and so forth". But that is like saying that an abstract entity called "arithmetic" contains such an open-ended sequence. Abstract entities can do things that organisms cannot; in particular, they contain infinite numbers of things. To say that we incorporate such an entity is already to beg the question against Davidson's claim that skills need not be backed up by mechanisms.

3. Determinate being

Another way that Fodor puts the issue between the Carnapians and the Wittgensteinians is to contrast "meaning realists"—people who think that "there are facts about meaning, and that they are suitable objects of scientific inquiry"—with those who think that "meaning arises from our practices of interpretation, so that there no more needs to be a single right answer to 'what does 'pinochle' mean?' than there is to 'what does *Hamlet* mean?'" Wittgensteinians, he says, think that "trying for a science of meaning would be silly; like trying for a science of games. Or of Tuesdays." ("A science of Tuesdays" in LRB, 20 July 2000, p. 22)

Fodor deplores what he calls "sympathy with this Wittgenstein-Goodman-Kuhn-Derrida sort of picture", and prays to God that "no miasmal mist from Harvard has seeped up [sic] the Charles to MIT". ("It"s all in the mind", TLS June 23, 2000). He associates this mist with what he calls "linguistic idealism" which he says is popular among philosophers "like Rorty, Putnam, Kuhn and Derrida" (CP, 22-23) That sort of philosopher thinks that meaning realism is as pointless as, recurring to Ryle's examples, build realism or outlook realism.

Quine initiated the Harvard-based assault on meaning realism by claiming that beliefs and meanings could never be fitted into a physicalistic world view, precisely because the indeterminacy of

translation and of intentional ascription meant that was no way to tell whether meant the same thing by the same words or shared the same belief. Starting from the claim that there is no entity without identity, Quine concluded that beliefs and meanings had no place in an attempt to "limn the true and ultimate structure of reality".

Fodor argues that since beliefs and meanings *need* to be given such a place, and since Quine was right that there can be no entity without identity, we must be meaning realists. Brandom, by contrast, sees no need either to insist that there is a single right answer to the question "what does 'pinochle' mean?" or to follow Quine in inferring from the lack of such an answer to the claim that beliefs and meanings are somehow not an ontological par with electrons and neurons. Davidson's account of linguistic competence in "A nice derangement of epitaphs" is of a piece with Brandom's view that knowing the content of an assertion is a matter of determining its place in the game of giving for asking for reasons. For both philosophers, social skills of the sort required to have intelligent conversations do not require the application of criteria for sameness of meaning or of belief.

The three-cornered dialectic between Fodor, Quine and the Wittgensteinians can be summarized by saying that Fodor and Quine agree that only beings that exist are those for which there are context-free criteria of identity, whereas Davidson and Brandom do not. But a more illuminating way of describing this standoff is provided by the terminology used by Cornelius Castoriadis and, more recently, by Vincent Descombes. Descombes says that Castoriadis criticized

> **An intractable prejudice of philosophers and of all those whom they (often unknowingly) inspire: that everything that exists exists in a determinate form. Everything that exists is precise, determined, and apprehensible. If by chance something exhibits indetermination, laziness, or vagueness, then that thing has shown itself to be, if not utterly illusory, at least of inferior status.(MP, 240)**

Descombes expands on this point by saying that whereas representations can be more or less determinate, "Reality is neither determinate nor indeterminate". (p. 241)

For Descombes, the trouble with Fodor, Quine and others who think that entity requires identity is that they succumb to what he calls "the illusion of a general metaphysics". Such a metaphysics, he says

> **would have to take up unity and plurality, identity and difference, the individual and relations, all without taking any particular domain into consideration. In particular, it would have to make clear how the words "being" and "identity" are to be understood before philosophical inquiry is divided into 'regional ontologies", including the ontology of nature and the ontology of mind. …[But] how can one inquire into conditions of identity without taking into account the type of things one seeks to identify? (p. 241)**

Here Descombes makes a point that is also made by Brandom. Descombes says that "the concept 'thing' is not meant to be applied in an exercise of enumeration". Brandom says that

> **establishing a criterion of identity [and so a sortal] is not only sufficient for countability; it is necessary as well. Unsortalized 'things' or 'objects' cannot be counted. There is no answer to the question how many things there are in this room; there is one number of books, another of molecules, another of atoms, another of subatomic particles….Counting is intelligible only by reference to a sortal concept." (ME, p. 438)**

But the fact that only what has been sorted out can be counted does not entail that all sortals pick out sets of countable items. As Descombes points out, the fact that one does not begin to know how to answer the question "How many representations are there in Gericault's painting *The raft of the Medusa*?" is not because representations are fuzzier than other things, but because unsortalized representations are as bad as unsortalized things. So to answer the 'how man?' question, Descombes says, "one would have to be able to enumerate the things represented in the painting". The same problem, he points out, arises for a brain.

Suppose that someone who, like Fodor, believes that the brain contains representational states wants to identify the representations present in the visual cortex of a brain attached to an eye that is focused

on an approaching predator. The answer he gets will depend upon whether he treats the brain as representing colors, or shades of colors, or patterns of colors, or middle-sized physical objects, or light waves, or environmental dangers, or chemical changes in the retina. To say that the brain is a computational device is not yet to decide which of many input-output functions the device is programmed to run. There are as many such functions as there are alternative descriptions of its environment and of its behavior. To be a meaning realist is to hold that one such function is the right one. But it is dubious that empirical evidence can decide between alternative functions.

Seen from the Descombes-Brandom angle, the Quinean idea that there are not really such things as meanings, or that meanings have a low ontological status, because an indefinitely large number of different translation manuals would do equal justice to the evidence is like saying that there is not really such a thing as visual perception because an indefinitely large number of input-output functions would do equal justice to the observed relations between the visual cortex and its environment. The impulse to infer from the existence of alternative descriptions to ontological second-rateness is found only in those who are subject to what Descrombes calls "the illusion of a general metaphysics".

So the Wittgensteinian response to the Quine's doctrine of the indeterminacy of translation is to say that there is nothing special about language, or the intentional. As Chomsky was, ironically enough, the first to argue, that indeterminacy is not distinct from the ordinary under-determination of theory by evidence. Such under-determination is no more puzzling than the familiar fact that if you have found it useful to have two distinct vocabularies available for describing the same chunk of space-time it is unlikely that any given sentence formulated in the one vocabulary in the one will entail or be entailed by any sentence formulable in the other. For if such entailments did exist, the vocabularies would long since have collapsed into one another.

Philosophers who, like myself, get their kicks from inhaling the Wittgenstein-Goodman-Kuhn-Derrida miasma claim to know quite well why "rabbit" translates "Gavagai" better than "rabbit-stage". We are confident that it has nothing to do with inspecting two entities called "meanings" and ticking off their resemblances and dissimilarities. It has

everything to do with relative ease of acquisition of the social skills promoted by bilingualism. Analogously, we can see the utility of sometimes treating the organism as responding to visually presented middle-sized physical objects, sometimes to aesthetic values, and sometimes to light waves. without worrying about what it is "really" responding to. The latter question seems to us as bad as the one about what *Hamlet* is *really* about, or of whether the "real" subject of *Winnie the Pooh*, is child abuse, or the self-consuming character of literary artifacts, or the collapse of late capitalism.

If we interpret "what is really going on" as "what it is useful to describe what as going on as in order to achieve our purposes," then Fodor's claim that we need to attribute semantic properties to brain states can be rephrased as "For the purposes of cognitive science, we need to decide whether to treat whether to treat the organism, or its central nervous system, as running this program rather than that--the one that takes these features of its environment as inputs rather than those, and produces output in the form of behavior described in this way rather than that". The multiplicity of input-output functions available for describing the same state of a computer parallels the multiplicity of contexts into which Frederick Crews puts the text of *Winnie the Pooh*. If cognitive scientists agree on which such function to select, it is presumably because they agree that the purposes of their discipline—the prediction and control of human and animal behavior—are best served by this choice.

This amounts to saying that the question of whether Wittgensteinians hesitancy to adopt what Chomsky calls "the standard methods of the natural sciences", and their doubt that it will be helpful to attribute semantic properties to brain states, boils down to doubt about whether doing so will do anything for purposes of prediction and control. The issue is neither ontological nor methodological but practical. You do not get to claim the prerogatives of a natural science simply by performing experiments and formulating hypotheses to explain their results. Even the alchemists could do that. To be recognized as such a science you have to give promise of what Bacon called "the improvement of man's estate".

Wittgensteinians doubt that postulating a "a fixed biologically-determined function that maps evidence available into acquired

knowledge, uniformly for all languages" will let us learn more languages faster, or that it will do much to explain why some people have so much trouble learning any foreign language. On the other hand, they are quite willing to predict that someday both pedagogy and diagnosis of therapy will be greatly improved by our ability to tweak neurons. For Wittgensteinians are as good physicalists as Carnapians. They are equally committed to the view that you cannot alter somebody's psychological state without altering—somewhere, somehow—her brain state. What they doubt is that there is a profitable level of inquiry in between folk psychology and neurology.

They may well be wrong about this. Cognitive science may well produce practical payoffs of a kind that were not afforded by psychology before the so-called "computational revolution"—a revolution that is supposed to have taken place about forty years ago. So what is puzzling is how philosophers managed to wind up taking sides in a straightforwardly empirical question about what lines of research will pay off. How, one might ask, did philosophy ever get into the act in the first place? Shouldn't they have learned, from cautionary tales about those who had a priori doubts about Dalton and Mendelejev, to just wait and see how things pan out?

4. Descombes on the location of the mind

I take Descombes' answer to this latter question to be: long before the computational revolution, philosophers were disagreeing about sort of thing mind is, and, more specifically, where it is. They have, Descombes says, split on whether the mind is

> **within or without. Within, according to the mentalist heirs of Descartes, Locke, Hume and Maine de Brian and among whom one can also place the phenomenologists and the cognitivists. Without, according to the philosophers of objective mind and the public use of signs, for example Peirce and Wittgenstein. (p. 2)**

This contrast between Descartes, Locke and Hume on the one hand and Hegel and Wittgenstein on the other highlights another link between Descombes and Brandom. Brandom views Hegel and Wittgenstein are precursors of his own inferentialist, "social practice", approach to mind and language, the approach that he opposes to the representationalism

of Locke and Hume. He and Descombes are in rough agreement about which philosophers were the good guys and which the bad.

The big difference between the two is that instead of Brandom's representationalist-inferentialist opposition, Descombes divides the mighty dead into sheep and goats with the help of an internalist-externalist one. He says that one of the two principal aims of *The mind's provisions* will be to argue for externalism—for the thesis that internalism, also known as "the classical philosophy of the subject" is wrong to hold that a thought is mental because inside the skull, whereas a book, being outside, is merely physical. A book, like a nation, counts as a bit of what Hegel called "objective Geist" and both are examples of what Peirce called "a sign".

The trouble with this way of putting the matter is that it leads Descombes to conclude, in his final chapter, that if we did get a scientific psychology, if cognitive science did get beyond the drum-beating stage and began to produce some useful results, it would nevertheless remain "unable to tell us about the mind, i.e., about thoughts". The reason for this inability, Descombes says, is that

> **mental vocabulary is deeply historical, which is another way of saying that there are historical conditions of meaning. A subject's words and thoughts have the meaning that they must be given in his world and cannot be disassociated from that world….The world that must be known in order to know what the subject thinks is not just a natural world….[I]t is a cultural world, one that contains institutions like the calendar, money, banks, and the game of chess.**

It seems to me better to avoid saying that cognitive science would be unable to tell us about "the mind". It would be simpler to say that it cannot tell us about culture. I think that the Wittgensteinian strategy should be to avoid the question "where is the mind located?" and instead to split "the mind" in two. It would be a good idea for them to stop talking about "the mind" and just say that two quite different things have been designated by that term: one is the abilities hard-wired into the anthropoid brain and the other is the set of social skills that we call "culture". "Mind" in the solipsistic, Cartesian, inside-the-brain sense is, indeed the brain—just as the cognitive scientists insist. "Mind"

in the sense of objective spirit, the sense in which (to use Descombes' example) books and paintings are mental entities, is obviously not the brain. It is the social world in which we find calendars and chessboards.

This Solomonic strategy of cutting the mind in two and giving a half to each claimant may seem too quick and dirty a way of dealing with the debates currently raging in philosophy of mind and language. But I shall try to illustrate its advantages by showing how nicely the distinction between the two halves of the mind fits with a distinction that has often been invoked by cognitive scientists: that between hardware and software.

5. Brain as hardware, culture as software

To see the parallelism with the software-hardware distinction, consider the following fiction: In the twenty-fourth century, everybody has a computer tucked in his pocket---a computer which is to today's models as they are to Babbage's calculating engine. Everybody carries around exactly the same model, an exact replica of the prototype that appeared two hundred years earlier, and that has been robotically duplicated in enormous quantities and at trivial cost ever since. If your computer goes wonky, you throw it away and pick up a new one exactly like it.

Because of the original manufacturer's monopolistic greed, if you try to open up your computer it self-destructs, and if you try to get into the robotic assembly line the robots kill you. Though software gets better and better every year, hardware remains completely unchanged. Nobody any longer has any idea how the computers work, of what goes on inside the black boxes. Some philosophers speculate that computers might work the same way brains do, but nobody really knows.

But suddenly it becomes clear that self-destruct mechanism has ceased to function on the most recently produced computers. So people start tearing their computers apart and doing reverse engineering. They discover that it is all done with 1's and 0's: both program instructions and input take the form of binary numbers. They reconstruct the various levels of programming language that were built into the machine's operating system. Now, they say, they know all about how computers work.

The analogy I want to draw is obvious. Culture is to software as the brain is to hardware. The brain has long been a black box, but with the help of nanotechnology we may soon be able to pick it apart, axon by axon, and set about projects of reverse engineering. This will probably produce new pedagogical techniques and new therapies. But it is a wide-upon question whether it will do anything for the prospect of a unified science. For all the reasons offered by Davidson, Arthur Collins, Lynn Baker, Helen Steward, and other philosophers, we are not going to be able to map *beliefs* onto neural states, though we might perhaps be do that for occurrent thoughts, as well as for surges of hunger or lust.

If there is nothing interesting to say about how changes in belief and changes in the meanings of words are related to neurological mechanisms, it is hard to see how study of what Chomsky calls "the brain/mind" can be expected to interact with, and be integrated into, culture. For what would be the parallel with the way in which post-Mendelejev chemistry interacted with pharmacology, textile manufacturing, the mining industry, and various other areas of technology? That breakthrough in chemistry gave us a sense that now, at last, we saw how microstructure and macrostructure fitted together, It was the exaltation produced by this experience that lay behind Carnap's "unity of science" movement. But the hardware-software and the brain-culture relations are not examples of the micro-macro relations. They are examples of the relation between a tool and its manifold uses.

To put it in a nutshell: we already knew that our brains could be programmed to do a lot of different things, just as the computer users in my fiction knew that their computers could. Discovering the micro-structural details of the brain might enable us to do different things with it than we had previously done, or it might not, just as further information about hardware design might or might not facilitate the production of improved software. But whether it does or not, it is hard to see that our self-image would change, or that we would begin to see the relation between the natural sciences and the human sciences in a new way.

It is sometimes suggested that the discovery of how the brain works will tie in with evolutionary biology so as to enable us to

understand better what our brain has been tailored to do, and thus to understand "human nature" better. But the analogy between biological evolution and software development suggests that what the brain was originally tailored to do may be irrelevant to what it is currently being used for. The fact that the first software breakthroughs were in the area of missile targeting, and of information search and retrieval, tells us nothing in particular about what present-day computers are good for. The fact that the first stages in the evolution of the mammalian brain were dictated by the needs of hunter-gatherers tells us equally little about what to do with the resulting product.

Steven Pinker tells us that "The mind is a system of organs of computation, designed by natural selection to solve the kinds of problems our ancestors faced in their foraging way of life". (p. 21) Descombes would presumably reply that that is true of the brain but not of the mind. I would urge, however, that Pinker's sentence is misleading even if we substitute "brain" for "mind" as its subject. It is misleading for the same reason that it would be misleading to say that my laptop is a designed to track missiles and comb data bases. Different hypotheses about the original designers were up to are not going to make any difference to what we do with the artifacts we inherited from them.

This pragmatic attitude is the one that Wittgensteinians strike when confronted with books like Steven Pinker's and E. O. Wilson's, which tell us that linguistics, cognitive psychology, and evolutionary biology are conspiring to change our self-image. These books tell us that what we had previously thought to be cultural will turn out to be biological. But this claim would sound convincing only if we come to think that evolution has created a brain that unfortunately cannot be programmed in a certain way—a brain so constructed that certain input-output functions cannot be realized by it. To acknowledge the hegemony of the biological would be to admit that some seemingly promising cultural initiative should not be attempted, because biology blocks it.

It is hard to imagine what the argument for such blockage would look like. For it would amount to saying: do not even try to modify our social practices in the proposed way, because we know in advance that it will not work. The experiment should not be tried, for its outcome is foredoomed. This would be analogous to saying "Stop trying to

transmute base metals into gold, for Dalton and Mendelejev have shown why this will never happen." But presumably the only way in which a cognitive scientist like Pinker or Wilson could back up such a claim would be to show, on physiological grounds, that a certain belief cannot be held or that a certain desire could not be had.

This would be hard, not only because of the arguments against the localizability of intentional states that I mentioned earlier, but because of the further argument that if a belief could really not be held by anybody, then nobody would ever have been able to propose that it be propagated. (In fact, however, socio-biologists rarely try to argue in this way. The customary rhetoric of socio-biology has been effectively parodied by Michael Williams: "This unfortunate trait became embedded in our genes in the course of countless millennia spent roaming the African savannah, but its effects can be mitigated by a simple change in the US tax code.")

6. Semantic holism, historicism and linguistic idealism

So much for my claim that it is as hard to make the brain relevant to culture as to make hardware relevant to software. I want now to conclude by taking up the connections between the historicism common to Descombes and Brandom and "linguistic idealism", a philosophical view that Fodor says lies at the bottom of the slippery slope down which semantic holists descend. Fodor rightly thinks of me as someone who holds this view. I should like to read both Brandom and Descombes as tumbling down the same slope along with me.

Toward the end of *The mind's provisions*, Descombes says that "Mental vocabulary is deeply historical, which is another way of saying that there are historical conditions of meaning" (p. 244). This parallels Brandom's view that the content of concepts becomes explicit in the course of history—in the Hegelian way in which bud becomes explicit in the flower and the flower in the fruit. Semantic holism tells us that meaning, or content, changes when inferential relations do. History tells us that, at least in the case of the most interesting concepts, these relations have changed a lot in recent times. Brandom thinks Hegel right in treating concepts—that is to say, uses of words, as quais-persons, capable of maturation.

Representationalists like Fodor and Chomsky, who see semantic holism as causing its proponents to lose touch with reality. Subject as they are to what Descombes calls "the general illusion of metaphysics", they think that the world holds constant and human beings come to represent it better and better by having fewer representations of unreal things and more representations of real things. So they become impatient when they find Descombes saying, in a passage I have already quoted, that

> **A subject's words and thoughts have the meaning that they must be given in his world and cannot be dissociated from that world…The world that must be known in order to know what the subject thinks is not just a natural world. (MP, 244)**

Cognitive scientists typically attend to the sort of world that Descombes here calls "natural". They set aside the differences between the world of twenty-first century America and that of the African savannah. This is because the transactions between organism and environment that they study are with the environment described in terms that have not changed much in the course of history---words like "brown", "cow", "and", "with," "red", and "yellow". They are interested in what our children share with chimpanzees rather than in what they share with Plato. As Steven Pinker puts it,

> **When Hamlet says "What a piece of work is man"…we should direct our attention not at Shakespeare or Mozart or Einstein pr areem Abdul-Jabbar, but at a four-year-old carrying out a request to put a toy on a shelf. (*How the mind works*, p. 4)**

From the point of view of linguistic idealism, nature is what is described by words whose use has not changed much in the course of history and culture as what is described by those whose use has changed quite a lot—a difference of degree, to be sure, but nonetheless important for that. This is, to be sure, a difference of degree, but though Plato, revived for purposes of a trans-generational colloquy, could quickly learn that "and" means *kai* and that "cow" means *bous*, it would him some weeks to get clear about why "atom" does not mean *atomon* and why "democracy" does not mean *demokrateia*. To put the contrast another way, nature is what, if you imagined a chimpanzee putting her

beliefs into words, she would talk about. Culture is the topic of all those beliefs that it never occurs to us to attribute to chimpanzees.

Notice that if we restrict ourselves to nature, the question of what input-output functions to treat the brain as realizing becomes easier. For the vocabularies we use to describe those functions will bc limited to those useful in explaining how the brain evolved, just as the reverse engineer's descriptions of the programming languages being run by her computer will be limited to those useful in explaining how the original engineers might have gone about their work.

But if we avoid what Fodor calls "meaning realism" we shall not ask whether this vocabulary describes what the brain is "really" doing, any more than we shall ask whether nature is really made up of cows and dogs as opposed to atoms and void. To say that there is a matter of fact about how the brain works will have no more metaphysical resonance than saying that there is matter of fact about the superiority of Wittgenstein to Carnap, or of Gericault to David. What counts as matter of fact will be determined by what beliefs suit our human purposes better.

One would only be inclined to representationalism over inferentialism or to meaning realism over laissez-faire pragmatism, or to preer Carnap to Wittgenstein if one were subject to what Descombes calls the "general illusion of metaphysics". Then one would think that there is a vocabulary that describes the way things really are, and that attributing truth to our beliefs means claiming that there is an isomorphism between the way mental representations are laid out and the way way things really are laid out. Linguistic idealists are the ones who insist that this view, though not happily describe as an illusion, is one that has outlived its usefulness.

The argument about whether it has or not will not be settled by the eventual fate of cognitive science. But splitting the field of application of the term 'mind' between those who wish to adopt what Chomsky calls "the standard methods of the natural sciences" and those who do not might advance the argument a little. There is an obvious sense in which the brain is the mind and an obvious sense in which it is nothing of the kind, just as there is an obvious sense in which the brain

is a mass of whirling unobservable particles and a sense in which it is not.

If we could see the unity of knowledge as the ability of people using different vocabularies to stay out of each other's way, rather that our ability to see their activities in a single synoptic vision, people on both sides of the controversy might make a little more progress. The general illusion of metaphysics, like the "idea idea", was generated by the desire for such a vision. Carnap's vision of the unity of science sustained this desire. But if we adopted Descombes' suggestion to have only partial ontologies, and if we follow up on Brandom's claim that "entity" and "identity" are sortal notions, this desire may weaken. Then we could replace questions about the place of intentionality in a world of particles with questions about the place of natural science in culture. Philosophy would then be not so much a matter of seeing how things fit together but a set of suggestions about how various social practices can coexist peaceably.

Richard Rorty
March 5, 2003