INTRODUCTION

There was a time, not so long ago, when psychologists shunned questions about the philosophy of mind in the belief that they were tedious and unanswerable compared to the problems which could be settled by experiments carried out in the laboratory. To some extent their attitude was justified by the fact that until recently philosophical speculation seemed vague and unhelpful, added to which discussion was often conducted without paying any attention to the discoveries of experimental psychology.

But as Professor Dennett points out, both philosophy and psychology have taken a turn for the better, and there are already signs of a fruitful *rapprochement*. 'Psychology has become "cognitive" or "mentalistic" and fascinating discoveries have been made about such familiar philosophical concerns as mental imagery, remembering and language comprehension'.[1] Conversely, philosophy has become noticeably less metaphysical, and by asking questions about the language which we use to describe mental processes it is at least in a position to provide helpful methods for thinking about the design and significance of psychological experiments.

Professor Dennett's career epitomises this change in the history of ideas. Trained as a philosopher both in England and the United States, his interest in logic and language has played an important part in the contribution he has made to the conversations which are now going on about the meaning and methodology of the cognitive sciences. The problem to which Professor Dennett has directed much of his attention is the traditional puzzle of the relationship between brains and minds, between a universe of discourse which includes words such as 'neurone', 'circuits', 'synapses' and 'neuro-transmitters'; and one which uses so-called intentional concepts such as belief and meaning. There are those who insist that the gap between these two domains is logically unbridgeable and that some of the traditional efforts to mediate between them have resulted in conceptually absurd proposals such as the well-known theory of an immaterial substance halfway between mind and matter. An alternative theory is the so-called reductionist one, which argues that with improvements in neurophysiological knowledge the language of mental states will eventually be replaced by an accurate account of so-called brain states. Philosophers who support this theory look forward to the time when it will be possible to characterise someone's beliefs by reading the nervous inscriptions written within the substance of the brain. What makes Professor Dennett's analysis so interesting is the way in which he breaks down the logical implications of such proposals, and instead of supplying yet another theory about the relationship between brains and minds, he provides a versatile and productive method for weighing the value of any theory which might be offered. In the interview that follows, Professor Dennett has restored philosophy to its proper place as the handmaiden of the exact sciences.

[1]Daniel C. Dennett: *Brainstorms*, Harvester, 1978.

DISCUSSION

MILLER Nowadays there seems to be a division between those who believe that the most profitable way of understanding how the mind works and how the brain works is to go at it from the top down, by analysing and examining and experimenting with the whole person. And there are those, on the other hand, who say that it's much more profitable to try and analyse the whole process from the ingredient elements upwards. Where would your money be? Do you think that it's going to be more profitable to go from above-down, or from below-up? Or is this in fact even a reasonable or understandable distinction?

DENNETT The top-down, bottom-up division is perfectly real and the subject of debate among people in this field. The neuroscientists, generally, are in favour of an abstemious bottom-up approach, in which one says: 'We're all agreed that the brain is an information processing organ of some sort, and we can tell that the *atomic* information processors, the basic building blocks, are the neurons. So now if we can just understand how each neuron processes information, we can then build molecules, larger structures of neurons. In effect: combining the effects of the individual neurons, we will determine what these large structures process in the way of information and gradually we'll work our way up – not really analysing but synthesising: building larger and larger structures of brain tissue which can be seen to process information.' This, they say, is the only proper way of doing the science involved.

Whereas the other group, the so-called top-down cognitive scientists, say that this is grasping the wrong end of the stick. You'll never make much progress that way. Rather you should start with the whole person, and say: here's a whole person, and these are his cognitive talents. This is what he knows, this is what he believes, this is what he can see and these are the problems he can solve. Now let's try to decompose this whole person into assemblies or sub-systems, which in complex interaction can explain the talents of the whole person. Then we can decompose each sub-system in turn into its component sub-systems, and so on, until we have finally reached components that can be securely identified by the neuroanatomists.' That is the top-down strategy.

Now, both strategies ought to work in principle, because both ought to end up having completed exactly the same task. In the end both sides want to understand the relationship between the brain and the mind. And you can either start with the mind and work down, or you can start with the bits of brain and work up. So if you compare this with the analogy of building a trans-continental railroad, you *do* start at both ends, and plan to meet somewhere in the middle. I would bet, however, that most of the track is going to be laid by the people who are working from the top down, rather than from the bottom up. For a very simple reason: top-down is much easier, as it turns out.

MILLER Now exactly why is it easier? One would have thought that there is certainly something easy about sticking a needle into the brain, finding a neuron, and then actually recording the action potentials and then doing it again to another neuron. There's something easy about the hardware of such a situation. It's ticklish, but it's easy.

DENNETT Yes, each little bit of data-gathering is easy enough, but the trouble is there's so much of it, and you don't know what further questions to ask. Those who think that the bottom-up strategy is the best way to proceed should be reminded of a system which is, in fact, much more simple than a human brain; a commercial chess-playing computer of the sort you can buy in the stores for Christmas. There it is, a little box, with a chessboard on the top, and it plays chess. It plays pretty good chess. How does it do it?

Well exponents of the bottom-up approach would say: 'Let's get our microscope and start taking the thing apart and looking at the individual tiny microchips and the relays on them and see if we can understand how each little part works, and then build up from there.' The trouble is if you start at that level of detail, you don't know what questions to ask. It's virtually inconceivable that somebody who didn't already know how a chess-playing computer was organised, and how computers in general were organised at a much higher level, could ever work his way up from a physical analysis of those little tiny parts.

For one thing, when would one ever get to the level at which one would be describing chess-playing, which after all is what the thing does? The top-down theorist on the other hand says: 'Let's think about it in a more or less *a priori* fashion first. The thing plays chess – what do you *have to know* to play chess? Well you have to know the rules, for one thing, so we'll bet that somewhere in this there's a representation of the rules – at least implicitly. That information is somehow in the system. You also have to know the principles of chess play. So somewhere we're going to find explicitly or implicitly represented, in something or other in there, the rules and principles of chess. Moreover, to play chess one has to know the position of the pieces on the board, so the system has to know what its move has been, and has to know where all the pieces are.'

MILLER And all of these questions can be asked of the machine, without ever asking oneself questions about the sort of material out of which these representations are going to be made.

DENNETT Not only can you be ignorant of the *sort* of material, you can even be ignorant of the *structure* of the material. That is, you can treat the question first purely at the level of the information – the beliefs and desires if you will – of any agent that can play chess. What information and what goals does such an agent require, whether it's an electronic agent, a sort of pseudo-agent, or a full-blooded human agent? What is the *content* of the strategic information, the tactical information, and the information about the rules that would enable *anything* to play chess, except by sheer magic?

So the top-down advocates asks those questions first. They then postulate interacting sub-assemblies of information processors that might do the job. Then one can test these postulates by designing and actually building models, and gradually working down to the level where one can finally tell somebody doing the bottom-up approach: 'Look for something in the brain that does the following rather fancy jobs; it should have the following sorts of features.' At some point (they hope) they are going to be able to say: 'now we've described some sub-assemblies that have certain talents; we bet you bottom-up people can find those – maybe you already have. If you haven't we can at least give you hints about what to look for.'

MILLER So you're saying that unless one has some sort of idea about what that machine is out to do, or what it's set up to do, one wouldn't actually know what to start looking for at the bottom level.

DENNETT Well, of course, the people who do bottom-up research can engage in informed speculation; they frame hypotheses about what the functions are going to be that they should look for. But that's not really their line of country, and they haven't developed a perspicuous vocabulary for describing such functions; they haven't trained themselves to think in terms of information processing models. Everybody says the brain's job is information processing, but that's a very tricky notion.

MILLER But presumably even the fact that they approach it with that sort of theoretical preconception means that they have imported, perhaps without knowing it, ideas from the top-down approach, in that they're already looking at the elements with a view to there being an information processing device, which itself is a top-down idea, isn't it?

DENNETT Yes. We shouldn't over-dramatise the division here. And after all there are people doing brilliant work who are trying very hard to pay attention to both ends of the track. And each strategy – and that's what they are; they aren't really competing methodologies, they're just competing strategies – each has its own foibles. The weakness of the bottom-up strategists is, as I've already said, that they don't know the right questions to ask; they don't know what hypotheses to test. The complementary weakness of the top-down people is that very often, they don't know enough about what the constraints on the hardware finally are going to be. They don't know enough about what nerve tissue can and can't do, so they are apt to design a model which might work in some other world, but couldn't really be done by the brain, given the way the brain is actually composed.

So one should, of course, try to keep alert to both of these. But people *do* specialise, and as I say, I'm putting my bet on the top-down approach to be much more dramatically fruitful in the foreseeable future, simply because it generates more fruitful hypotheses to test.

MILLER Why does that approach generate so many serviceable hypotheses about the brain?

DENNETT Well, it's because the idea of information (which is still not very well formulated) is nevertheless a very powerful generator of experiments. It's powerful because it allows you to abstract from messy details, in a certain way. When you're ignorant of a lot of messy details, if you can find a way of