



confusion over consciousness.

Part of the problem is its reflexive nature. Consciousness cannot be studied in the same way as the things I am conscious of. If I turn my attention to it, it ceases to be consciousness and becomes just another thought or experience.

Another problem is its ever-changing quality. What it means to be conscious is essentially what it is like being me, here and now. And yet as soon as I think that, here and now are already gone: I, and the world around me, have changed. That is why William James, one of the 19th-century pioneers of psychology, used the phrase 'the stream of consciousness'. It is a stream that is flowing; it feels unbroken but never repeats itself exactly. Science searches for patterns and regularity, so it is hard to know where to start with something whose very nature lies in change.

The best starting place with such problems is always to re-examine the question. Asking 'What is consciousness?' is, according to many people, a meaningless question, for consciousness is not a 'thing' at all.

A more tractable question is, 'Why do we have consciousness?' It might be possible to imagine animals (or even machines) that could carry out all the actions I carry out and yet be blissfully unaware of anything at all. Those in pain, or who enjoy the occasional oblivion of alcohol, might even think this would be preferable. So why did consciousness evolve? Evolution must have had some reason for making us conscious - mustn't it?

John Crook, an ethologist at the University of Bristol, has argued that human consciousness is distinctive because we are conscious of being conscious; we have a sense of having personal power over our actions and we have self-identity. These together with language, arose, he argues, because our ancestors began using tools. Only an animal that needs to know who owns and uses which tools needs such a strong sense of personal identity.

By contrast the psychologist and writer, Nicholas Humphrey, argues that we evolved self-awareness for social reasons: to hold together a cooperative society, early humans needed to understand their fellows - to become, in Humphrey's words, *Homo psychologicus*. The best way to do that, Humphrey argues, was to use a 'privileged picture' of themselves as a model for what it is like to be another person. So, according to Humphrey, we learnt to look into our own mental processes so as to be able to predict (and hence control?)

the desires or actions of others.

This connection between consciousness and representing the self is made even more explicit by Richard Dawkins, a zoologist at the University of Oxford. He says: 'Perhaps consciousness arises when the brain's simulation of the world becomes so complete that it has to include a model of itself.' And yet there is something wrong here. The brain is not really modelling itself. It is not an image of neurons and glial cells that is central to my self-consciousness. Surely people who have never learnt about the structure and function of the brain are just as conscious as neurosurgeons.

Was Homer unaware of himself?

According to all these theories, self-awareness would have appeared very early in the evolution of humans. By contrast, an American psychologist, Julian Jaynes, proposed the controversial thesis that consciousness appeared only in historical times.

The early Greeks and the Hebrews of the early Old Testament period did not, according to Jaynes, experience themselves as thinking beings. Instead of attributing their verbal images to internal processes they attributed them to the voices of the Gods. He cites the Iliad, written about 3000 years ago, as having no references to mental concepts such as mind, thoughts, feelings, or even self. People were not to be blamed or held responsible for their actions: they were only instruments of external forces.

This interesting idea puts our views of ourselves in a new light. If the Gods were only a theory invented to account for behaviour, perhaps our present-day concept of self is just another such theory invented to explain ourselves to ourselves. It may be a better theory than the Gods but is it really accurate? Isn't it just another useful illusion?

The problem with all these theories is that none of them directly addresses the question of consciousness itself. Why couldn't a system evolve with a good concept of self, attributing behaviour to internal processes - even with psychological skills modelling other people's (non-conscious) behaviour - without actually feeling like anything? These theories may show how self-identity or the ability to predict behaviour arose, but they say nothing directly about awareness. I can still imagine an animal or machine that does all these things and is completely unaware.

Perhaps, you might say, it is not just passive awareness that distinguishes us from unconscious automata, but the fact that we can

consciously control what we do. Recent research, however, reveals that this too may be (partly or wholly) another illusion.

Michael Gazzaniga, a neuroscientist at Cornell University in New York State, has studied 'split brain' patients - the left and right sides of their brains have been surgically separated (often as a treatment for severe epilepsy). In most people, the ability to use language is by and large localised in the left brain; but one of his patients had some verbal ability in both halves, although only the left could produce speech. When a written command, such as 'laugh', was presented only to the part of the visual field linked to the right brain, he laughed. When asked 'Why did you laugh?' he answered, using the left side of the brain, but simply fabricated a reason: 'Oh, you guys are really something!' The left hemisphere had apparently observed the laughter and tried to account for it somehow. This may be no quirk of split brains. Most of our reasons for action may be totally unavailable to conscious introspection. Our verbal selves may make up plausible reasons for the actions they observe their body making.

Research on split brains has revealed much more about the nature of consciousness. In some cases, each half brain displays separate desires, intentions and even hopes for the future - and a sense of self. It is easy to think that splitting the brain has split an originally single consciousness, but Gazzaniga believes that the surgery only reveals a general principle: that human minds are multiple entities consisting of many subsystems. It is only the ability to put things into words that creates 'a personal sense of conscious reality out of the multiple systems present', he believes.

So is the idea of a single conscious 'will' totally false? Here again some recent research provides a clue. Libet turned his attention from conscious sensation to voluntary action and devised an ingenious experiment based on earlier work. Using electrodes fixed to the scalp, other researchers had found that just before someone does something voluntarily the electrical potential of the scalp shifts to the negative, a phenomenon called the readiness potential.

Electrodes pick up readiness potential a second or more before any apparently voluntary movement begins. We might assume, then, that the conscious decision to act must come before the readiness potential, if a conscious decision to act is what starts the train of events.

Libet set out to test this. He asked subjects to flex their wrist or fingers at any time they felt the 'urge' or desire to do so. To

measure the timing of that urge or desire, he asked them also to watch a revolving spot and to report its clock position at the time they felt the urge to move.

Consciousness is an afterthought

His findings were consistent and surprising. The readiness potential came first and the desire to move about 400 milliseconds (nearly half a second) later. The implication seems to be that even in apparently spontaneous voluntary acts, an unconscious brain event happens well before any conscious desire or decision to act. Again, the consciousness seems to be an afterthought.

Libet went on to show that subjects could 'veto' the action within a period of between 100 and 200 milliseconds before the action would have started. This, he argues, still leaves some potential role for the 'will'. Other scientists have argued, however, for more extreme interpretations of his data. The experiments could mean that we do not directly experience an intention to act at all.

Rather, we might infer an urge or desire to act after the process has already begun unconsciously. Because this urge precedes the action, we can keep up the illusion that it causes the action - but it is only an illusion. According to this view, consciousness has no active role at all.

This dispute reflects the long philosophical argument over whether mental events, anything from the desire for ice cream to the sensation of pain, can cause physiological events in the brain. Libet maintains they can. He thinks there is still room for the 'will', which can intervene to stop physiological processes when required. Others prefer to use his findings as evidence that they cannot. The important question for us is whether science is on the verge of turning a difficult philosophical problem into an empirical one.

Many researchers have argued that consciousness can be modelled by computing systems, and many debates focus on whether consciousness is then associated with the highest level of such a system. Such discussions may help to explain aspects of the working brain. But no such theory can explain why the contents of one part (whichever part it may be) should have the quality of feeling like something while the rest do not.

The American philosopher Thomas Nagel made this crucial point about consciousness in his now famous paper entitled 'What is it like to

be a bat?', published in his collection of essays *Mortal Questions*. His point was that when we say something is conscious, we essentially mean that there is something it is like to be that thing. Now if humans are complex information-processing systems, why should there be 'something it is like to be' some levels of that system but not others? Whichever part of the system you choose, the essential mystery remains untouched.

A radical solution is to sweep this question away and say that all mental models are conscious. Mental models are the internal representations of the world which computers, as well as animals and humans, use to control their behaviour. You could not get around without a good model of your own body and the world around it. I am suggesting that it is not like anything to be the skin, blood or bones of a bat, but it is like something to be the bat's model of a bat. This is equivalent to saying that there cannot be such things as not-conscious mental models.

The major objection to this theory is that it does not seem to us as though all the many models our brain constructs are conscious. But the answer all hinges upon who that 'us' is. What am 'I' who is to be conscious of all that activity? We can suggest an answer. 'I' am only one of the models in the system, a model of 'self in the world' built largely by language. 'I' am a self-image, a body image, a construct of a human being. This makes it perfectly obvious why all the rest of the system appears to be unconscious. It is unconscious to 'me' but not to itself.

Before we start thinking of human systems as peopled by infinite conscious models, we should reflect what it would be like to be most of those models. I think the answer is not much. For example, the models in the lower levels of perceptual processing entail no concept of self, action or an external world. Their consciousness would be correspondingly limited. Only the complex model of 'self in the world' sustains full reflexive awareness - consciousness of being conscious. It is this which seems to be 'me'.

Most of us have only one of these (unless you include any 'dream selves'). Presumably people with so-called multiple personalities have several models of self, each conscious in its own way and competing for dominance. This situation is abnormal, but now we can see it as only an extreme form of the normal case. Every human brain constructs multiple conscious models, but there is only one 'me', my model of self. This makes sense, too, of the phenomenon of selective attention. When 'I' turn my attention to something, that thing seems to come into my awareness. From a computational point of view, we

might say that the system has incorporated the model of that thing into the model of self. So 'I' become aware of it, and it of 'me'.

In this view, there need be no top and bottom, no one consciousness controlling the rest. Consciousness is not something that controls. It is not a thing, a place, a substance or a part of the system. Indeed, it has no function. It is only what it is like being a mental model.

This approach transforms the whole evolutionary question. There is no definite point at which consciousness arises, nor any purpose for it at all. Any organism that constructs representations (as even the most primitive will do) will have correspondingly primitive consciousness. Those that model a self will also be conscious of self. As evolution progresses, the quality of consciousness depends on the kinds of models constructed. So perhaps Jaynes had a point. Once the voices of the Gods were replaced by a concept of an active self, self-consciousness took another step forward. And what will our next step be? On this view, the mystics' search for higher consciousness or the Buddhists' training for 'enlightenment' may be no idle fancy but steps into new mental models. Mind and meditation

At the heart of Buddhist training lie the skills of meditation and mindfulness. Pursued far enough, these are supposed to transform consciousness and free the trainee from suffering. Does it help us to understand this transformation to ask how the models of reality are changed?

Most of us, claimed the Buddha, are constantly distracted by sights, sounds and ideas. We live in a whirl of confusion and fantasy: more in the past or future than in the present moment. Struggling to find happiness, we cling desperately to our concept of self and the things we think will make that self happy. But this, paradoxically, is precisely what leads to suffering.

By contrast, mindfulness means living in the present moment, every moment, constantly alert. In meditation, this skill is practised in quiet sitting. Any thoughts that arise are let go again. They may come again and again, but there is no clinging, no leap into the building of fantasies. They just come and go until eventually the mind is still.

These contrasting ways of being can be understood in terms of the mental models that are constructed. In the normal way, attention shifts from one thing to another. Surprising events grab the

attention: other chains of thought wait to be finished as soon as there is a gap. So there is never any peace. This is efficient in using all available processing capacity, but what does it feel like to be the models in such a system? I suppose it feels like most of us do feel - pretty confusing. The only thing that gives it any stability is the constant presence of a stable self model. No wonder we cling to it.

On the other hand, being mindful means not following every association; not stacking up ideas to be completed; not rehearsing, planning or even selecting. The models become ever simpler and all processing capacity is not immediately used up. More and more can be linked into one model because there is less and less being modelled. And what would it be like to be the models in such a system?

Obviously very different. Is this why everything seems brighter and more 'real' during meditation? Why the quietest sound is clearly heard and the beating of one's heart is constantly present? If so, testable predictions might follow. Long-term practitioners might show greater awareness of normally inaccessible low-level modelling, for example.

In the end, it is not even necessary for the model of self to be constructed at all. Imagine what it would be like to be a system that processes incoming information but builds no further constructions upon it, not even any self to observe. With no modelled distinction between self and other, I imagine the world would all seem one. (This condition might make survival problematic, however.)

People who have this training and experience find it very hard to describe what happens, which is no wonder when 'they' (their model of self) have been dissolved. It may, in fact, be impossible to describe using the assumptions of our usual language. But perhaps a new possibility now arises: that science might at least develop the concepts and language needed to understand mystical experiences in terms of information processing.

Buddhism teaches the doctrine of 'no-self'. There is no thing that is conscious but just consciousness itself. Cognitive psychology reveals that the self is a mental model; if consciousness is just what it is like being a mental model, there is no self being conscious of anything, but just a series of changing models. So can I now answer my question: 'Why am I aware of anything?' The answer seems to be that because I am only a mental model, it could not be otherwise.

