

|    |  |    |    |  |     |
|----|--|----|----|--|-----|
| 17 | Teletransporter<br>Jolyon Troschianko  | 74 | 21 | The old hag<br>Jolyon Troschianko                        | 102 |
| 18 | Brain areas involved<br>in volition<br>Spence, S. A. and Frith, C. D.<br>(1999) 'Towards a functional<br>anatomy of volition. <i>Journal<br/>of Consciousness Studies</i> ,<br>6(8-9), 11-29 | 83 | 22 | Mapping states of<br>consciousness<br>Jolyon Troschianko | 109 |
| 19 | Libet's experiment<br>Jolyon Troschianko   | 87 | 23 | What is it like to be<br>a bat?<br>Jolyon Troschianko    | 118 |
| 20 | Wegner's theory of<br>conscious will<br>From D. Wegner, <i>The Illusion<br/>of Conscious Will</i> (MIT Press,<br>2002), p. 68  | 97 | 24 | Chimp with a mirror<br>Jolyon Troschianko                | 121 |
|    |  |    | 25 | Looking in the fridge<br>Jolyon Troschianko              | 130 |

The publisher and the author apologize for any errors or omissions in the above list. If contacted they will be pleased to rectify these at the earliest opportunity.

## Chapter 1


### Why the mystery?

#### The hard problem

What is consciousness? This may sound like a simple question but it is not. Consciousness is at once the most obvious and the most difficult thing we can investigate. We seem either to have to use consciousness to investigate itself, which is a slightly weird idea, or to have to extricate ourselves from the very thing we want to study. No wonder that philosophers and scientists have struggled for millennia with the concept, and that scientists rejected the whole idea for long periods and refused even to study it. The good news is that, at the start of the 21st century, 'consciousness studies' is thriving. Psychology, biology, and neuroscience have reached the point when they are ready to confront some tricky questions: What does consciousness do? Could we have evolved without it? Could consciousness be an illusion? What is consciousness, anyway?

This does not mean that the mystery has gone away. Indeed, it is as deep as ever. The difference now is that we know enough about the brain to be ready to confront the problem head on. How on earth can the electrical firing of millions of tiny brain cells produce this – my private, subjective, conscious experience?

If we are going to get anywhere with understanding consciousness, we have to take this problem seriously. There are many people who



1. No one has yet succeeded in bridging the fathomless abyss, the great chasm or the explanatory gap between inner and outer, mind and brain, or subjective and objective.

claim to have solved the mystery of consciousness: they propose grand unifying theories, quantum mechanical theories, spiritual theories of the 'power of consciousness', and many more, but most of them simply ignore the yawning chasm, or 'fathomless abyss', between the physical and mental worlds. As long as they ignore this problem they are not really dealing with consciousness at all.

This problem is a modern incarnation of the famous mind-body problem with which philosophers have struggled for more than two thousand years. The trouble is that in ordinary human experience there seem to be two entirely different kinds of thing, with no obvious way to bring the two together.

On the one hand, there are our own experiences. Right now I can see the houses and trees on a distant hill, hear the cars down on the main road, enjoy the warmth and familiarity of my own room, and wonder whether that scratching noise is the cat wanting to be let in.

All of these are my own private experiences and they have a quality that I cannot convey to anyone else. I may wonder whether your experience of green is the same as mine or whether coffee has exactly the same smell for you as it does for me, but I can never find out. These ineffable (or indescribable) qualities are what philosophers call 'qualia' (although there is much dispute about whether qualia exist). The redness of that shiny red mug is a quale; the soft feel of my cat's fur is a quale; and so is that smell of coffee. These experiences seem to be real, vivid, and undeniable. They make up the world I live in. Indeed, they are all I have.

On the other hand, I really do believe that there exists a physical world out there that gives rise to these experiences. I may have doubts about what it is made of, or about its deeper nature, but I do not doubt that it exists. If I denied its existence I would not be able to explain why, if I go to the door, I shall probably see the cat rushing in – and if you came by you would agree that there was now a cat trailing muddy footprints across my desk.

The trouble is that these two kinds of thing seem to be utterly different. There are real physical things with size, shape, weight, and other attributes that everyone can measure and agree upon, and then there are private experiences – the feeling of pain, the colour of that apple as I see it now.


Throughout history most people have adopted some kind of dualism: that is the belief that there are indeed two different realms or worlds. This is true of most non-Western cultures today, and surveys suggest that it is true of most educated Westerners as well. The major religions are almost all dualist: Christians and Muslims believe in an eternal, non-physical soul, and Hindus believe in the Atman or divine self within. Among religions, Buddhism alone rejects the idea of a persisting inner self or soul. Even among non-religious people, dualism is prevalent in Western cultures. Popular New Age theories invoke the powers of mind, consciousness, or spirit, as though they were an independent force;

and alternative therapists champion the effect of mind on body, as though these were two separate things. Such dualism is so deeply embedded in our language that we may happily refer to 'my brain' or 'my body'; as though 'I' am something separate from 'them'.

In the 17th century the French philosopher René Descartes (1596–1650) formally proposed the best-known dualist theory. Known as Cartesian dualism, this is the idea that the mind and the brain consist of different substances. According to Descartes, the mind is non-physical and non-extended (i.e. it takes up no space or has no position), while the body and the rest of the physical world are made of physical, or extended, substance. The trouble with this is obvious. How do the two interact? Descartes proposed that they meet in the tiny pineal gland in the centre of the brain, but this only staves off the problem a little. The pineal gland is a physical structure and Cartesian dualism provides no explanation of why it, alone, can communicate with the mental realm.

This problem of interaction bedevils any attempt to build a dualist theory, which is probably why most philosophers and scientists completely reject all forms of dualism in favour of some kind of monism; but the options are few and also problematic. Idealists make mind fundamental but must then explain why and how there appears to be a consistent physical world. Neutral monists reject dualism but disagree about the fundamental nature of the world and how to unify it. A third option is materialism and this is by far the most popular among scientists today. Materialists take matter as fundamental, but they must then face the problem that this book is all about. How do you account for consciousness? How can a physical brain, made purely of material substances and nothing else, give rise to conscious experiences or ineffable qualia?

This problem is called the 'hard problem' of consciousness, a phrase coined in 1994 by the Australian philosopher David Chalmers. He wanted to distinguish this serious and overwhelming difficulty from what he called the 'easy problems'. The easy problems, according to



2. Descartes explained reflex responses to pain in terms of mechanical responses and the flow of 'animal spirits' in tiny tubes. But when it came to conscious experiences he proposed that they were part of a quite different mental world, connected to the physical body through the pineal gland in the centre of the brain.

Chalmers, are those that in principle we know how to solve, even if we have not yet done so. They include such problems as perception, learning, attention, or memory; how we discriminate objects or react to stimuli; how sleep differs from being awake. All these are easy, he says, compared with the really hard problem of experience itself.

Not everyone agrees with Chalmers. Some claim that the hard problem does not exist; that it depends on a false conception of consciousness, or on drastically underestimating the 'easy' problems. The American philosopher Patricia Churchland calls it a



'hornswoggle problem', arguing that we cannot, in advance, decide which problems will turn out to be the really hard ones. It arises, she claims, from the false intuition that if we explained perception, memory, attention, and all the other details, there would still be something left out – 'consciousness itself'.

These are important objections. So before we go any further we must be clearer about what, if anything, 'consciousness itself' might mean.

## Defining consciousness

What is it like to be a bat? This curious question looms large in the history of consciousness studies. First asked in the 1950s, it was made famous by the American philosopher Thomas Nagel in 1974. He used the question to challenge materialism, to explore what we mean by consciousness, and to see why it makes the mind-body problem so intractable. What we mean, he said, is *subjectivity*. If there is something it is like to be the bat – something *for the bat itself*, then the bat is conscious. If there is nothing it is like to be the bat, then it is not.

So think, for example, of the mug, or pot, or plastic ornament on your table. Now ask – what is it like to be the mug? You will probably answer that it is like nothing at all; that mugs cannot feel, that china is inert, and so on. You will probably have no trouble in opining that pots and mugs are not conscious. But move on to worms, flies, bacteria, or bats and you may have more trouble. You do not know – indeed, you cannot know – what it is like to be an earthworm. Even so, as Nagel points out, if you think that there is something it is like to be the worm then you believe that the worm is conscious.

Nagel chose the bat as his example because bats are so very different from us. They fly, live mostly in the dark, hang upside-down from trees or in damp caves, and use sonar, not vision, to see the world.

## Defining consciousness

There is no generally agreed definition of consciousness, but the following gives some idea of what is meant by the word.

**'What it's like to be . . .':** If there is something it is like to be an animal (or computer, or baby) then that thing is conscious. Otherwise it is not.

**Subjectivity or phenomenality:** Consciousness means subjective experience or phenomenal experience. This is the way things seem to me, as opposed to how they are objectively.

**Qualia:** The ineffable subjective qualities of experience, such as the redness of red or the indescribable smell of turpentine. Some philosophers claim they do not exist.

**The hard problem:** How do subjective experiences arise from objective brains?

That is, they emit rapid bursts of high-pitched squeaks while they fly and then, by analysing the echoes that come back to their sensitive ears, learn about the world around them.

What is it like to experience the world this way? It is no good imagining that you are a bat because an educated, speaking bat would not be a normal bat at all; conversely, if *you* became a normal bat and could not think or speak then you would not be able to answer your own question.

Nagel argued that we can never know and from this concluded that the problem is insoluble. For this reason he is dubbed a *mysterian*. Another mysterian is the American philosopher Colin McGinn, who argues that we humans are 'cognitively closed' with respect to

understanding consciousness. That is, we have no hope of understanding consciousness, just as a dog has no hope of being able to read the newspaper he so happily carries back from the shops. Psychologist Stephen Pinker agrees: we may be able to understand most of the detail of how the mind works, yet consciousness itself may remain forever beyond our reach.

Not many people share Nagel's pessimism, but his question has proved helpful in reminding us what is at stake when we talk about consciousness. It is no good talking about perception, memory, intelligence, or problem-solving as purely physical processes and then claiming to have explained consciousness. If you are really talking about consciousness, then you must deal in some way or another with subjectivity. Either you must actually solve the hard problem and explain how subjectivity arises from the material world, or alternatively, if you claim that consciousness is identical to those physical processes, or is an illusion or even that it does not exist at all, you must explain why it *appears* so strongly to exist. Either way, you can only claim to be dealing with consciousness if you are talking about 'What it is like to be . . . '.

This essential meaning of the term consciousness is also called phenomenality, or phenomenal consciousness, terms coined by American philosopher Ned Block. Block compares *phenomenal consciousness*, which is what it is like to be in a certain state, with *access consciousness*, which refers to availability for use in thinking, or guiding action and speech. Phenomenal consciousness (or phenomenality, or subjectivity) is what Nagel was talking about and is the core of the problem of consciousness.

With these ideas in mind, we are ready for one of the central disputes in consciousness studies. This concerns the following question: Is consciousness an extra ingredient that we humans have in addition to our abilities of perceiving, thinking, and feeling, or is it an intrinsic and inseparable part of being a creature that can perceive and think and feel? This really is the key question on which

the rest depends, and you might like to decide now what you think about it, for the implications either way are quite striking.

On the one hand, if consciousness is an extra added ingredient then we naturally want to ask why we have it. We want to ask what consciousness is for, what it does, and how we got it. On this view, it is easy to imagine that we might have evolved without it, and so we want to know why consciousness evolved, what advantages it gave us, and whether it evolved in other creatures too. On this view, the hard problem is indeed hard; and the task ahead is to answer these difficult questions.

On the other hand, if consciousness is intrinsic to complex brain processes and inseparable from them, then it is nonsensical to ask most of these questions. On this view (which in some versions is called functionalism), there is no use in asking why consciousness evolved, because any creature that evolved to have intelligence, perception, memory, and emotions would necessarily be conscious as well. Also there would be no sense in talking about 'consciousness itself' or about 'ineffable qualia', for there is nothing extra that exists apart from the processes and abilities.

On this view, there really is no deep mystery, and no hard problem. So the task is quite different; it is to explain why there *seems* to be such a problem and why we *seem* to be having ineffable, non-physical, conscious experiences. It is here that the idea of consciousness as an illusion comes in, for neither consciousness nor the hard problem are what they seem, and so we must explain how the illusion comes about.

If the implications of this dichotomy seem hard to grasp, a thought experiment might help.

## Zombie

Imagine someone who looks exactly like you, acts like you, thinks like you, and speaks like you, but who is not conscious at all. This other you has no private, conscious experiences; all its actions are carried out without the light of awareness. This unconscious creature – not some half-dead Haitian corpse – is what philosophers mean by a zombie.

Zombies are certainly easy to imagine, but could they really exist? This apparently simple question leads to a whole world of philosophical difficulties.


On the 'yes' side are those who believe that it really is possible to have two functionally equivalent systems, one of which is conscious while the other is unconscious. Chalmers is on the 'yes' side. He claims that zombies are not only imaginable but possible – in some other world if not in this one. He imagines his zombie twin who behaves exactly like the real Chalmers but has no conscious experiences, no inner world, and no qualia. All is dark inside the mind of zombie-Dave. Other philosophers have dreamed up thought experiments involving a zombie earth populated by zombie people, or speculated that some real live philosophers might actually be zombies pretending to be conscious.

On the 'no' side are those who believe the whole idea of zombies is absurd, including both Churchland and American philosopher Daniel Dennett. The idea is ridiculous, they claim, because any system that could walk, talk, think, play games, choose what to wear, enjoy a good dinner, and do all the other things that we do, would necessarily have to be conscious. The trouble is, they complain, that when people imagine a zombie they cheat: they do not take the definition seriously enough. So if you don't want to cheat, remember that the zombie has to be completely indistinguishable from a normal person on the outside. That is, it is no good asking the zombie questions about its experiences or

testing its philosophy, for *by definition* it must behave just as a conscious person would. If you really follow the rules, the critics say, the idea disappears into nonsense.

It should now be easy to see that the zombie is really just a vivid way to think about the key question: Is consciousness a special added extra that we conscious humans are lucky to have, or is it something that necessarily comes along with all those evolved skills of perceiving, thinking, and feeling? If you believe that it's an added extra, then you can believe that we might all have evolved as zombies instead of as conscious people – and even that your neighbour might be a zombie. But if you believe that it's intrinsic and inseparable from the skills we humans have, then zombies simply could not exist and the whole idea is daft.

I think the whole idea is daft. Nevertheless, it remains extremely alluring, largely because it is so easy to imagine a zombie. Yet how easy something is to imagine is not a good guide to its truth. So let's consider a rather different aspect of the same problem – the question of whether consciousness does anything.



3. The idea of the philosopher's zombie leads only to confusion



The phrase 'the power of consciousness' is common in popular discourse. The idea is that consciousness is some sort of force that can directly influence the world – either by acting on our own bodies, as when 'I' consciously decide to move my arm and it moves – or, more controversially, in things like psychic healing, telepathy, or 'mind over matter'. Like the zombie, this 'power' is easy to imagine. We can visualize our conscious mind somehow reaching out and influencing things. But does this idea make any sense? As soon as you remember that consciousness means subjectivity or phenomenality, then the idea begins to seem less plausible. How could 'what it's like to be' something be a force or power? How could my *experience* of the green of that tree cause something to happen?

One way to explore whether consciousness could be a power or force is to ask what would happen if you took it away. Obviously, if consciousness has any power at all, what would be left could not be a zombie because the zombie must, by definition, be indistinguishable from a conscious person. So you would be left with someone who was different from a conscious person because they could not . . . what?

Perhaps you think consciousness is needed for making decisions, but we know a lot about how the brain makes decisions and it does not seem to need an extra added force to do so. Also, we can make computers that make decisions without a special consciousness module. The same goes for seeing, hearing, controlling movements, and many other human abilities. Perhaps you think it is needed for aesthetic appreciation, creativity, or falling in love, but, if so, you would have to show that these things are done by consciousness itself rather than by the workings of a clever brain.

All this leads to the awkward notion that perhaps consciousness does nothing, and other oddities point the same way. For example, think about people catching cricket balls, playing table tennis, or interrupting fast-flowing conversations. These quick actions all

seem to be done consciously, but is it the consciousness itself that makes them happen? In fact, as we shall see, such actions happen too fast, and they are coordinated by parts of the brain that appear not to be involved in conscious experience.

Could consciousness, then, be completely powerless? One version of this idea is epiphenomenalism – the idea that consciousness is a useless by-product or epiphenomenon. This is a very curious notion because it entails consciousness actually existing but having no effects on anything else. And if it has no effects at all it is hard to see how we could end up worrying about it – or even talking about it.

But epiphenomenalism is not the only way of understanding consciousness as powerless. An alternative is to say that all creatures like us that can see, feel, think, fall in love, and appreciate a fine wine will inevitably end up believing they are conscious, imagining zombies are possible, and thinking that consciousness does things. The bottom line for this kind of theory is that we are deluded; we feel as though consciousness is a power or added ability, but we are wrong. If this theory needs a name, we might call it 'delusionism'.


I think this is the right way to think about consciousness, but it implies that our ordinary assumptions about consciousness are deeply misguided. Could we really be so wrong? And why should we be? Perhaps we should take a closer look at some of those assumptions and ask how reliable they are.

## The theatre of the mind

The most natural way to think about consciousness is probably something like this. The mind feels like a private theatre. Here I am, inside the theatre, located roughly somewhere inside my head and looking out through my eyes. But this is a multi-sensational theatre. So I experience touches, smells, sounds, and emotions as well. And I

can use my imagination too – conjuring up sights and sounds to be seen as though on a mental screen by my inner eye or heard by my inner ear. All these are the ‘contents of my consciousness’ and ‘I’ am the audience of one who experiences them.

This theatre imagery fits happily with another common image of consciousness – that it flows like a river or stream. In the 19th century, the ‘father of modern psychology’, William James (1842–1910), coined the phrase ‘the stream of consciousness’ and it seems apt enough. Our conscious life really does feel like a continuously flowing stream of sights, sounds, smells, touches, thoughts, emotions, worries, and joys – all of which happen, one after another, to me.



4. I feel as though I am somewhere inside my head looking out – that I experience the outside world through my eyes and ears, imagine things in my mind’s eye, and direct my arms and legs to walk me down the street and post the letter. But the brain cannot work this way. This is Dennett’s mythical Cartesian theatre.

This way of conceiving of our own minds is so easy, and so natural, that it hardly seems worth questioning. Yet when we get into an intellectual muddle, as we seem to have done with the problem of consciousness, it is sometimes worth challenging our most basic assumptions – in this case, these apparently innocent analogies.

The strongest challenge comes from philosopher Daniel Dennett. He argues that while most people are happy to reject the idea of Cartesian dualism, they still retain strong vestiges of dualist thinking in the form of what he calls the Cartesian theatre. This is not just the analogy of the mind with a theatre, but the notion that somewhere in the mind or brain there must be a place and time at which everything comes together and ‘consciousness happens’; that there is some kind of finishing line in the brain’s activities, after which things mysteriously become conscious or ‘enter consciousness’.

This has to be false, claims Dennett. To begin with, there is no centre in the brain which could correspond to this notion, for the brain is a radically parallel processing system with no central headquarters. Information comes in to the senses and is distributed all over the place for different purposes. In all of this activity there is no central place in which ‘I’ sit and watch the show as things pass through my consciousness. There is no place in which the arrival of thoughts or perceptions marks the moment at which they become conscious. There is no single location from where my decisions are sent out. Instead, the many different parts of the brain just get on with their own jobs, communicating with each other whenever necessary, and with no central control. What, then, could correspond to the theatre of consciousness?

It is no good, adds Dennett, to shift from thinking of the theatre as an actual place, to thinking of it as some kind of distributed process, or widespread neural network. The principle remains the same and is still wrong. There simply is no place or process or anything else that corresponds to the conscious bit of the brain’s activities, leaving



all the rest unconscious. There is no sense in which the input is brought together to be displayed 'in consciousness' for someone to see or hear, and no little person inside who acts on what they see. The brain is not organized that way, and it wouldn't work if it were. Somehow we have to understand how this feeling of being a conscious self having a stream of experiences comes about in a brain that really has no inner theatre, no show, and no audience.

Dennett coined the term 'Cartesian materialist' to describe those scientists who claim to reject dualism but still believe in the Cartesian theatre. Note that both these terms, Cartesian theatre and Cartesian materialism, are Dennett's and not Descartes'. Few, if any, scientists admit to being Cartesian materialists. Yet, as we shall see, the vast majority assume something like a stream of consciousness, or treat the mind as an inner theatre. They may, of course, be right, and if they are, then the task of a science of consciousness is to explain what that metaphorical theatre corresponds to in the brain and how it works. But I rather doubt that they are. Exploring a little more about how the brain works may help us to see why.

## Chapter 2

# The human brain

## The unity of consciousness

The human brain is said to be the most complex object in the known universe. Relative to body weight, human brains are larger than those of any other species, and by a long way. They are about three times larger than you would expect by comparing them with those of our closest relatives, the other great apes. A human brain weighs nearly one and a half kilograms and consists of over a billion neurons (nerve cells), with many billions of interconnections. Out of all these connections come our extraordinary abilities: perception, learning, memory, reasoning, language, and – somehow or another – consciousness.

We know that the brain is intimately involved in consciousness because changes in the brain cause changes in consciousness. For example, drugs that affect brain function also affect subjective experiences; stimulation of small areas of the brain can induce specific experiences such as hallucinations, physical sensations, or emotional reactions; and damage to the brain can drastically affect a person's state of consciousness. This much we know for sure, but what remains a mystery is why we should be conscious at all.

In some ways the brain does not seem to be designed the right way to produce the kind of consciousness we have. Most