

# MINDFULNESS, LEARNING AND THE BRAIN<sup>1</sup>

Guy Claxton

*University of Bristol, UK*

**ABSTRACT:** I have tried to sketch an approach to the complex phenomena that go by the name of ‘mindfulness’ that both does justice to this complexity and depth, and also offers a way of thinking about mindfulness in evolutionary, ecosocial and neural terms: terms that enable us to ask questions like: where did mindfulness come from? What kind of consciousness is it? What was it for, before it was co-opted by spiritual and therapeutic kinds of discourse and practice? And how do brains do it? In essence, I am suggesting that human brains seem to have developed, for good evolutionary reasons, a degree of facility with imaginative empathy and as-if identification; and that mindfulness capitalises on this to create what is probably a uniquely human form of learning—or rather unlearning.

**KEY WORDS:** mindfulness; learning; consciousness; brain.

*The highest activities of consciousness have their origins in physical occurrences in the brain, just as the loveliest melodies are not too sublime to be expressed by notes.*

This paper looks at mindfulness through the lenses of cognitive neuroscience. To do this, I need to start by laying out a series of propositions that are just plain obvious to me. They form the pretheoretical framework for the discussion that follows. They need spelling out because I have learned that, though they are pretty

---

<sup>1</sup>This is a speculative piece. Research shows that one of the most important enabling conditions for creativity of any sort is to have at least one person in your life who doesn't think you're nuts. I am lucky enough to have had two, for this exploration, who were not only supportive but actively incisive and helpful. They are psychologist Susan Blackmore, and Buddhist scholar and teacher Stephen Batchelor. Heartfelt thanks to them both.

Author for correspondence to Guy Claxton, Graduate School of Education, University of Bristol, 35 Berkeley Sq., Bristol, BS1 8JA UK; e-mail: Guy.Claxton@bristol.ac.uk.

unexceptionable in scientific circles, they are not accepted as self-evident by everyone. People who hold different views often do so vigorously. For example, some people dispute that consciousness is a phenomenon exclusively associated with brains of a certain kind. Some of them believe that consciousness is the primary stuff of the whole universe, rather like 'energy' in physics. It's not just that rocks 'have' consciousness; they are, in some way, *made* of consciousness. We may need a word for such primordial stuff, but I think it's needlessly confusing if we call it 'consciousness.' I've learned that it can be a waste of time to try to argue about these propositions. An ocean of philosophical ink has been spilt trying to defend and justify these assumptions, and trying to undermine other people's. My preference is to just build the platform, and then explore what can be seen from the vantage point it creates. So I'll just state them, to be clear, and then work within them. Ultimately the proof of the pudding is in whether this presuppositional framework leads in interesting directions. We will see. Here are my axioms.

First, as I've already said, consciousness is produced by and exclusively associated with certain kinds of nervous tissue: complex embodied relational *brains* (CERBs). When I go to sleep, it's in my brain that changes occur as I lose consciousness, and again as I become conscious again as I start to dream. Drugs and accidents that happen to people's heads have effects on consciousness that are more direct and dramatic than traumas to toes or kidneys. I don't know if rocks feel pain, so it seems perverse to suppose they do. I know I do. A rock falling on my toe generates pain. I feel it in the toe. But I *stop* feeling it if someone does something to my brain, like puts me under a general anaesthetic. This effect is much faster and more drastic than any toe-therapy ever is. (I *know* people can feel pain in bits of the body that aren't there any more, so-called 'phantom limbs.' We'll come to that.)

Second, consciousness is a property of *embodied* brains. Philosophers talk about what it would be like to be a disembodied brain in a vat, but they can't possibly know. Brains come with bodies, evolved with bodies, and are an essential element of the complicated systems that bodies are. Brains are so deeply embedded in and continuously affected by what is going on in the body that it makes no sense to treat them as if they were separable. After all, a lot of consciousness, like the pain in my toe, *feels* as if it is happening in the body, and that may turn out to be important in

understanding consciousness, and even mindfulness. The Buddhist *satipatthana sutta* pays enormous attention to mindfulness of the body, and the Buddha, though no neuroscientist, may have had his reasons.

Third, it is probably the case that consciousness is associated with brains that are essentially *relational* and *ecological*—embedded in and continuously affected by both their social and physical milieux. Brains are built relational—for example, they automatically prime themselves to copy what respected others around them are doing—and they automatically build models of how those around them look, sound and react. This capacity for *intersubjectivity* may well be crucial for consciousness of all kinds, including mindfulness.

Fourth, the *vocabulary of consciousness* is hopeless confused. People use the words in different ways. By conscious I shall mean ‘having phenomenal experience’: the sensory world, images, memories, thoughts, feelings, and so on. It is what you lose when you fall asleep and regain when you wake up. I shall use ‘aware’ to mean the same thing. Sometimes I want to use the phrase ‘unconsciously aware,’ to refer to things that I must have registered, and which are affecting me behind the scenes, but which are not themselves conscious. But some people balk at that, so I shall try to use the word ‘registered’ for such events. ‘Unconscious’ and ‘pre-conscious’ I shall use to refer to neural events that are not themselves conscious. ‘Self-conscious’ and ‘mindful’ I’ll explain later. None of these terms necessarily implicate language, by the way. People who talk as if language were somehow essential to consciousness either haven’t looked out of their eyes recently, or are using ‘consciousness’ in a very different way.<sup>2</sup>

Fifth, we need to clear up what ‘mind’ means. It is at least three-ways ambiguous, and we get into trouble if we don’t keep clear in which sense we are using it. First, ‘mind<sub>1</sub>’ is a verb meaning ‘to care.’ To mind about something is to care about it. A child-minder is not just aware of the child; they are taking care of it. Second, ‘mind<sub>2</sub>’ means consciousness, as in ‘I’ll bear it in mind,’ ‘My mind’s gone blank,’ or ‘Mind out!’ Third, ‘mind<sub>3</sub>’ refers to the hypothetical ‘organ

---

<sup>2</sup>To illustrate the differences of opinion on such issues, I’d like to quote Susan Blackmore’s reaction to this paragraph. ‘I disagree with you... the usual idea of consciousness entails someone who is experiencing a stream of conscious experiences, and this is false; people only think it must be that way because they have language. Without language, consciousness is totally different.’ I agree that the quality of consciousness is significantly conditioned by beliefs that have been conveyed through language; but all I am asserting here is the self-evident truth that phenomenal awareness is not itself *necessarily* linguistic.

of intelligence,' as in 'He's got a brilliant mind' or 'Her mind wasn't up to it.' Two confusions in particular keep recurring. One is the idea that you can have mind<sub>2</sub> or mind<sub>3</sub> without mind<sub>1</sub>. Mind<sub>2</sub>—consciousness—is, I shall assume, always concerned, directly or indirectly, with what concerns us, what matters. And mind<sub>3</sub>—the organ of consciousness—is likewise never truly dispassionate. It's there to help us survive and flourish. (The way the brain functions is always relative to what it happens to be up to that moment.) The other confusion is the idea that mind<sub>2</sub> and mind<sub>3</sub> are the same thing. They aren't. The 'organ of intelligence' is the CERB, the embodied brain, not 'conscious awareness.'

Sixth, consciousness always has *content*. One is conscious *of* things: objects, events, impressions, feelings, memories.... Actually there is no distinction between consciousness and the things. Consciousness is *not* like a lightbulb that can illuminate the objects in a room, but which could still be shining if there was nothing there. That's a bad analogy. Occasionally people have weird experiences when they feel their consciousness is contentless, but that is on a par with the phantom limbs I mentioned above. The fact that it is possible to feel pain as if it were in a nonexistent finger does not 'prove' that consciousness has nothing to do with the body, nor does an experience of Divine Light mean that consciousness is not essentially intentional (that is, *about* things, things that matter). Weird phenomena of consciousness (such as 'contentlessness') are interesting but not very important: what Zen practitioners call *makyo*. (I know this assumption is contentious, and I'll have lost some of you, especially the Hindus, here. Sorry, it can't be helped.)

Seventh, to repeat, mind<sub>3</sub>, the *organ of intelligence* is the complex embodied relational brain, the CERB, not consciousness. Everybody knows that often their best ideas just 'occur' to them, they 'pop into their heads' from 'out of the blue.' We sleep on things and in the morning they are clear. We don't know what we think till we hear what we say. And so on. Yet the Cartesian myth persists that consciousness—and especially conscious deliberation—is where my intelligence operates, as if I were the Managing Director of the mind, sitting in the brightly-lit executive office making important decisions and barking orders to menial glands and muscles. This is another bad analogy.

A better one is a computer. All the processing goes on invisibly. It is done by the central processor. The screen is not (for the purposes of this analogy) the smart bit of my laptop. Every so often, things

appear on the screen that reflect, in a very complicated, nonstraight-forward way, what is going on in the ‘works.’ But there is no person-like agent either looking at the screen, or reacting to it, or telling the computer what to do. A ‘person’ is what we call one terminal of a highly complicated and continually interacting carbon (rather than silicon) based network. It is the swirl of activity in that network—physiological, ecological and social—that is registered, largely without conscious awareness, at any particular human terminal, and which influences how the arms, legs, lips and hormones, as well as the screen, located at that terminal, behave.

Eighth, there is no easy way to say this one, but I have to come clean. Consciousness is ‘epiphenomenal.’ However much it seems like it, it doesn’t *do* anything. It cannot reach back down into the biocomputer and somehow alter the course of what is going on. For reasons which we shall probably never know, certain transient states of a CERB are accompanied by conscious experience. Just as there is a very good reason for evolution to have invented milk, but no apparent reason for it to be white, so there are good reasons why evolution has given *homo sapiens* a CERB—the kinds of organism we are would flounder and quickly die out without one—but no real reason why CERBs create that momentary flickering will o’the wisp called consciousness.

There is no logical way of disproving ‘epiphenomenalism,’ despite many valiant tries. Whatever you might suggest consciousness is a cause of, or what it might be *for*, you could just as well say ‘that is what the-complicated-brain-state-that-underlies-consciousness causes, or is for.’ Note for philosophers: I am not asserting the familiar functionalist belief that consciousness *is the same thing as* the underlying processes. Obviously unconsciously flickering patterns of neurons and the conscious experience of a sunset are not the same thing. I am just saying that when that underlying flickering is of the kind that generates conscious experience, no additional causal power is generated, over and above that of the brain state, by the corollary generation of a conscious state. (Much clever conscious wrangling has been generated by brains that have been programmed by experience to find this conclusion aversive.)

## SO WHAT CAN WE SEE FROM HERE?

That, very roughly, is my pretheoretical framework; the platform from which I am going to launch my exploration of what mindfulness

is, and how it ‘works.’ If you haven’t given up in the face of this ‘obviously’ ‘demeaning’/‘reductionist’/‘nonsensical’ approach, join me up in my observation post, and let’s see what we can see.

The most pressing problem that I seem to have created for myself is this. If there is no ‘me’ who is observing and responding—no ghostly operator sitting at the terminal, outside the CERB network itself, weighing things up, making judgements and instigating courses of action—how come it feels so powerfully as if there is? The first thing to realise is that this is not a paradox at all. That sense of being an efficacious center of operations is simply appearing on the screen, not sitting in the chair in front of it. Let’s call this image, in my case, ‘Screen-guy.’ What commonly appears on the screen of many human biocomputer terminals is an image of someone (‘Screen-man’ or ‘Screen-woman’) *looking at a screen and making decisions and banging keys that initiate actions.*

Screen-guy is tapping on the on-screen keys... and sometimes, he is pleased to see that his key-taps have the effect on the contents of the on-screen screen he is watching that he ‘wanted’ or ‘expected.’ And then CERB (for some reason) makes Screen-guy *think* that he was the instigator. But ‘he’ never is: the association between taps and screen-changes is always mediated by the CERB biocomputer, and the ecosocial network of which it is an indissociable part. Sometimes CERB serves up contingencies that strengthen the illusion of direct cause-and-effect control between Screen-guy’s taps and changes on the screen; and sometimes it serves up things that are surprising or nonsensical or shocking.

If Screen-guy *notices* that there is a much looser relationship between his taps and the display than he believes, he gets scared, because he *thinks*—or rather, CERB makes him think—things are getting out of control. He ‘realises’ he has less control in fact than he ‘thinks’ he needs, if things are going to go at all well. He cannot see that what he thought of as control was just predictions, all along. So he panics. He either contorts himself—frowning his brow, talking to himself, and so on—in a futile attempt to re-establish the (sense of) control he never really had. Or he feels depressed and helpless. Or he distracts himself from the disconcerting evidence, and pretends to himself that all is well, and that there was, at most, a temporary glitch in his ability to manage his affairs.

The apparent necessity for all of these manoeuvres arises only because Screen-guy does not realise that he *is* Screen-guy, and

mistakes himself for a real man in a real seat looking at a real screen with real keys he can hit and real levers he can pull to make things happen. Under this misapprehension, there is no rest for Screen-guy. While 'he' is under the spell of this illusion, he has the unenviable choice either to notice all the disconfirming evidence of his lack of control, and feel insecure or depressed or worryingly ineffectual; or not to notice, and operate in a perceptual world that is badly skewed, holed and attenuated.

Either way, he's in trouble. When he chooses the first option, he feels unsettled, and has to get busy shoring himself up (with the props of property, memories, photo albums and familiar routines that enable the fallibility of prediction, and therefore the illusory nature of control, to be masked). When he chooses the second, he has to edit out of his experience much of what makes life vivid and fun. All forms of waywardness of mind and body—dreams, unworthy (though entertaining) thoughts, persistent tunes, even an inexplicably pounding heart—become potential threats to the frail illusion of control, and thus need to be denied or denigrated. All this—each Screen-man and Screen-woman's individually constructed portfolio of unnecessary sufferings and futile responses—is what Buddhists call *dhukkha*.

## MINDFULNESS TO THE RESCUE

Of the two pseudo-solutions to the pseudo-problem of *dhukkha*, anxiety is preferable to denial, because the felt discomfort of Screen-guy reflects a CERB-level recognition that all is not well. The more successful a person is at denial and distraction, shrinking their lifestyle and skewing their perception to edit out the discomfort, the more they become oblivious to the fact that there is anything wrong. If I don't see that, I am effectively strangling my own life—and will continue to do so until some inescapable reminder of my powerlessness comes along, as it eventually does. (The Buddha's trinity of wake-up calls—sickness, old age and disability, and impending death—will penetrate the sleep of delusion in the end, if a redundancy, a car crash or a junkie child doesn't do it first.) With the direct experience of anxiety, Screen-guy may be looking for a solution in the wrong place, but at least he is looking. The trouble, however, is that perception tinged (or saturated) with such anxiety may be unable to fix its gaze on the 'problem' long enough to see clearly what the problem really is. Whenever Screen-guy catches sight of the corner of the problem, habit or guilt or despair

may lead to him flinching away before he has had a good look. So I simply don't get to the point where I see how much of my trouble is home-made.

The first function of 'mindfulness' is simply to strengthen the courage and skill to look: to look *into* the discomfort rather than away from it. The courage and determination may arise for a range of reasons. Perhaps some painful or perplexing issue has forced itself on me so strongly that I am willing to try anything—even giving up old habits of aversion and distraction. Chronic pain or a mid-life crisis will sometimes do it. Perhaps I meet someone who seems to have stared into life and become less, not more, *dhukkha*-ridden as a result. Could be a monk; could be a person at a party, whose being is registered by, and thus affects, my CERB, so that Screen-guy experiences a surge of bravery and hope. Or perhaps, as a result of some happy accident, I find myself trapped in a weekend retreat, and too stubborn or too embarrassed to leave.

For whatever reason, Screen-guy starts to experience himself differently. He starts to notice what is on *his* screen more accurately and honestly, and to begin to appreciate the nature and scale of the problem more fully. He begins to hear clearly, maybe for the first time, the incessant self-serving, self-protective chatter of the voice in his head. He notices more forcibly the small flickers of fear or greed or jealousy that continually derail his peace of mind. He sees how, in a second, his mind<sub>2</sub> has turned a twinge in his hip into the fear of a chronic disability. And he begins to observe the attendant evasions and cover-ups he may have habitually edited out for years. This kind of mindfulness training does two things. It strengthens the ability to hold one's gaze, and to observe the minutiae of one's life, in all their detail. And at the same time it shifts the reaction to these observations from horror, through shame, and on to amusement and eventually greater acceptance. Your Screen-man or woman comes to observe (assuming your CERB is not too far down the slippery slope of self-destruction)<sup>3</sup> that all kinds of wild and wacky thoughts and feelings come and go without you actually going mad or running amok. And cumulatively those observations lead you to ease up on the futile attempt to control the uncontrollable.

As the scum on the surface of life becomes both more obvious and less aversive, so a more basic problem begins to come into focus: the

---

<sup>3</sup>People who suffer from serious ('clinical') mental and emotional problems tend to be those who feel they have to fight most desperately against their own waywardness. See Bentall (2004).

problem that mind<sub>2</sub>—the conscious circus that Screen-guy fondly imagined he was the ringmaster of—turns out to be far more deeply wayward than he thought. Mind<sub>2</sub> very definitely has a mind of its own. I see, infuriatingly, that I cannot hold my attention on the rise and fall of my breath for two seconds before it slips the leash and is running out of my control, scenting lunch, remembering an old flame, rewriting yesterday's galling phone conversation, whatever. The illusion that 'I' am in charge is rudely shattered, time and again, and the assumption that 'identity' entails 'control' has to be challenged increasingly head-on.

This is where mindfulness comes to my aid for the second time, helping to make the problem more precise. For Screen-guy's 'sense of self' is composed of a variety of different components of 'self' that are bundled together more loosely than I had imagined. Mindfulness helps to pull them apart, and let me see where 'I' stand with respect to the different elements. There is Me as a Body (I do have a degree of control over *my* toes which I do not have over yours). There is Me as a Center of Experiences; Me as a Center of Actions; Me as a Collection of Traits; Me as a String of Memories; Me as the Owner of Things; and so on. We say I Can, I Did, I Have, I Am, I Was, I Decided, I See... and the language leads me to believe that I and Me refer to a real hub—a single center of a circle from which all these different attributes radiate. But, as Derek Parfit has argued, there is no such hub: only a collection of self-ingredients (Parfit, 1987). Just as when we say 'it is raining' there isn't any 'it' (apart from the rain), so when we say 'I made up my mind,' there isn't any 'I' apart from the appearance of a decision. And just as through words and thoughts we can conjure an effective ghost in the corner of the bedroom, so through words and thoughts we come to feel that we can almost see—indeed almost are—the ghost in the biological machine. The less we scrutinise this conglomerate, the more we feel it to be real and unitary; and the more we feel that, the more powerful is the need to live up to its standards and within its constraints. The more earnestly we have to shore up the feeling of Being In Control. As the senses of self adhere and stick together, so this coagulation of the self makes more trouble.

The penetrating oil of mindfulness can begin to free up and tease apart an assemblage of self-ideas that has become seized. For example, the sense of myself as *observer*—the 'I' who sees, hears, feels, experiences—is not the same 'self' as the sense of myself as the

*controller* or *instigator*—the ‘I’ who *plans* things and *does* things. Mindfulness practice can pull these apart, and enable me to inhabit the vantage point of the observer, for example, whilst disidentifying with the vantage point of the instigator.

In terms of my previous analogy, what mindfulness practice can do is get the CERB bio-terminal to create a slightly more complicated screen, on which there is a second Screen-guy, ‘Screen-guy<sub>2</sub>’, who is watching Screen-guy<sub>1</sub> as he taps on the on-screen keys and watches the on-screen monitor. The biocomputer now shifts the vantage point, generating the sense that ‘who I am’ is Screen-guy<sub>2</sub> rather than Screen-guy<sub>1</sub>: the Watcher rather than the Tapper, and this shift, when it happens, helps to prise my sense of ‘who I am’ away from the identification with the Tapper. If ‘who I am’ can be Screen-guy<sub>2</sub>, the Watcher, observing Screen-guy<sub>1</sub>, the Tapper, then, from that vantage point, ‘I’ can be *interested* in Screen-guy<sub>1</sub>’s activities, but much less *invested* in them. And if I can shift my vantage point in this way, it follows that I cannot really *be* the Tapper. In fact, if I follow the logic, I can’t *really* be either of them. Learning to toggle between *immersion* in my mind’s activity, and calmer *observation* of those same activities, enables me to begin to see that these are *both* screen-identities; not ‘real’ ones. The feeling of ‘over-identification’ with *any* screen-man can be diluted. That can be useful for people who are at risk of jumping out of the frying-pan of frantic Tapping into the fire of compulsive Watching—a not-infrequent mistake amongst meditators, who mistakenly think that their error was to identify with the *wrong* Screen-person, rather than with any of them.

## THE NEURAL BASIS OF IDENTIFICATION AND PERSPECTIVE-TAKING

This business of being able to shift vantage points is crucial to understanding how mindfulness works, so it might be helpful to speculate a little more about why we have it, and how we develop it. Evolutionarily, Nicholas Humphrey, Richard Dawkins and others have argued that, as human tribes became larger and more stable, so learning to predict the interests and habits of our neighbours became a highly desirable ability, to facilitate both competition and cooperation. To do this, we had to learn to watch them, and to distil their behaviour over time into neural models of their idiosyncratic

portfolios of preferences, fears and dispositions that we could then 'run' in order to predict their likely reactions to events, and therefore potentially assist them or outwit them. This ability to infer and capitalise on these working models of other people's habits became a universal habit in its own right (see e.g. Humphrey, 1986).

Nothing new in principle was involved in brains making these kinds of models. They had already been good at such observational and predictive learning for a long time. Such pattern detection was, in large measure, what they evolved to do. Before<sup>4</sup> they were building models of people, they were busy extracting all sorts of useful, predictive regularities and building models of the kind of world they lived in—the kinds of objects, scenarios and so on that typified the eco-niches in which they found themselves.

The process of building and using such predictive models of the world did not need to be conscious or deliberate. Both learning and use could, in principle, have proceeded quite intuitively. An articulate scholar, observing these processes from the outside, might describe it as people 'reading each others' minds,' or ascribing hopes, desires and intentions to them. But it need not have seemed anything like that to people themselves. Tennis players, receiving serve, adjust their footwork and swing in an instant, on the basis of barely conscious observations of the servers' stance and toss, filtered through knowledge derived from precious encounters (or from watching videos of their earlier matches). There is much in their own minds, and in the sensory information they are making use of, that does not itself become conscious. Equally, parents and children learn to read and accommodate to each other without thinking, or even noticing what they are doing.

However, conscious awareness entered the picture somewhere, and there is reason to suppose that the brain's generation of consciousness is associated with occasions of surprise and/or satisfaction relative to these models. On the basis of my predictive models, my brain generates a background map of expectations: what is likely to happen, what I fear might happen, and what is currently desirable. These expectations weave together into a moment-by-moment picture of what is stable, predictable, problematic or advantageous, against which actual events, as they unfold, are compared. What matters most is when something desired (or feared) and expected turns up (a 'satisfaction' or a known 'threat'); or when something unexpected

---

<sup>4</sup>Both evolutionarily and developmentally.

happens. It is only against the backdrop of such a map of likelihood that Matches that Matter, Surprises and Disappointments can stand out—and it is precisely those occurrences that seem to attract conscious awareness. As such events are attended to and dealt with, so they become absorbed into the continually updated ‘briefing paper,’ underwriting a revised set of possible satisfactions, threats, surprises and disappointments, that gets us ready to act and react in the smartest possible way. We don’t see What’s There, in some kind of neutral way; we see events relative to, and filtered through, a set of Affordances for Advantageous Action.

It seems as if human brains have developed the ability to turn this background orienting knowledge into two different, complementary kinds of map, that neuroscientist Jeffrey Gray called the *egocentric* and the *allocentric* maps (Gray, 2004). They differ in whether ‘I’ am represented as the ‘origin,’ the invisible center, of the map, or as an object, with its own location and characteristics, on the map itself. In the egocentric map, the world radiates out from ‘me,’ as it were, and everything else is placed, physically and psychologically, relative to my eyes, my mouth, my hands, my legs, and also to my idiosyncratic portfolio of interests, desires and neuroses. I need this map, obviously, for grasping, chasing, fleeing, eating and so on. But the brain has also developed the advantageous trick of being able to de-center from my own physical (and psychological) point of view, and adopt a view that is not so strongly ‘captured’ by my current position. From this more neutral vantage point, less strongly skewed by whatever happens to be my currently active motivational agenda, for example, I can see more possibilities. The egocentric map gives me my ‘first-person view’ of things, we might say, while the allocentric map gives me a ‘third-person perspective.’ These two complementary maps are actually represented in different locations and pathways in the brain, and we are able, it seems, to toggle between them, using the strengths of each to offset the limitations of the other.

Thus my brain creates different ‘worlds,’ depending on what set of assumptions and desires, and what perspective, it is currently adopting. But in order to do this, and especially to adopt the allocentric perspective, it needs a flexible model of *me*. Just as I observe other people, and distil out generalisations about their habits and preferences, so, presumably, I can observe and remember my own behaviour over time, and the brain can use these records, in the same way as it does for others, to distil generalisations about my own traits. My

brain creates a rudimentary model of the kind of person its owner is, containing my characteristic (but continually changing) powers and limitations, hopes and fears, likes and dislikes. This model of Me enables me to be surprised, threatened or satisfied by my own thoughts, feelings and actions. And the more accurate this implicit self-knowledge is, the better I can calibrate my actions, taking account not just of what I think *you* might do, but also of *my own* state and my own dispositions. When I am, as it were, ‘inside’ this model, it forms the psychological origin of my egocentric worldview. When I am outside it, capable of looking at it from an external vantage point, it becomes a ‘me’ that can be observed and thought about from the point of view of a different ‘I.’

This third-person perspective may be some kind of generalised other—a ‘bird’s-eye-view’—or it might be that of a particular other. Brains can, it seems, ‘get inside’ the models they have created of other people, and then look outward through their eyes, and with their motivational agendas. We all develop this ability more or less well, of course, but the ability for imaginative identification—for conscious empathy—seems a fundamental endowment of brains (e.g. Thompson, 2001). We have developed the ability to ‘run’ these models internally, and see how a person would behave under different circumstances. In these improvised home-movies, we are able to make explicit those aspects of other people that are implicit in the neural model, but which had not previously been drawn out. So while inside their ‘shoes,’ I might still retain the interest of my own perspective, trying to uncover how they might ‘jump,’ so I can the better coordinate or compete. Alternatively, while I am fully inside their minds, I might look back at myself, seeing me through their eyes. Inside my neural model of them—if I am as sophisticated as most of us are—I find I have constructed (more or less crudely; more or less accurately) their model of me.

We may have to be careful that, when we think we are being ‘mindful’ of ourselves, in a neutral way, we are not inadvertently choosing a third-person perspective to inhabit that is imbued with its own values, biases and judgements. Indeed, many people find it hard, when they begin to practice a form of mindfulness meditation, to be merely ‘self-aware’ without lapsing into the compulsively judgemental form of self-awareness that we call ‘self-consciousness,’ where the vantage point from which we are looking back at ourselves is actually that of a critical parent or teacher. One of the virtues of the

nonjudgmental therapist or ‘guide’ is precisely to enable our brains to build a new ‘model of another’ from which to look back at ourselves—perhaps the first one in our lives to contain relatively few (overt or covert) demands, desires and expectations of its own. To be seen clearly, yet without judgment, by another, is to have the opportunity to learn to construct that observant, compassionate vantage point for ourselves.

## FINAL THOUGHTS

I have tried to sketch an approach to the complex phenomena that go by the name of ‘mindfulness’ that both does justice to this complexity and depth, and also offers a way of thinking about mindfulness in evolutionary, ecosocial and neural terms: terms that enable us to ask questions like: where did mindfulness come from; what kind of consciousness is it; what was it for, before it was co-opted by spiritual and therapeutic kinds of discourse and practice; and how do brains do it? In essence, I am suggesting that human brains seem to have developed, for good evolutionary reasons, a degree of facility with imaginative empathy and as-if identification; and that mindfulness capitalises on this to create what is probably a uniquely human form of learning—or rather unlearning. Being able to shift from the perspective of Screen-guy<sub>1</sub>—the Tapper—to that of Screen-guy<sub>2</sub>—the Watcher—is simply an extension of something that evolution originally developed to allow us to get inside each others’ minds. Our own ‘self’ is not a unitary construct; we are a ‘community of selves.’ And thus we can use our facility for imaginative identification intra- as well as inter-psychically.

## REFERENCES

- Bentall, R. (2004). *Madness explained*. London: Penguin.
- Gray, J. (2004). *Consciousness: Creeping up on the hard problem*. Oxford: Oxford University Press.
- Humphrey, N. (1986). *The inner eye*. London: Faber & Faber.
- Parfit, D. (1987). Divided minds and the nature of persons. In C. Blakemore & S. Greenfield (Eds.), *Mindwaves: Thoughts on intelligence, identity and consciousness*, Oxford: Blackwell.
- Thompson, E. (2001). Empathy and consciousness. *Journal of Consciousness Studies*, 8, 1–32.