Mark Solms[1]

Précis of The Hidden Spring: A Journey to the Source of Consciousness

Here is a 15-point précis of my book, *The Hidden Spring*. The commentaries which follow it discuss the book itself, not only this précis. The précis is provided for the benefit of readers who are not familiar with the book. The publication of *The Hidden Spring* was preceded by a preliminary communication in this journal (Solms and Friston, 2018), which readers seeking more detail about the technical arguments, only briefly alluded to here, may fruitfully consult.

1.

The great nineteenth-century physiologist Johannes Müller believed that animate organisms 'contain some non-physical element or are governed by different principles than are inanimate things' (Bechtel and Richardson, 1998). His pupils, Helmholtz, Brücke, du Bois-Reymond, Ludwig, and others, disagreed; they argued that 'no other forces than the common physical and chemical ones are active within the organism' (Du Bois-Reymond, 1842/1918). Brücke's pupil, Sigmund Freud, tried to establish a natural science of the mind on this basis, in which mental life could be reduced to 'quantitatively determinate states of specifiable material particles' (Freud, 1895). He failed in his project, lacking the methods, and abandoned it in 1896.

Correspondence:
Email: mark.solms@uct.ac.za

[1] Director of Neuropsychology, Neuroscience Institute, University of Cape Town, South Africa.

2.

A century later, Francis Crick (1994) declared that 'you, your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behaviour of a vast assembly of nerve cells and their associated molecules'. He exhorted us to try again to discover the physical basis of consciousness, and he attempted to do so himself. Like many others who followed in his footsteps, Crick used vision as his model example. This was an unfortunate choice, as we shall see.

3.

In response, David Chalmers (1995) argued that our search for the physiological mechanism of consciousness was an 'easy' problem. In line with Frank Jackson's (1982) 'knowledge argument', he pointed out that the functional mechanism of vision does not explain how and why there is something it is like to see. For Chalmers (like Thomas Nagel before him), the 'hard' problem revolves around this something-it-is-like-ness of experience: 'An organism has conscious mental states if and only if there is something that it is like to *be* that organism — something it is like *for* the organism' (Nagel, 1974). The hard problem, accordingly, is this: why and how do objective neurophysiological functions produce the subjective qualia of experience?

4.

To ask how objective things produce subjective things is to speak loosely, and it risks making the hard problem appear harder than it is. Objectivity and subjectivity are observational perspectives, not causes and effects. Neurophysiological events can no more produce psychological events than lightning can produce thunder. They are dual manifestations of a single underlying process. The cause of both lightning and thunder is electrical discharge, the lawful action of which explains them both. Physiological and psychological phenomena must likewise be reduced to unitary causes, not to one another. This is merely a restatement of a well-known position on the mind–body problem: that of dual-aspect monism.

5.

Crick's model example of consciousness was unfortunate because vision (even cortical vision) is not an intrinsically conscious function.

The same applies to most other perceptual and cognitive functions; it is readily possible to perceive without awareness of what is perceived and to learn without awareness of what is learned (Kihlstrom, 1996). Chalmers therefore correctly pointed out that the functional mechanism of vision does not explain what it is like to see. On this basis, he reasonably asked: 'Why is the performance of these functions accompanied by experience? Why doesn't all this information-processing go on "in the dark," free of any inner feel?' (Chalmers, 1995). Science's failure to answer this profound question raises the possibility that consciousness is not reducible to the laws of physics: 'We know that a theory of consciousness requires the addition of something fundamental to our ontology, as everything in physical theory is compatible with the absence of consciousness' (*ibid*.). So, Chalmers claimed that experience might be added to the fundamental properties of the physical universe. He speculated further that this property might be the subjective aspect of information: the being of information.

6.

However, as John Kihlstrom's review of the experimental evidence (just cited) reminded us, not all information processing feels like something, not even in the human brain. Chalmers' question about information processing going on in the dark 'free of any inner feel' may legitimately be asked of almost all cognitive functions, for the reasons already stated: it is readily possible to perceive without awareness of what is perceived and to learn without awareness of what is learned. But does the same apply to *affective* functions? How can you have a feeling without feeling it? Can neuroscience explain the functional mechanism of affect without explaining why and how it causes the organism to experience something? Surely the function of feeling positively predicts that it should feel like something, at a minimum, that it should be pleasurable or unpleasurable.

7.

This is not a semantic point. It is of the utmost interest to observe that cortical functioning is accompanied by consciousness if and only if it is 'enabled' by the reticular activating system of the upper brainstem. Damage to just two cubic millimetres of this primitive tissue reliably obliterates consciousness as a whole (Fischer *et al.*, 2016). Since, unlike the cortex, the functioning of this brain region is intrinsically conscious — indeed the generating of consciousness is *the* function of the reticular activating system — is it not a more appropriate place for us to seek the functional mechanism of consciousness? The reason we have not done so before is because the reticular activating system was historically assumed to generate only the quantitative ‘level’ of consciousness (‘consciousness as the waking state’, Zeman, 2001) and not its qualitative ‘contents’ (‘consciousness as experience’, *ibid*.).

That widely held assumption is no longer sustainable. The form of consciousness that is generated by the reticular activating system (together with the periaqueductal grey, PAG) has a qualitative content of its own. This is affect. The converging lines of evidence supporting this important conclusion are overwhelming, and they are based on a wide range of research findings. Here are some representative examples. Children born without a cortex, but in whom the upper brainstem is intact, display appropriate affective responses to adequate stimuli (Shewmon, Holmes and Byrne, 1999; Merker, 2007). Psychiatrically healthy patients in whom the PAG or individual reticular activating nuclei are stimulated electrically report intense affective states for the duration of the stimulus (e.g. sudden onset of suicidal depression; Blomstedt *et al*., 2008). Positron emission tomography of research participants experiencing intense affective states — such as fear, sadness, anger, and joy — reveals maximal activation of the upper brainstem during these states, and it also shows cortical *de*activation (Damasio *et al*., 2000). Pharmacological manipulation of the neuromodulators that are sourced in the reticular activating system have substantial effects upon affective states — such as depression, anxiety, and psychosis. The source cells for serotonin, noradrenaline and dopamine — the targets of mainstream antidepressant, anti-anxiety, and anti-psychotic medications, respectively — are in the raphe nuclei, locus coeruleus complex, and ventral tegmental area, all of which are located in the upper brainstem.

By contrast with these findings, which converge on the conclusion that affect is generated in the brainstem, predictions from cortical theories which claim that it is generated in the insula (Craig, 2009) or the prefrontal lobes (e.g. LeDoux and Brown, 2017) are roundly disconfirmed by clinical studies which show that affect is both subjectively and objectively preserved — indeed increased — in neurological patients with total obliteration of these cortical structures (Damasio, Damasio and Tranel, 2013; Solms, 2021). The same applies to decorticate mammals. Neonatally decorticate rats, for example, stand, rear, climb, hang from bars, and sleep with normal postures. They groom, play, swim, eat, and defend themselves. Either sex is capable of mating successfully when paired with normal cage mates. When they grow up, the females show the essentials of maternal behaviour, which, though deficient in some respects, allow them to raise pups to maturity (see Merker, 2007, for review).

Since it is generally accepted that cortical consciousness is contingent upon brainstem consciousness, and since it is now evident that brainstem consciousness is affective, we must conclude that *affect is the foundational form of consciousness*. Sentient subjectivity (in its elementary form) is literally constituted by affect.

8.

I have argued that the functional mechanism of vision does not explain how and why there is something it is like to see, but that the same does not apply to affect. What is the functional mechanism of affect? It is a form of homeostasis. Using the analogy of lightning and thunder, homeostasis is the equivalent of electrical discharge: it explains how affective modulation looks (physiologically) and how it feels (psychologically). Affect valences biological needs hedonically, so that increasing and decreasing deviations from settling points (increasing and decreasing homeostatic 'error signals') are felt as unpleasure and pleasure, respectively. Deviations from biologically preferred states — i.e. the viable states of each species — are registered by the organism, for the organism, as *needs* (*cf.* Nagel, 1974, above). Needs are 'a measure of the demand made upon the mind for work in consequence of its connection with the body' (Freud, 1915). Each category of bodily need — of which there is a great variety — has an affective quality of its own, and each triggers an action that is predicted to return the organism to its viable bounds. The same applies to emotional needs, which are no less homeostatic than bodily ones (e.g. fleeing corrects 'danger' signals [fear], attack corrects 'frustration' signals [rage], crying and searching corrects 'separation' signals [panic]). These error-correcting actions take the form initially of innate reflexes and instincts (i.e. predictions that evolved through natural selection), which are automatically executed and therefore don't necessarily entail consciousness. However, the capacity for *feeling* by the organism of fluctuations in its own need states enables it to make here-and-now *choices.* This facilitates survival in *uncertain* contexts. Unpredictable situations are far from rare in nature. The evolved capacity for choice therefore bestows enormous adaptive advantages; it enables the organism to *feel its way* through life's problems. Choices must be rooted in a value system. The value system of all living things stipulates that it is 'good' to survive — and reproduce — and 'bad' not to do so. Accordingly, biological goodness is felt as pleasure (decreasing uncertainty) and badness as unpleasure (increasing uncertainty). This seems to be the biological function of feeling.

Take respiratory control, for example. The homeostatic regulation of your blood gases — i.e. maintaining your preferred oxygen to carbon dioxide ratio — becomes conscious when you do not have a ready-made prediction to maintain the physiologically viable range. In your rush to escape from a carbon-dioxide-filled room, for instance, how do you know which way to turn? You have never been in this situation before (in any burning building, let alone this specific one) so you cannot possibly predict what to do. Now you must choose whether to go this way or that, up or down, etc. So, you make your choices by feeling your way through the problem: the feeling of suffocation alarm (the error signal) waxes or wanes, depending upon whether you are going the right way or wrong — that is, depending upon whether the availability of oxygen increases or decreases as you move. This reveals the biological function not only of feeling itself but also of voluntary action (as opposed to automatized behaviour). Consciousness has everything to do with the capacity for voluntary action (in fact, voluntary action may be used as an objective marker of sentient intentionality).

9.

Not all organisms — let alone all homeostats — are equipped with the mechanisms just described. This is partly because valence is a continuous variable (e.g. an error signal of 8/10 is greater than one of 5/10). This is easily regulated automatically, in the same way as domestic temperature control systems are. However, since complex organisms have *multiple* needs, each of which must be met in its own right, and, indeed, on a context-dependent basis, they cannot be reduced to a common denominator. For example, 5/10 of fear plus 9/10 of sleepiness cannot be treated as 14/20 of total need, where sleeping alone might be the appropriate response. Likewise, fear trumps sleepiness in some contexts but not in others. So, the needs of complex organisms like ourselves must be coded as categorical variables. This enables the resolution of them to be computed in a context-dependent fashion (i.e. it enables them to be compartmentalized, which is the standard statistical-mechanical solution to excessively complex calculations — *cf.* the 'combinatorial explosion').

Categorical variables are distinguished qualitatively, not quantitatively. Hence, thirst feels different from sleepiness feels different from separation distress feels different from fear, etc. The needs of complex organisms which can act differentially, in flexible ways, in variable contexts, are therefore 'colour-coded' or 'flavoured'. This provides at least one mechanistic imperative for qualia.

10.

But why are qualia felt *subjectively*, by the organism and for the organism? Homeostasis evolved naturally through the fundamental physical process of self-organization (Ashby, 1947). The critical property of self-organizing systems is that they continue to exist as systems by sequestering themselves from the entropic forces that surround them, through the formation of Markov blankets (Friston, 2013). Such blankets distinguish the system from the not-system, and thereby maintain its structural and functional integrity, by registering entropic threats to its survival as 'sensory' states of their blankets, and by responding proactively to such threats via the 'active' states of their blankets. The sensory-motor activity of such systems is proto-intentional, in the sense that the actions have an aim and a purpose: they must always minimize the entropy of the system. The system's entropy is quantified via its variational free energy (*ibid.*). 'Free energy' is the portion of the system's total energy that is not being put to effective work; it is unused energy. An information processing system's free energy is calculated from the average difference between the sensory outcomes that were predicted by its generative model to flow from its actions, and the sensory outcomes that actually flow from those actions. In other words, it is a measure of the *efficiency* of the system's predictive model. Every complex self-organizing system must develop a predictive model of the consequences of its actions in the world, based on past experience; how else can it minimize its free energy?

Notably, however, the actual state of the external world is radically hidden from Markov-blanketed systems; it can be represented only vicariously via states of the blanket. This fact bestows a *point of view* upon the system. The attribution of subjectivity to self-organizing systems is justified by their selfhood. A system's point of view is available only to the system, for the system. One system can never sense the internal states of another system — as internal states — not only because the external world is hidden from it but also because the internal states of other systems are internal to those systems alone. This obviously does not imply that all self-organizing systems are sentient. Nevertheless, it makes it meaningful to speak of the viewpoint (the subjectivity) of such systems.

This is important because Claude Shannon's (1948) conception of 'information', which introduced the concept into physics, is a mathematical theory of *communication*. Communication assumes a source and a receiver of information: an asker of questions the answers to which are registered by the questioner in the yes/no format of bits (i.e. 1s vs. 0s; *cf.* Wheeler, 1990). Most information processing models in cognitive science seem to lack question-askers. The same applies to Chalmers' (1995) notion that *all* information has an internal aspect to it. Such theories beg the question: where is the receiver — the subject — of the information processing? Self-organizing systems, by contrast, are obliged to ask questions; their very survival depends upon it. They must chronically ask: 'What will happen to my free energy if I do that?' The answers they receive determine their confidence in the current prediction, in line with the waxing and waning of uncertainty, as described above. This is why I claim in *The Hidden Spring* that not all information processing ('integrated' or otherwise) is conscious; sentience appears to be a property of only some information processing systems with very specific properties, namely those systems that must ask questions of their surrounding world in relation to their existential needs.

11.

A distinction must be drawn between needs and feelings. An organism's needs cannot all be felt at once. This is because it cannot act upon all its needs simultaneously. A selection must be made. This selection occurs mainly at the level of the midbrain PAG, since all needs (all homeostatic error signals in the vertebrate brain) converge there. Here they are prioritized by what Bjorn Merker (2007) calls a 'midbrain selection triangle', a mechanism whereby all the needs converging on the PAG are ranked — as they must be — in relation to current opportunities (displayed in the form of a two-dimensional 'saliency map' in the adjacent superior colliculi). The opportunities just mentioned represent the fluctuating contexts within which needs must be evaluated.

The resultant ‘decision’ triggers an affectively flavoured action program, which then unfolds in an expected context, elaborated over a deep hierarchy of predictions (encoded by the organism’s predictive model, upwards from the brainstem reflexes and subcortical instincts through the long-term non-declarative and then the declarative memory systems, and ultimately the short- and ultra-short-term memory systems of the expanded mammalian forebrain).

The actions that are generated in this way in relation to the prioritized need (the currently felt one) are voluntary, which means they are subject to here-and-now choices rather than to the pre-established algorithms that govern responses to the non-prioritized needs. For example, when the need to escape danger (felt as fear) is prioritized over the need to urinate, you might wet yourself without giving it a second thought. The making of choices is buffered in working memory: a ‘global workspace’ which contextualizes the currently prioritized affect, not at the level of a crude saliency map but rather as ‘a fully articulated, panoramic, three-dimensional world composed of shaped solid objects: the world of our familiar phenomenal experience’ (*ibid.*). The cortex — which generates stabilized ‘mental solids’ — specializes in articulating this expected world, which it ‘holds in mind’ in relation to each currently prioritized need.

The organism’s voluntary choices are made on the basis of fluctuating *precision weighting* (also known physiologically as ‘arousal’ — see Pfaff, 2005 — or ‘post-synaptic gain’, which is modulated mainly by the reticular activating system) of the exteroceptive and proprioceptive error signals that are rendered salient by a prioritized need, with the aim of minimizing uncertainty (maximizing confidence) in its current predictions as to how that need can be met. As previously stated, increasing confidence in a current prediction is ‘good’ and decreasing confidence (i.e. increasing confidence in the attendant error signals) is ‘bad’. This is how consciousness becomes *applied to* cognition. Unconscious cognition, by contrast, proceeds on the basis of monotonous precision weightings of the multiple other action programs that are predicted to meet the non-prioritized needs. Consciousness, all of it, is *felt uncertainty*. The major distinction between affective and cognitive consciousness concerns what this uncertainty is *about* (*cf.* Brentano, 1874). In a nutshell, affective consciousness is about the state of the subject, while cognitive consciousness is about the state of its objects: ‘I feel like this’ vs. ‘I feel like this *about that*’.

It is worth noting, with respect to Crick's research programme, that there can be no objects of consciousness (e.g. visual ones) in the absence of a subject of consciousness. As stated, the subject of consciousness is literally constituted by affect.

12.

The research agenda of the twentieth-century behaviourists discerned the laws of learning (e.g. operant conditioning) which underwrite all cognition. Edward Thorndike (1911) famously concluded: 'If a behaviour is consistently accompanied by rewards it will increase, and if it is consistently accompanied by punishments it will decrease.' This was his Law of Effect. Bizarrely, Thorndike assumed that the rewarding and punishing properties of stimuli reside in the stimuli themselves, rather than in the subject: in the receiver of information — the question-asker — the organism as active agent. This is because the behaviourists ruled out of science the first requirement of any science of mind, namely the requirement for its investigators to sometimes adopt the viewpoint of the system they are investigating. This is a move that is required for the simple reason that the mind is subjective (what else could it be?). The behaviourists thereby ruled the mind out of mental science. They excluded the psyche from psychology. To the extent that cognitive neuroscience still maintains this epistemological stricture, it can do no better.

The hard problem will remain forever unsolved if we continue to exclude subjectivity from the neuroscience of consciousness. Accordingly, in *The Hidden Spring*, I supplement the standard perspective of cognitive neuroscience with that of affective neuroscience (Panksepp, 1998), and I replace the Law of Effect with a Law of Affect: 'If a behaviour is consistently accompanied by pleasure it will increase, and if it is consistently accompanied by unpleasure it will decrease.' To be clear: the waxing and waning of levels of statistical confidence in an action plan (i.e. fluctuations in the system's expected free energy and its constituent precisions — its confidence values) can only ever be felt *by the organism*, for the organism. Moreover, as already explained, the adoption of the organism's perspective is justified precisely by its selfhood.

Due to the Law of Affect, reliably need-satisfying (i.e. pleasure-generating) choices result in long-term adjustments of the organism's generative model. This is 'reconsolidation' (Nader, Schafe and LeDoux, 2000), which may be described as *predictive-work-in-*

progress. It really is nothing more than learning from experience, with the emphasis falling on 'experience'. If conscious affect measures the demands made upon the mind for work (see above), then conscious cognition is the work so demanded: it is the mind experiencing its way through its problems.

The goal of cognitive consciousness over longer timescales seems to be the establishment of ever deeper (i.e. less uncertain, and therefore less likely to demand conscious experience) predictions as to how our needs may be met. This entails increasing systems consolidation — the reduction of predictions from declarative to non-declarative memory systems — which enables a reduction of complexity in the predictive model, which is important to facilitate generalizability.

We living organisms aspire to automaticity — to absolute confidence — but we can never achieve it (of course) in our endlessly unpredictable worlds. To the extent that we fail, we suffer feelings. Since we can never achieve perfect predictive models — perfectly efficient models with zero free energy — our 'default' affect is what Panksepp (1998) called SEEKING — a drive which is mediated by the dopaminergic medial forebrain bundle. This is the source of our ceaseless proactive engagement with uncertainty, with the aim of resolving it in advance. When this affect is prioritized, we feel (at a minimum) a sense of interest in the world. Karl Friston calls it 'epistemic foraging'.

13.

These few points summarize my conception of the causal basis of consciousness — in both of its manifestations, physiological and psychological — i.e. what it looks like and what it feels like. The functional mechanism of consciousness, I am claiming, can be reduced to physical laws, such as Friston's Law, which states: 'All the quantities that can change, i.e. that are part of the system, will change to minimize free energy.' These laws, which are explicated more fully in my book, are no less capable of explaining how and why proactively resisting entropy (i.e. avoiding oblivion) feels like something to the organism, for the organism, than other scientific laws are capable of explaining other natural things. Consciousness is part of nature, and is mathematically tractable. As Galileo said: 'The book of Nature is written in the language of mathematics.'

14.

It is conceivable that the biological control task that consciousness evolved to perform could equally well have been performed by some other (non-conscious) mechanism. However, the same applies to legs: the task of ambulation could be performed equally well by other mechanisms, and indeed it is in many species. Nevertheless, what we ourselves have is legs, because they have worked so well for us. As the great neurologist Jean-Martin Charcot once said: 'Theory is good, but it doesn't stop things from existing.' We must be careful not to set the bar for a scientific explanation of consciousness higher than we have set it for all other problems in biology.

15.

All known conscious systems are alive, but not all living systems are conscious. Likewise, all living systems are self-organizing, but not all self-organizing systems are alive. If the argument laid out here is correct, then, in principle, an artificially conscious self-organizing system can be engineered. When Richard Feynman died, the following statement was found on his blackboard: 'What I cannot create, I do not understand.' The creation of an artificial consciousness is, accordingly, the ultimate test of any claim to have solved the hard problem. To do so would realize the wildest dreams of the Helmholtz school of physiology. However, we must question our motives for attempting to engineer consciousness, accept collective responsibility for the potentially dire consequences, and proceed with extreme caution.

References

Ashby, W. (1947) Principles of the self-organizing dynamic system, *Journal of General Psychology*, **37**, pp. 125–128.

Bechtel, W. & Richardson, R. (1998) Vitalism, in Craig, E. (ed.) *Routledge Encyclopedia of Philosophy*, 9, pp. 639–643, London: Routledge.

Blomstedt, P., Hariz, M., Lees, A., Silberstein, P., Limousin, P., Yelnik, J. & Agid, Y. (2008) Acute severe depression induced by intraoperative stimulation of the substantia nigra: A case report, *Parkinsonism and Related Disorders*, **14**, pp. 253–256.

Brentano, F. (1874) *Psychologie vom empirischen Standpunkte*, Leipzig: Duncker and Humbolt.

Chalmers, D.J. (1995) Facing up to the problem of consciousness, *Journal of Consciousness Studies*, **2** (3), pp. 200–219.

Craig, A.D. (2009) How do you feel — now? The anterior insula and human awareness, *Nature Reviews Neuroscience*, **10**, pp. 59–70.

Crick, F. (1994) *The Astonishing Hypothesis: The Scientific Search for the Soul*, New York: Charles Scribner's Sons.

Damasio, A., Grabowski, T., Bechara, A., Damasio, H., Ponto, L.L.B., Parvizi, J. & Hichwa, R.D. (2000) Subcortical and cortical brain activity during the feeling of self-generated emotions, *Nature Neuroscience*, **3**, pp. 1049–1056.

Damasio, A., Damasio, H. & Tranel, D. (2013) Persistence of feelings and sentience after bilateral damage of the insula, *Cerebral Cortex*, **23**, pp. 833–846.

Du Bois-Reymond, E. (1842/1918) Letter to Hallmann, in Du Bois-Reymond, E. (ed.) *Jugendbriefe von Emil Du Bois-Reymond an Eduard Hallmann, zu seinem hundertsten Geburtstag, dem 7. November 1918*, Berlin: Reimer

Fischer, D., Boes, A., Demertzi, A., Evrard, H., Laureys, S., Edlow, B., Liu, H., Saper, C., Pascual-Leone, A., Fox, M. & Geerling, J. (2016) A human brain network derived from coma-causing brainstem lesions, *Neurology*, **87**, pp. 2427–2434.

Freud, S. (1895) Project for a scientific psychology, *Standard Edition of the Complete Psychological Works of Sigmund Freud*, 1, pp. 283–397, London: Hogarth.

Freud, S. (1915) Instincts and their vicissitudes, *Standard Edition of the Complete Psychological Works of Sigmund Freud*, 14, pp. 117–140, London: Hogarth.

Friston, K. (2013) Life as we know it, *Journal of the Royal Society Interface*, **10**, 20130475. doi: 10.1098/rsif.2013.0475.

Jackson, F. (1982) Epiphenomenal qualia, *Philosophical Quarterly*, **32**, pp. 127–136.

Kihlstrom, J. (1996) Perception without awareness of what is perceived, learning without awareness of what is learned, in Velmans, M. (ed.) *The Science of Consciousness: Psychological, Neuropsychological and Clinical Reviews*, pp. 23–46, London: Routledge.

LeDoux, J. & Brown, R. (2017) A higher-order theory of emotional consciousness, *Proceedings of the National Academy of Science*, **114**, pp. e2016–e2025.

Merker, B. (2007) Consciousness without cerebral cortex: A challenge for neuroscience and medicine, *Behavioral and Brain Sciences*, **30**, pp. 63–68.

Nader, K., Schafe, G. & LeDoux, J. (2000) Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval, *Nature*, **406**, pp. 722–726.

Nagel, T. (1974) What is it like to be a bat?, *Philosophical Review*, **83**, pp. 435–450.

Panksepp, J. (1998) *Affective Neuroscience: The Foundations of Human and Animal Emotions*, New York: Oxford University Press.

Pfaff, D. (2005) *Brain Arousal and Information Theory*, Cambridge, MA: Harvard University Press.

Shannon, C. (1948) A mathematical theory of communication, *Bell System Technical Journal*, **27**, pp. 379–423.

Shewmon, D., Holmes, G. & Byrne, P. (1999) Consciousness in congenitally decorticate children: Developmental vegetative state as self-fulfilling prophecy, *Developmental Medicine and Child Neurology*, **41**, pp. 364–374.

Solms, M. (2021) *The Hidden Spring: A Journey to the Source of Consciousness*, London: Profile, New York: Norton.

Solms, M. & Friston, K. (2018) How and why consciousness arises: Some considerations from physics and physiology, *Journal of Consciousness Studies*, **25** (5–6), pp. 202–238.

Thorndike, E. (1911) *Animal Intelligence*, New York: Macmillan.

Wheeler, A. (1990) *A Journey into Gravity and Spacetime*, New York: W.H. Freeman.
Zeman, A. (2001) Consciousness, *Brain*, **124**, pp. 1263–1289.