'An astonishing achievement. Nick Chater has blown my mind – as well as assuring me that my brain just doesn't work the way I think it does. I haven't been able to stop talking about the ideas in this book'

Tim Harford, author of *Fifty Things That Made the Modern Economy* and *The Undercover Economist*

'This is a remarkable book. Every other book about the mind will tell you either why we're so dumb, or why we're so smart. Chater offers a single elegant theory to explain both: why our minds so often let us down and confound us, at the same time that they far surpass our current attempts to build intelligence in machines'

Josh Tenenbaum, Professor of Cognitive Science and Computation at MIT

'A fascinating intellectual journey … While the mind may indeed be flat in the sense it is devoid of unconscious ruminations, reading this book leaves us with a much deeper, transformed understanding of our own thoughts and feelings, and of how we perceive the definitively non-flat world in which we live'

George Loewenstein, author of *Exotic Preferences: Behavioral Economics and Human Motivation*

'The implications for understanding human decision-making are profound. Everyone who enjoyed *Thinking, Fast and Slow* must read this book'

Gus O'Donnell, former Cabinet Secretary and Chair of the Behavioural Insights Team Advisory Board

ISBN 978-0-241-20844-1

FSC

NICK CHATER THE MIND IS FLAT

patients reported strange fragments of conscious experience, such as the smell of burnt toast, when a piece of cortex was electrical stimulated, but no anomalous conscious experience at all as Penfield removed entire areas of brain.

This perspective explains, too, why patients with visual neglect, where the cortex corresponding to a large area of the visual field may be damaged or entirely inoperative, can none the less be entirely unaware of their deficit. We are consciously aware, perhaps, only of the specific task on which we are currently focused. So, if engaged in fruit-picking, a person with visual neglect will focus their attention only on visual information in parts of the visual cortex which are intact, and link with memory and action systems through the coordinating power of structures deep in the brain, just as for a person with normal visual processing. In deep brain areas, conscious experience may be entirely normal. They will not, of course, pick or describe fruit whose visual positions project into the 'blind' area of visual cortex – so their visual phenomenology, while entirely normal moment by moment, will be restricted to, say, the right half of their visual field.

Our brain is fully engaged with making sense of the information it is confronted with *at each moment*. Consciousness, and indeed the entire activity of thought, appears to be guided, sequentially, through the narrow bottleneck: deep, sub-cortical structures search for, and coordinate, patterns in sensory input, memory and motor output, one at a time. The brain's task is, moment by moment, to link together different pieces of information, and to integrate and act on them right away. Our brain will, of course, lay down fresh memories as this processing proceeds; and draw on the richness of memories of past processing.

So our *no background processing* slogan is reinforced. Or, at least, if there are brain processes which are scurrying about behind the scenes, contemplating, evaluating and reasoning about matters that we appear not to be thinking about at all, then neuroscience has found no trace of them. The brain appears, instead, be concentrating on making sense of immediate experience, and generating sequences of actions, including language (whether spoken aloud or inner speech), through the narrow bottleneck of conscious thought. This is

why it can only integrate and transform to solve *one problem at a time*.

We now have some tentative answers about how the cooperative style of brain computation operates. In the Penfield/Merker vision of the brain, both the questions that the brain faces and the answers that it provides are represented in sub-cortical structures, including

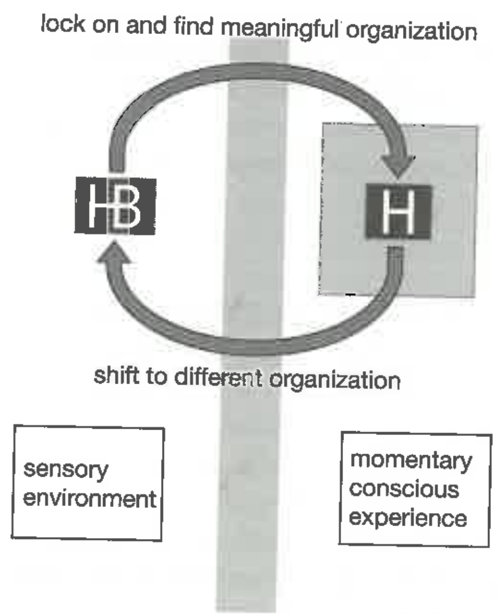

Figure 29. The cycle of thought. *Upper arrow*: The brain locks onto and organizes a fragment of the visual stimulus; we are conscious of, and can report, this organization. *Lower arrow*: But the brain – and the eye – is restlessly struggling to break free of its current organization and lock onto a different fragment of the image. This cycle is so rapid and fluent that we can have the sense of awareness of a complex object – or even an entire scene, in full detail and colour. Our stream of consciousness is of successive sensory organizations – consciousness is entirely confined to the shaded box. We have no conscious access to the information being picked up by our senses (the left-hand side of the figure) or how that information is interpreted (the curved arrows); or how our brain shifts to lock onto different information, e.g. by shifting our attention (with or without moving our eyes).

the thalamus, which serves as a relay station between the cerebral cortex and our sensory and motor systems – essentially as a gateway between the hemispheres of the brain and the outside world. And, we might suspect, both the questions and their answers are concerned primarily with the organization of sensation and movement; and the rich interconnections between these structures and the cortex provide the networks of cooperative computation that are able to solve the problems posed by these sub-cortical structures. Yet while the cortex is crucial in processing visual information, planning movement and drawing on memories, we are aware of the results of the vast cooperative enterprise across the brain only in so far as the results of such computations reach the sub-cortical 'gateway' structures – these, and not the cortex itself, are the locus of conscious experience.

FOUR PRINCIPLES OF THE CYCLE OF THOUGHT

Here is an outline of how the mind works, in the form of a picture (Figure 29) and four proposals. The first principle is that *attention is the process of interpretation*. At each moment, our brain 'locks onto' (or, in everyday terms, pays attention to) a target set of information, which the brain then attempts to organize and interpret. The target might be aspects of sensory experience, a fragment of language, or a memory.[9] Figure 29 illustrates a case where the brain has momentarily locked onto the 'H' in complex stimulus. Moments later, it might alight on the 'B'. Crucially, our brain locks onto one target at a time. This implies that the line shared between the 'H' and 'B' in the stimulus is, at any given moment, interpreted as belonging to one or the other – but not both. Following Penfield and Merker's conjectures, we might suspect that such information is represented in sub-cortical structures buried deep in the brain, with connections across the entire cortex, so that the full range of past experience and knowledge can be brought to bear on finding meaning in the current target. Remember that we can lock onto, and integrate, information of all kinds. We can use any and all pieces of information and great leaps of ingenuity

and imagination to find meaning in the world, but we can only create one pattern at a time.

Our second principle concerns the nature of consciousness and is that our only conscious experience is our interpretation of sensory information. The result of the brain's 'interpretations' of sensory input is conscious – we are aware of the brain's interpretation of the world – but the 'raw materials' from which this interpretation is constructed, and the process of construction itself, are not consciously accessible. (Conscious experience is the organization represented in deep brain structures, with inputs from across the cortex – we are not directly conscious of cortical activity itself.) So in Figure 29, we perceive the 'H' or the 'B', but have absolutely no awareness of the process through which these were constructed.

Perception always works this way: we 'see' objects, people, faces as a result of a pattern of firing from light-sensitive cells triggered by light falling on our retina; we 'hear' voices, musical instruments and traffic noise as a result of picking up the complex patterns of firing of vibration-detecting cells in our inner ear. But we have no idea, from introspection alone, where such meaningful interpretations spring from – how our brain makes the jump from successive waves of cacophonous chatter from our nervous system to a stable and meaningful world around us. All we 'experience' is the stable and meaningful world: we experience the result, not the process.[10]

So far, we have focused on our consciousness of meaning in sensory information. The third principle is that we are conscious of nothing else: *all conscious thought concerns the meaningful interpretation of sensory information*. But while we have no conscious experience of non-sensory information, we may be conscious of their sensory 'consequences' (i.e. I have no conscious experience of the abstract number 5, although I may conjure up a sensory representation of five dots, or the shape of symbol '5'). Deep brain regions are, after all, relay stations for conveying sensory information to the cortex – so if they are locations of conscious experience, then we should expect conscious experience to be sensory experience, and nothing more.

The claim that we are aware of nothing more than meaningful organization of sensory experience isn't quite as restrictive as it sounds. Sensory information need not necessarily be *gathered* by our senses, but may be invented in our dreams or by active imagery. And much sensory information comes, of course, not from the external world but from our own bodies – including many of our pains, pleasures, and sensations of effort or boredom. We are conscious of the sounds or shapes of the words we use to encode abstract ideas; or the imagery which accompanies them. But we are not conscious of the abstract ideas themselves, whatever that might mean. I can imagine (just about) three apples or the symbols '3', 'iii' or 'three'; and I can imagine various rather indistinct triangles and the word 'triangle'. But I surely can't imagine, or in any sense be conscious of, the abstract number 3; or the abstract mathematical concept of 'triangularity'. I can hear myself say 'Triangles have three straight sides' or 'The internal angles of a triangle add up to 180 degrees' – but I surely don't have any additional conscious experience of these abstract truths.

Similarly, as we have seen already, it is a mistake to think that we are conscious of any of our beliefs, desires, hopes or fears. I can say to myself 'I'm terrified of water' or I can have visions of myself struggling desperately as a rip-tide pulls me out to sea. But it is words and images that are the objects of consciousness – not the 'abstract' belief. Just in case you doubt this viewpoint, reflect on what beliefs you are conscious of *right now*. How many are there, exactly? Can you feel when one belief leaves your consciousness or a fresh belief 'comes into mind'? I suspect not.[11]

Now we can tie together our three proposals into a fourth principle. I have proposed that an individual conscious thought is the process of the creation of a meaningful organization of sensory input. So what is the stream of consciousness? Nothing more than a succession of thoughts, an irregular cycle of experiences which are the results of sequential organization of different aspects of sensory input – the shifting contents of the right-hand box in Figure 29. This fits with the Penfield/Merker story about the brain: sub-cortical structures deep in the brain form a 'crucible' onto which the resources of the whole cortex can be focused to impose meaning on fragments of sensory information – but only one pattern can be placed in the crucible at a time.

Note, in particular, that the cycle of thought is sequential: we lock

onto and impose meaning on one set of information at a time. Now of course your brain can control your breathing, heart-rate and balance independent of the cycle of thought – to some extent at least (we don't topple over when particularly engrossed in a problem). But the brain's activities beyond the sequential cycle of thought are, we shall see, surprisingly limited – we can manage, roughly speaking, just one thought at a time.

From this point of view, many of the strange phenomena we saw in Part One fall into place:

- The brain is continually scrambling to link together scraps of sensory information (and has the ability, of course, to gather more information, with a remarkably quick flick of the eye). We 'create' our perception of an entire visual world from a succession of fragments, picked up one at a time (see Chapter 2). Yet our conscious experience is merely the output of this remarkable process; we have little or no insight into the relevant sensory inputs or how they are combined.
- As soon as we query some aspects of the visual scene (or, equally, of our memory), then the brain immediately locks onto relevant information and attempts to impose meaning upon it. The process of creating such meaning is so fluent that we imagine ourselves merely to be reading off pre-existing information, to which we already have access, just as, when scrolling down the contents of a word processor, or exploring a virtual reality game, we have the illusion that the entire document, or labyrinth, pre-exist in all their glorious pixel-by-pixel detail (somewhere 'off-screen'). But, of course, they are created for us by the computer software at the very moment they are needed (e.g. when we scroll down or 'run' headlong down a virtual passageway). This is the sleight of hand that underlies the grand illusion (see Chapter 3).
- In perception, we focus on fragments of sensory information and impose what might be quite abstract meaning: the identity, posture, facial expression, intentions of another person, for example. But we can just as well reverse the process. We can focus on an abstract meaning, and create a corresponding sensory

image: this is the basis of mental imagery. So just as we can recognize a tiger from the slightest of glimpses, we can also *imagine* a tiger – although, as we saw in Chapter 4, the sensory image we reconstruct is remarkably sketchy.

- Feelings are just one more thing we can pay attention to. An emotion is, as we saw in Chapter 5, the *interpretation* of a bodily state. So experiencing an emotion requires attending to one's bodily state as well as relevant aspects of the outer world: the interpretation imposes a 'story' linking body and world together. Suppose, for example, that Inspector Lestrade feels the physiological traces of negativity (perhaps he draws back, hunches his shoulders, his mouth turns down, he looks at the floor) as Sherlock Holmes explains his latest triumph. The observant Watson attends, successively, to Lestrade's demeanour and Holmes's words, searching for the meaning of these snippets, perhaps concluding: 'Lestrade is jealous of Holmes's brilliance.' But Lestrade's reading of his own emotions works in just the same way: he too must attend to, and interpret his own physiological state and Holmes's words in order to conclude that he is jealous of Holmes's brilliance. Needless to say, Lestrade may be thinking nothing of the kind – he may be trying (with frustratingly little success) to find flaws in Holmes's explanation of the case. If so, while Watson may interpret Lestrade as being jealous, Lestrade is not experiencing jealousy (of Holmes's brilliance, or anything else) – because experiencing jealousy results from a process of interpretation, in which jealous thoughts are the 'meaning' generated, but Lestrade's mind is attending to other matters entirely, in particular, the details of the case.
- Finally, consider choices (see Chapter 6). Recall how the left hemisphere of a split-brain patient fluently, though often completely spuriously, 'explains' the mysterious activity of the left hand – even though that hand is actually governed by the brain's right hemisphere. This is the left, linguistic brain's attempt to impose meaning on the left-hand movements: to create such meaningful (though, in the case of the split-brain patient, entirely illusory) explanation requires locking onto the activity of the left hand in order to make sense of it. It does not, in particular, involve locking onto any hidden inner motives lurking within the right hemisphere (the real controller of the left hand) because the left and right hemispheres are, of course, completely disconnected. But notice that, even if the hemispheres were connected, the left hemisphere would not be able to attend to the right hemisphere's inner workings – because the brain can *only* attend to the meaning of perceptual input (including the perception of one's own bodily state), not to any aspect of its own inner workings.

We are, in short, relentless improvisers, powered by a mental engine which is perpetually creating meaning from sensory input, step by step. Yet we are only ever aware of the meaning created; the process by which it arises is hidden. Our step-by-step improvisation is so fluent that we have the illusion that the 'answers' to whatever 'questions' we ask ourselves were 'inside our minds all along'. But, in reality, when we decide what to say, what to choose, or how to act, we are, quite literally, *making up* our minds, one thought at a time.

8

The Narrow Channel of Consciousness

If thoughts are a cycle, then it follows that we have thoughts one at a time. More specifically, we can only focus on, and attempt to impose meaning on, just one set of information at a time. But the brain does many things at once. Most of us can, as the saying goes, walk and chew gum at the same time; and also walk, chew gum and be shocked by an overheard conversation. But if our mind is locked onto the conversation, it will not simultaneously be locked onto the walking or the gum-chewing: these activities, like the control of our breathing and our heart-rate, will, in a very real sense, be mindless. These processes are precisely those that do not involve interpretation (i.e. our best attempt to imaginatively apply anything and everything that we know in order to make sense of the information that is currently in our mental 'focus'). Such mindless, automatic processes turn out to be very limited, both in what they can do and how well they can perform (though with the occasional surprising exception, as we'll see later).

Yet if the mind is able only to lock onto one set of information at a time, does this mean that we are effectively oblivious to anything we are not currently paying attention to? Not quite. First, automatic processes such as gum-chewing and walking can continue uninterrupted, and these require the processing of some sensory information – about the terrain in front of us, our posture, limb positions and muscle activity to make sure that we don't topple over, or sensory information about the inner world of our own mouths, if we are not inadvertently to bite our own tongues. Secondly, there is the question of vigilance, even for information we are not currently attending to. Remember that the periphery of the retina is continually monitoring

the signs of motion, flashes of light, or other abrupt changes; our auditory system is alert, to some degree at least, to unexpected bangs, creaks or voices; our bodies are 'wired' to detect unexpected pains or prods. In short, our perceptual system is continually ready to raise the alarm – and to drag our limited attentional resources away from their current task in order to lock onto a surprising new stimulus. But these 'alarm systems' don't themselves involve the interpretation and organization of sensory input; instead they help *direct* our attempts to organize and interpret sensory input. So we do not know what it is that has attracted our attention until we have locked onto the unexpected information and attempted to make sense of it.[1]

This means that we are sometimes oblivious to information to which we are not attending, even though it may be in plain view. Such 'inattentional blindness' seems highly counter-intuitive, but turns out to be all too real. The perceptual psychologists Arien Mack and Irvin Rock asked the participants in their experiments to fixate on a small cross in the centre of a computer screen. Then, a much larger cross appeared on the screen – and the task was to judge whether the horizontal or vertical arm of the large cross was longer. As can be seen in Figure 30, this is a fairly subtle discrimination, requiring careful attention. The large cross disappeared after one fifth of a second, before being replaced by a random black-and-white 'mask' pattern (on the right in Figure 30). Previous studies have shown that the mask will obliterate further visual analysis of the cross. The mask was just used to control the amount of time that people could look at the cross. If no mask was used, and the screen simply went blank, the participants would still potentially have access to an after-image of the cross on their retinas.

The viewer initially fixates on a small central cross; then comes the 'critical stimulus', with its large central cross. The viewer's task is to report whether the vertical or horizontal 'arm' of the cross is the longer. After one fifth of a second, the critical stimulus is obliterated by a 'mask'.

The key moment in Mack and Rock's experiment came on the third or fourth trial, when they introduced an additional object, for example a black or coloured blob a couple of degrees away from the point of fixation (and hence projected near to, although not actually