From
The Future of the Brain: Essays by the World's Leading Neuroscientists
Jeremy Freeman and Gary Marcus (eds.), Princeton University Press

CONSCIOUSNESS, BIG SCIENCE, AND CONCEPTUAL CLARITY

Ned Block

With enormous investments in neuroscience looming on the horizon, including proposals to map the activity of every neuron in the brain, it is worth asking what questions such an investment might be expected to contribute to answering. What is the likelihood that high-resolution mapping will resolve fundamental questions about how the mind works? I will argue that high-resolution maps are far from sufficient, and that the utility of new technologies in neuroscience depends on developing them in tandem with the psycho-neural concepts needed to understand how the mind is implemented in the brain.

Using high school geometry, we can understand why a rigid round peg won't fit into a square hole in a board; mapping every single particle in the peg and board would be of little use without the high school geometrical account, as Hilary Putnam once noted. Similarly, a map of the activation of every neuron in the brain will be of no use without a psychological level understanding of what those activations are doing. For this reason, advocates of high-resolution mapping have advocated a "functional brain map." It is easy to add the word "functional," but massive quantities of data alone cannot produce theoretical breakthroughs in understanding the mind at a psychological level. Using the example of consciousness, I will discuss one of the obstacles to constructing a functional brain map that explains how neural activations function to underlie human psychology and how the obstacle can be circumvented without high-density brain imaging. The obstacle is the measurement problem of finding consciousness in the brain.

The Measurement Problem

The measurement problem of finding consciousness in the brain depends on the fundamental distinction between consciousness and cognition. Consciousness is what it is like to have an experience. Cognition includes thought, reasoning, memory, and decision, but all of these cognitive processes can occur unconsciously. Consciousness and cognition can causally interact, and of course cognition can be conscious, but they fall on opposite sides of a joint in nature. I will focus on the difference between conscious perception—what it is like to have a perceptual experience—and perceptual cognition—the processes in which perceptual experiences play a role in thought, reasoning, and the control of action. If an experimenter wants to know whether a subject in an experiment has consciously seen, say, a triangle, the subject has to do something, for example, say whether a triangle was present. For a subject to categorize what was seen as a triangle requires computational processes, say retrieving a representation of a triangle from memory and comparing the conscious percept with the memory trace, and there will be a further cognitive process of deciding whether to respond, and then if the decision is to respond, enumerating and deciding among candidate responses and generating a response. Further, one of the cognitive processes that can occur during a conscious percept of a triangle is a decision whether to further attend to the triangle, and subsequently the top-down attentional processes themselves. Since these cognitive processes are all in service of cognitively accessing the perceptual information and applying that information to a task, let us lump these cognitive processes all together as processes of *cognitive access*. The measurement problem, then, is how to distinguish the brain basis of consciousness from the brain basis of cognitive access.

Note that the measurement problem is distinct from David Chalmers's "hard problem" of consciousness, the problem of explaining why the brain basis of an experience of red is the brain basis of *that* type of experience rather than the experience of green or no experience at all. The hard problem depends on a prior notion of "brain basis" of the experience of red. We should be able to say *what* the brain basis of the experience of red is even if we cannot explain *why* that brain basis is the basis of that experience rather than another experience.

Why is the measurement problem a problem? Cognitive neuroscientists have identified many specialized circuits in the brain. The methodology is simple: compare the circuits that are active in, say, face perception with those that are active in other kinds of perception or when there is no perception. This methodology has resulted in the identification of the "fusiform face area" and two other linked face areas. Why can't neuroscientists just use the same idea applied to consciousness: compare what is happening in the brain during a conscious percept with what is happening in the brain during a comparable unconscious percept? One useful procedure involves presenting the subject with a series of stimuli that are at the threshold of visibility. Given the probabilistic nature of visual processing, the subject sometimes does and sometimes does not see threshold stimuli consciously. The stimuli remain the same, only the consciousness changes, so the perceptual processes common to both conscious and unconscious perception can be distinguished from the processes underlying consciousness of the stimulus. This is the "contrastive method." The problem is that, as just noted, we can only tell the difference between conscious and unconscious perception on the basis of the subject's response. So when we compare conscious with unconscious perception, we inevitably lump together the neural basis of the conscious percept with the neural basis of the response to that percept. Since the neural basis of the response underlies the very cognitive processes that I have lumped together as "cognitive access," the contrastive method inevitably conflates the neural bases of conscious perception with the neural basis of cognitive access to the perceptual content. The problem has seemed so severe that many regard it as intractable, resigning themselves to studying what I have called "access consciousness," that is, an amalgamation of the machinery of consciousness together with the machinery of cognitive access.

Further, as Lucia Melloni and her colleagues have recently shown, there are always *precursors* to a conscious state that may not be part of the neural basis of consciousness (Aru et al. 2012). For example, whether one sees a stimulus or not depends not only on fluctuations in attention but also on fluctuations in spontaneous brain activity that occur before the stimulus that may set the stage for consciousness without being part of it. To solve the measurement problem we must manage to separate consciousness from the nonconscious processes that inevitably accompany it in the situations in which we know consciousness obtains.

Indeed, the measurement problem is even thornier than I have suggested so far. Consider, for example, a type of brain injury (involving lesions in the parietal lobe) that causes a syndrome known as visuo-spatial extinction. If the patient sees a single object on either the left or the right, the patient can identify it, but if there are objects on both sides, the patient claims not to see one of the items; if the brain damage is on the right, the patient will claim to not to see the item on the left because perceptual fibers cross in feeding to the brain. However, in one such case in which a patient identified as "GK" was presented with two objects, including a face on the left that he said he did not see, Geraint Rees showed him to have activation in the relevant face area (the "fusiform face area") to almost the same degree as when he reported seeing the face. How could we find out whether GK has a conscious face experience that he does not know he has? It may seem that all we have to do is find the neural basis of face experience in unproblematic cases and ascertain whether this neural basis obtains in GK when he says he sees nothing on the left. The problem is that subjects who report seeing a face differ from those who deny seeing a face in activation of the neural basis of cognitive access to seeing a face in the frontal and parietal lobes. So it seems that in order to answer the question about GK we must first decide whether the neural basis of cognitive access to seeing a face is part of the neural basis of the conscious experience of seeing a face. And this was the question we started with.

One might wonder whether it even makes sense for GK to have a conscious face experience that he does not know about. What makes the measurement problem so problematic is the possibility that some aspect of cognitive access is actually partly constitutive of consciousness itself. If cognitive access is partly constitutive of consciousness itself, then GK could not possibly have a face experience he does not know about. If we do not solve the measurement problem, we could record every detail of activation in the face circuit and other circuits in the brain without determining whether those activations are conscious or unconscious.

The measurement problem is particularly trenchant for consciousness, but aspects of the problem arise for other mental phenomena. Masses of high-resolution data about neural activations are no use without an understanding of what the neural activations are doing at a psychological level. Once we have a theory at the psychological level, high-resolution brain data may tell us whether the theory makes correct predictions. But without the theory at the psychological level, the data are of no use no matter how high the resolution.

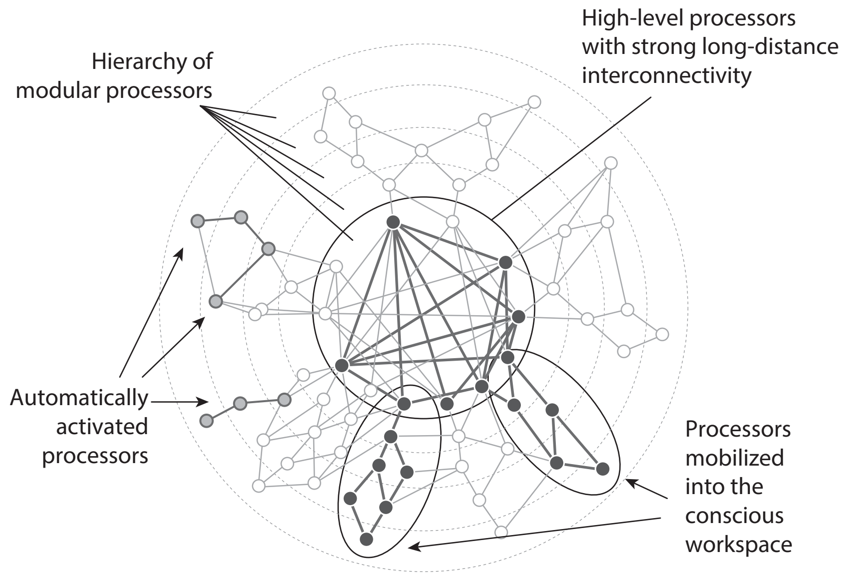

Figure 1. Diagram of the global neuronal workspace. Neural processors are symbolized by circles and connections between them by lines. Filled circles and bold lines indicate activation. The outer circles indicate sensory input, whereas the center indicates the areas in the front of the brain responsible for cognition. From Dehaene and Nacchache (2001). With permission of Elsevier.

Cognitive versus Noncognitive Theories of Consciousness

This issue—of whether cognitive access is part and parcel of consciousness—divides the field. Cognitive theories of consciousness say yes. Stanislas Dehaene, Jean-Pierre Changeux, and their colleagues (2011) have advocated a "global neuronal workspace" theory of consciousness. According to that theory, neural coalitions in the sensory areas in the back of the head compete with one another, the winners triggering "ignition" of larger networks via long-range connections to frontal areas responsible for a variety of cognitive functions. The activation of the central network feeds back to the peripheral sensory activations, maintaining their firing. Once perceptual information is part of a dominant coalition, it is available for all cognitive mechanisms and is said to be "globally broadcast" (see figure 1).

According to the global neuronal workspace theory, consciousness just is global broadcasting. Many philosophers and scientists hold versions of this view, including Sid Kouider, Daniel Dennett, and in a more attenuated form, Jesse Prinz. This is a cognitive theory of consciousness because the global workspace governs cognitive processes such as categorization, memory, reasoning, decision, and control of action. An alternative cognitive theory of consciousness David Rosenthal and Hakwan Lau (2011) hold emphasizes higher-order thought: a perception is conscious if it is accompanied by a thought about that perception. (The thought is higher order in that it is about another mental state.)

An opposed point of view, which Victor Lamme, Ilja Sligte, Annelinde Vandenbroucke, Semir Zeki, and I hold, is that activations in perceptual areas in the back of the head can be conscious without triggering global broadcasting. It is not part of our view that there can be conscious experience without any possibility of cognitive access, but only that there can be conscious experience without *actual* cognitive access. This point is shown in an experimental paradigm from Victor Lamme's laboratory illustrated in figure 2. The subject sees a circle of rectangles, then a gray screen, then another circle of rectangles. A line appears indicating the position of one of the rectangles. The line can occur with the second circle of rectangles as in A, or with the first circle as in B, or in the middle, as in C. The subject is supposed to say whether the indicated rectangle changes orientation between the first and second circle. Subjects can do this almost perfectly in B but are bad at it in A with a capacity of only four of the eight rectangles. The interesting case is C when the line appears during the gray screen. If the subjects are continuing to maintain a visual representation of all or almost all the rectangles (as they say they are doing), the difference between C and B will be small, and this is what is found. Subjects have a capacity of almost seven of the eight rectangles even when the line appears in the gray period 1.5 seconds after the first circle. The point illustrated here is that subjects can have a conscious experience of all the rectangles even though it is only possible to actually cognitively access half of them. Thus Victor Lamme and I argue that contrary to the views of those who favor a cognitive theory of consciousness, the neural basis of consciousness does not include the neural basis of actual cognitive access.

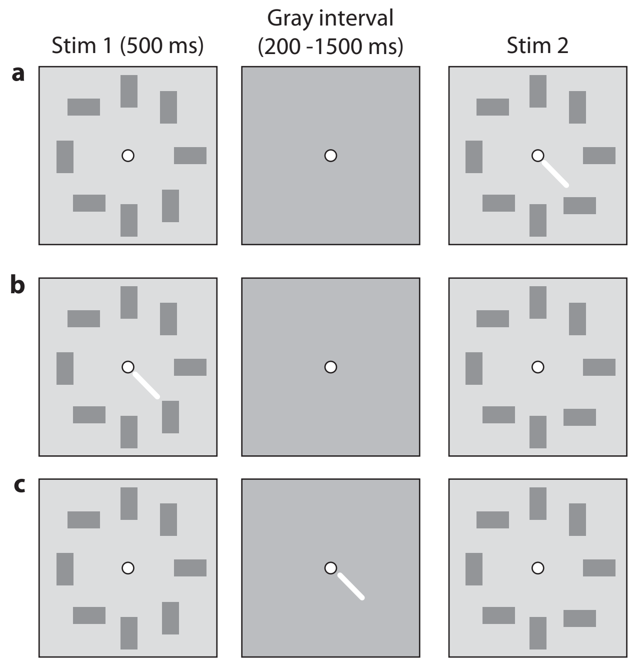

Figure 2. A perceptual task used in Victor Lamme's laboratory at the University of Amsterdam. A circle of rectangles is presented for half a second, then a gray screen for a variable period, then a new circle of rectangles. At some point in this process the subject sees a line that indicates the position of one of the rectangles. The subject's task is to say whether the rectangle at that position has changed orientation between the first and second circle of rectangles. From Lamme (2003). With permission of Elsevier.

As you might guess, this dispute has involved heavy polemics. In his 2014 book, Stanislas Dehaene says our point of view leads to dualism. He says, "The hypothetical concept of qualia, pure mental experience detached from any information-processing role, will be viewed as a peculiar idea of the prescientific era" (Dehaene 2014, 221). Of course, Lamme, Zeki, and I do not think that phenomenal consciousness has no information-processing role. We think that consciousness greases the wheels of cognitive access but can obtain without it.

The measurement problem under discussion is how it is possible for evidence to count one way or the other as between cognitive and non-cognitive theories of consciousness, given that our ability to find out whether a perception is conscious or not depends on cognitive processes by virtue of which the perception surfaces in the very behavior that provides evidence of consciousness. Some theorists have held that the measurement problem may be solved by new technology, a subject to which we now turn.

Transgenic Mice and the Optogenetic Switch

Global broadcasting involves not only feed-forward flow of activation but heavy feedback from frontal to sensory areas. Christof Koch and Nao Tsuchiya (2014) propose to use transgenic mice whose neural genes have been rendered sensitive to light, for example, by being infected with genetically altered viruses. In these mice, top-down feedback from frontal to sensory areas can be turned off optogenetically by light sources on the skull or optical fibers implanted in the brain. If there is no top-down attentional feedback there can be no "ignition" and no global broadcasting. Koch and Tsuchiya predict that without attentional feedback, the mice will be able to consciously see a single object with no distractors. On their view, top-down attention may only be required to single out an item in the visual field from other items. For example, one can detect a red "T" without top-down attention if it is the only visible object, but it takes top-down attention to detect a red "T" when the display also contains distractors: black "T"s and red "F"s.

Suppose their prediction is confirmed that the mice will be able to do a task without distractors but not when there are distractors. How we are supposed to know whether the mice whose top-down feedback has been deactivated by the optogenetic switch are doing their tasks *consciously*? Koch and Tsuchiya propose to use postdecision wagering in which the mice express their confidence in their choice by in effect betting on whether the choice is right or not. Here is how postdecision wagering works in people: the subject is given credits that are worth money. In each trial the subject makes a decision as to whether there was a stimulus present and then bets on whether that decision was right. There is a condition known as blind-sight in which destruction of parts of the lowest-level visual cortex render the subjects incapable of consciously seeing objects in the destroyed part of the visual field. Subjects can guess with very high degrees of accuracy what is presented, but they have the phenomenology of guessing, not of seeing. These blind-sight subjects bet very poorly in postdecision wagering since they have no idea which of their guesses are right, and that has suggested that betting can provide an index of conscious perception.

It turns out that animals can do something equivalent to betting to get more food pellets. And Koch and Tsuchiya say that one may be able to use postdecision wagering to test whether the optogenetic mice are consciously seeing the stimulus. High confidence would suggest conscious perception; low confidence unconscious perception. But won't the shutting off of top-down processes ruin wagering in the mice? Koch and Tsuchiya think that confidence may be mediated by different top-down processes from those involved in attention and global broadcasting and so may not be turned off by the optogenetic switch.

One way to think about this proposal is to try to imagine what it would be like to be an optogenetic mouse. Suppose you are a transgenic being whose optogenetic switch has been flipped so as to preclude top-down attention. And suppose Koch and Tsuchiya are right that you would have conscious experience. What would that experience be like? Without top-down attention, that experience would be a kaleidoscopic chaotic array of fragmentary perceptions in all sensory modalities with no sustained attention in one modality or on one thing rather than another. (Alison Gopnik has suggested that this is what it is like to be an infant in the first months of life since these infants have many more synapses and more myelination in sensory areas than in the frontal areas responsible for top-down attention.) Suppose that before the switch is flipped, you had been trained to respond to a red "T" either by itself or in a sea of black "T"s and red "F"s. Now the switch is flipped and you have a visual impression of the red "T" as part of "blooming buzzing confusion" of percepts in all sensory modalities. How much would you bet that your perception of the red "T" was accurate? It is certainly possible that the effect of the kaleidoscopic chaotic perception would be to lower one's confidence in any one percept.

Now suppose instead that the prediction of Koch and Tsuchiya is wrong—that when the optogenetic switch is flipped, it knocks out conscious perception as well as top-down attention. Without top-down signals there can be no global broadcasting. Still, the subject might be able to reliably guess whether there is a red "T" on the basis of unconscious perception as with the blind-sight patient. How would betting behavior be affected? All but one of the blind-sight patients that have been studied have had a partially blind and partially sighted field. The one human blind-sight patient whose entire visual field was blind was able to walk, with apparent confidence, through an obstacle-laden hallway. So it is hard to predict how confident a perceiver with only unconscious vision would be. In sum, betting might not correlate with consciousness once the optogenetic switch was flipped.

The upshot is that although the use of transgenic mice could make an important contribution, it would just be another line of evidence that cries out for interpretation.

Nonconceptual Representations and the Measurement Problem

Coming to grips with the measurement problem requires rethinking the basic ideas we are using. Here is a model of perception that appears in Tyler Burge's monumental *Origins of Objectivity* (2010).

Burge distinguishes between an attribute, say the circularity of the plate, and a perceptual representation, what he calls an "attributive," for example, a perceptual representation of circularity. The format of a perceptual representation is iconic and can be represented in words as "That X" where the "that" is an element that picks out an individual, the plate on the left in figure 3, and the "X" is a pure perceptual representation that picks out the circularity of the plate. The next stage to the right of the perception in figure 3 is a basic perceptual judgment in which the perceiver judges that the item is circular. Note: "That X" contains no concept, whereas "That is circular" contains the concept circular; and "That X" does not make a statement or judgment, that is, it does not say that anything is so or is the case. A basic perceptual judgment like "That is circular" is produced via the application of the concept of circularity to the percept to yield a structured propositional mental representation.

Why are we discussing percepts and concepts? Coming to grips with the measurement problem depends on understanding of the difference between two kinds of experiences: nonconceptual perceptions and conscious perceptual judgments involving concepts.

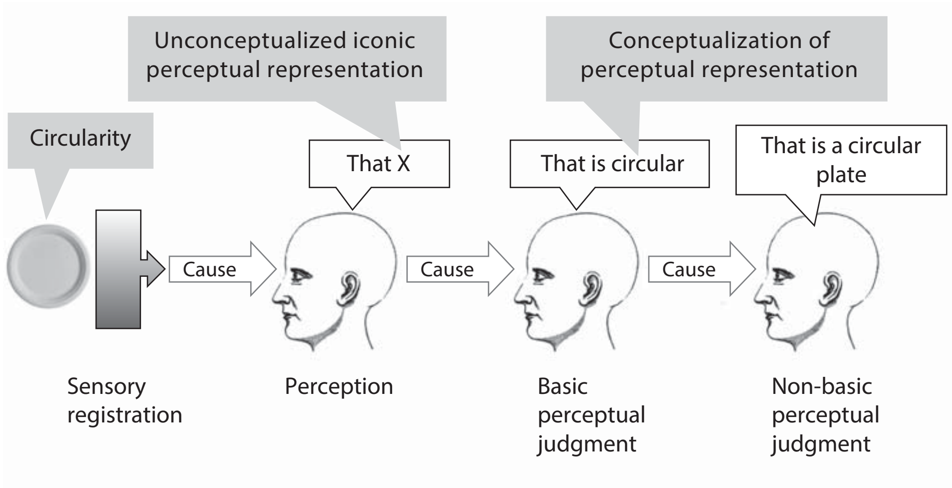

Figure 3. Burge's model of perception. © Ned Block 2013. Feed-backward influences have been omitted from this diagram. There are no top-down effects on the retina, but there are top-down influences at every other level.

What is a concept? As I am using the term "concept," a concept is a constituent of a thought or judgment that applies to something, as "circular" applies to the plate.

It is extremely important to keep separate concepts from what they are concepts of, a common confusion. For example, Bruno Latour infamously claimed that Ramses II could not have died of tuberculosis since Robert Koch discovered tuberculosis in 1882. He said, "Before Koch, the bacillus had no real existence. To say that Ramses II died of tuberculosis is as absurd as saying that he died of machine-gun fire." However, what did not exist before 1882 was not the tuberculosis bacillus, but rather the human concept of that bacillus. Many people died of tuberculosis before any humans had the concept of what killed them.

I mentioned one difference between percepts and concepts: format. Percepts are iconic; concepts are parts of thoughts or judgments that are "propositional"—they have a structure analogous to that of a sentence. Another difference is computational role: percepts are, to a first approximation, elements in a modular system, whereas concepts have a much wider role in thinking, inferring, deciding, and the like. But what is important here is not what the exact distinction is between percepts and concepts but rather that there is a joint in nature whose exact characterization is still an object of study.

In Burge's model of perception, there are two different items that could be thought of as aspects of conscious perception: the nonconceptualized percept itself and the basic perceptual judgment. A conscious percept may require little or no cognition. Perhaps a mouse could consciously perceive circularity even with no ability to think or reason about circularity. A conscious basic perceptual judgment by contrast is something that exists only in concept-using creatures, creatures that can think and reason. Although percepts can be unconscious as well as conscious, the distinction between a nonconceptual percept and a basic perceptual judgment can help in thinking about the measurement problem. One of the big advances in consciousness research in the 1990s was the realization by Francis Crick and Christof Koch that because the visual apparatus of many mammals is similar to our own, we can study perceptual consciousness in these animals even though they lack the linguistic capacities required for much of thought and reasoning. I now turn to a discussion of how the distinction may be relevant to actual experiments.

Simple Methodological Advance: Don't Ask for a Report

The familiar brain imaging pictures one sees in newspapers typically represent active brain areas. The imaging technology that produces these images—fMRI, PET, CAT—all localize spatially without much capacity to localize temporally. But in the study of conscious perception, time has proven to be as important if not more important than space. One useful technology is that of "event related potentials," or ERPs, in which electrodes placed on the scalp measure the temporally varying reaction to an event, say a visual stimulus. The brain reaction to a visual stimulus has a number of identifiable components, and researchers can and do ask which of these components correlate best with visibility of the stimulus. Stanislas Dehaene and other advocates of the global broadcasting approach have used ERP technology to find the neural basis of consciousness. And their efforts have provided evidence that the

ERP component that reflects visibility happens late in the process, when frontal concept representations have been brought into play—which is what the global broadcasting theory predicts. However, the methods Dehaene and his colleagues have used involve conceptualization of the stimulus. One study presented a target digit that was on the threshold of visibility, and the objective index of whether subjects saw it was whether the subjects could say whether the digit was larger or smaller than 5, a task that required the subject to conceptualize the seen shape in arithmetical terms and to perform an arithmetic operation, a conceptually loaded task. In another experiment, subjects had to report whether they saw the name of a number, again a task that required conceptualization of the stimulus. It is reasonable to object that what the ERP methods were revealing was not the pure percept but instead a perceptual judgment in which a concept was applied to the percept.

How can we avoid such a trap? Michael Pitts (2011) presented a series of 240 trials in which subjects saw a red ring with small discs on it. The subjects' task was to focus on the ring, looking for one of the discs to dim. Meanwhile, in the background of the ring, there were a myriad of small line segments that could be oriented randomly or, alternatively, some of the segments could be oriented so as to form one or another geometrical figure. About half the time, there was a rectangular background figure. After 240 trials of stimuli and responses about the discs were over, Pitts asked subjects to answer a series of questions that probed whether they had seen any figures in the background in the 240 trials, how confident they were about having seen these figures and what figures they saw. Those who were at least moderately confident of having seen a rectangle showed a different ERP profile from the others, and that profile differed markedly from what Dehaene and his colleagues had reported: the ERP components that correlated best with judged visibility of the rectangle came *before* global broadcasting, suggesting that subjects consciously experienced the rectangles prior to making the perceptual judgment that there was a rectangle. The activations were in perceptual areas and not in frontal areas responsible for conceptualization. The key innovation in this experiment was simple and low tech: the *relevant* conscious experience was not related to any task until *after* the perception was long gone, so the usual conflation of consciousness and cognition may not have occurred.

The idea of not asking the subject to do anything was used with an entirely different paradigm, binocular rivalry, by Wolfgang Einhäuser's lab (Frässle et al. 2014). Binocular rivalry is a phenomenon that was discovered in the sixteenth century in which two different images are presented to the two eyes. The subject's whole visual field is filled by one, then the other; the two interpretations of the world alternate with only momentary mixtures of the two images. For example, one eye may be fed a grid moving to the left and the other eye fed a grid moving to the right. The subject is aware of left motion, then right motion, then left motion, and so on. Many studies have shown that as the rivalrous percepts alternate, activations change both in the visual areas in the back of the head and in the global broadcasting areas in the front of the head, and many have taken this to support the global broadcasting theory of conscious perception. Plate 9 illustrates one of the first of these studies in which one eye is fed an image of a face and the other eye an image of a house. The percept alternates between face and house and allowed researchers to pinpoint a circuit in the brain that specializes in faces and another that specializes in houses (see plate 9).

In the original binocular rivalry experiments, subjects reported what they were seeing by pressing a button. The Einhäuser experiment used a new method of telling when the percept shifted that did not require the subject to respond. The new method involved small eye movements that tip the experimenter off as to whether the subject is perceiving leftward or rightward motion and, in another version, changes in pupil size. The subjects' button presses validate the eye movement method, but once the method is validated the subjects do not have to do any task. The interesting result was that when there was no task there was no differential frontal brain activity. All the differences in conscious perception were in the visual and spatial areas in the back and middle of the head. The authors conclude that previous results that showed frontal global workspace changes reflected the self-monitoring required to make a response, but that when no response was required, there was little or no monitoring. Stanislas Dehaene says in his 2014 book that when "the prefrontal cortex does not gain access to . . . [a] message, it cannot be broadly shared and therefore remains unconscious" (2014, 155). But what these experiments suggest is that perceptual representations can be consciously experienced even when not actually accessed—not broadcast in the global workspace—so long as they are accessible.

This study did use new technology, but it was behavioral technology—the use of eye movements and changes in pupil size to differentiate one percept from another. These results were combined with ordinary resolution brain imaging, but ordinary resolution can be good enough when you know what you are looking for.

So we have made enormous progress in solving the measurement problem, but that progress depended on conceptual clarity, behavioral technology, and low-tech brain imaging, not expensive high-resolution brain imaging. The lesson to be drawn is that isolating consciousness in the brain may depend more on being clear about what we are looking for than on massive investments in new technology. More broadly, high-resolution data are of no use without a theory of what brain activations mean at the psychological level. When we have substantive cognitive neuroscience theories—together with the sophisticated concepts embedded in such theories—testing these theories may require Big Science. But we cannot expect the theories and concepts to somehow emerge from Big Science. To paraphrase Immanuel Kant, concepts without data are empty; data without concepts are blind; "only through their unison can knowledge arise" (Kant 1787, 75).

References

Aru, J., T. Bachmann, W. Singer, and L. Melloni. 2012. "Distilling the Neural Correlates of Consciousness." *Neuroscience and Biobehavioral Reviews* 36: 737–46.

Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.

Cohen, M., and D. Dennett. 2011. "Consciousness Cannot Be Separated from Function." *Trends in Cognitive Sciences* 15 (8): 358–64.

Dehaene, S. 2014. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. New York: Viking.

Dehaene, S., and J. Changeux. 2011. "Experimental and Theoretical Approaches to Conscious Processing." *Neuron* 70: 200–227.

Dehaene, S., and L. Nacchache. 2001. "Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework." *Cognition* 79: 1–37.

Frässle, S., J. Sommer, A. Jansen, M. Naber, and W. Einhäuser. 2014. "Binocular Rivalry: Frontal Activity Relates to Introspection and Action but Not to Perception." *Journal of Neuroscience* 34 (5): 1738–47.

Kant, I. 1787. *Critique of Pure Reason*. Translated and edited by P. Guyer and A. Wood. Cambridge: Cambridge University Press, 1997.

Lamme, V. 2003. "Why Visual Attention and Awareness Are Different." *Trends in Cognitive Sciences* 7: 12–18.

Lau, H., and D. Rosenthal. 2011. "Empirical Support for Higher-Order Theories of Conscious Awareness." *Trends in Cognitive Sciences* 15 (8): 365–73.

Latour, B. 1998. "Ramses II est-il mort de la tuberculose?" *La Recherche* 307 (March): 84–85.

Pitts, M., A. Martinez, and S. A. Hillyard. 2011. "Visual Processing of Contour Patterns under Conditions of Inattentional Blindness." *Journal of Cognitive Neuroscience* 24 (2): 287–303.

Tong, F., K. Nakayama, J. T. Vaughan, and N. Kanwisher. 1998. "Binocular Rivalry and Visual Awareness in Human Extrastriate Cortex." *Neuron* 21 (4): 753–59.

Tsuchiya, N., and C. Koch. 2014. "On the Relationship between Consciousness and Attention." *Cognitive Neurosciences* 5. Forthcoming from MIT Press in a volume edited by M. Gazzaniga and G. Mangun.

Further Reading

Block, N. 2011. "Perceptual Consciousness Overflows Cognitive Access." *Trends in Cognitive Sciences* 15 (12): 567–75.

Kandel, E., H. Markram, P. Matthews, R. Yuste, and C. Koch. 2013. "Neuroscience Thinks Big (and Collaboratively)." *Nature Neuroscience* 14: 659–63.