The Myth of Double Transduction

Daniel C. Dennett

Do you remember Woodstock? I now know what it was like to be a cop at Woodstock. Even at Woodstock there were good cops and bad cops and I would like to be remembered as a good cop—one of those who went around saying, "Smoke what you like and have a good time and just try not to hurt anybody; it'll all blow over after a while. It may take 20 years, but go and let a thousand desert flowers bloom." The Tucson Conference was a great opportunity to expose to the bright desert air a great variety of ideas—most of which will turn out to be wrong—but that is the way we make progress on this topic.

My own view is skeptical. All through my youth I considered myself a sort of radical, but now I see that I am actually a sort of conservative. I say we do not need scientific revolutions. We have all the materials we need at hand—ordinary normal science. The same science that can explain immunology and metabolism and volcanoes can explain consciousness. What I am seeing is what we might call the "heartbreak of premature revolution."

Consider the cartoonists' convention of the thought balloon or thought bubble, a fine metaphorical way of referring to somebody's stream of consciousness. It is vivid and maybe even accurate and precise, but it is a metaphor for what is actually going on *somewhere*—I say in the brain, but others may want to locate the activity elsewhere. The problem of consciousness, as I understand it, is this: If a thought balloon is the metaphorical truth, what is the literal truth? What is actually happening for which a thought balloon is such a fine metaphorical representation?

The temptation is to view conscious observation as a very special sort of transduction. Transduction is a good term because it cuts crisply across the artificial and the natural. We have artificial transducers, such as photocells, and natural transducers, such as the rods and cones in the retina. The transducers take information in one medium and, at a boundary surface, transduce it; the same information is sent on in some other physical medium—by turning photons into sound or by turning photons into spikes of electrochemical activity in a neuron's axon. There is a temptation to think that consciousness *is* a very special sort of transduction—the sort that we call observation.

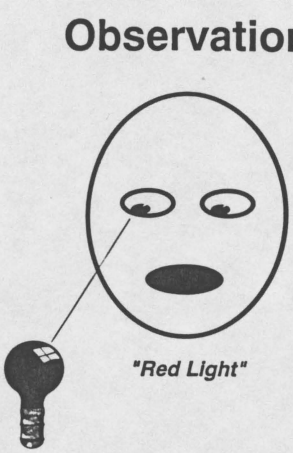

Figure 7.1 A conscious observer reports awareness of a red light.

Figure 7.1 shows a conscious observer seeing a red light and letting us know that he's seen it by saying "Red light!" Some people think that consciousness is a fundamental division in nature, which divides things that are conscious from things that are unconscious. The things that are conscious (sentient) engage in this very special sort of transduction. Of course, just saying "Red light!" under this condition does not guarantee that a person is conscious. After all, we might have an experimental subject who stands all day saying "Red light, red light, red light, red light." The fact that after someone flashes a red light in the subject's eyes, the subject says "Red light!" is no indication that the subject is conscious of the red light.

The implication is just as dubious if we think about other things that the subject might do. Suppose a person is driving along the highway and a red stoplight appears. The person's foot descends on the brake pedal, but does that response give us evidence of the driver's consciousness? Jeffrey Gray (Chapter 25) talks about this sort of issue in his example of playing tennis. His view is clear: Reacting appropriately to such a stimulus would not show consciousness. We could readily design a simple AI gadget—not even near the cutting edge of AI—that could respond to a bright red stoplight by pressing on the brake pedal of an automobile or respond to a tennis ball in flight by swinging a racket. There certainly would not have to be conscious transduction; there would not have to be *any* special observation going on in such a case. That is the way we are inclined to think about consciousness—that it is a very special sort of transduction. Consider a diagram from Frisby (1979).

Figure 7.2 illustrates one of the ways people convince themselves that there is a special sort of late transduction event occurring in the brain, a transduction event which *is* consciousness. In the figure, light is transduced at the retinas and then (in the medium of neuronal pulses) the information

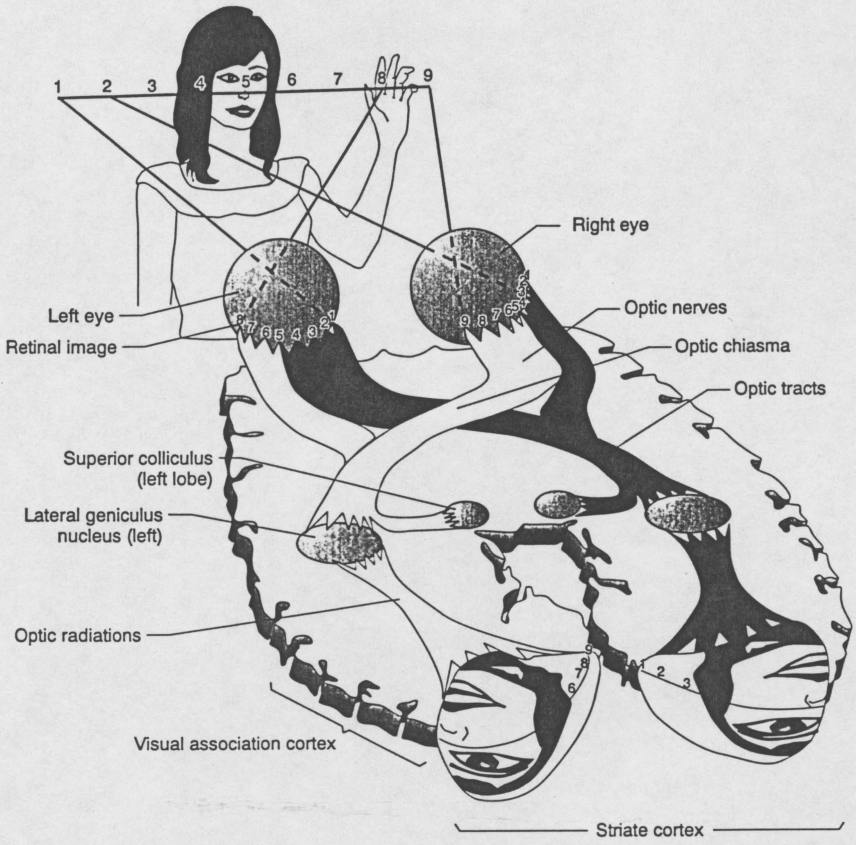

Figure 7.2 Transduction of a visual image from retina to striate cortex. (Adapted from Frisby 1979, by Dave Cantrell.)

works its way back via the lateral geniculate nucleus to the cortex. As Frisby's illustration vividly shows, by the time the information gets back to the occipital cortex, to visual area 1 (V1), the information seems distorted. Somebody who looked at your cortex while you were looking at the woman in the figure would not see the image depicted; nevertheless, the patterns of stimulation on the cortex would be approximately as shown—distorted, inverted, and twisted. A natural reaction is to say, "Well, that's interesting, but that's surely not how it seems to *me* when I look at the woman—so I guess the *seeming to me* must happen at some later point in the process. There must be some place in which *how it seems to me* is restored or put together in some later transduction."

The transduction at the retina, into neuronal pulses, seems to have taken us into an alien medium, not anything we recognize as the intimate medium with which we are familiar. "The activity in V1 is not in the *Me*dium," one might say. "It may be a medium of visual information in my brain, but it's not "moi." It's not the medium in which I experience consciousness." So the idea takes root that if the pattern of activity in V1 looks like *that* (and that is

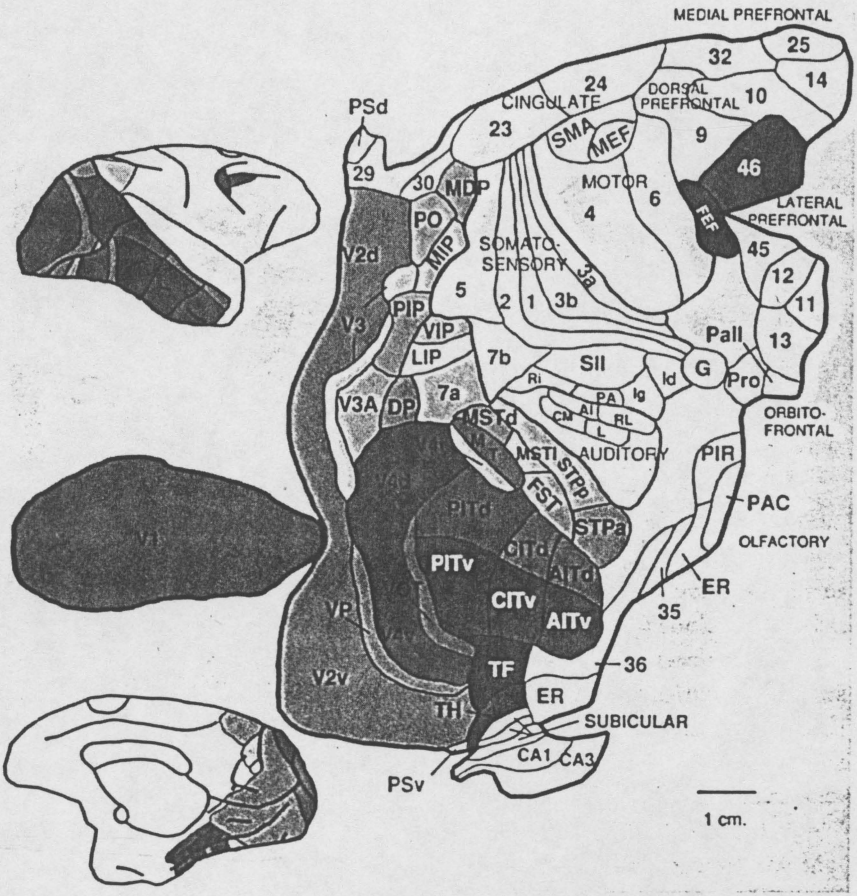

Figure 7.3 Visual areas in the cerebral cortex of the macaque monkey. (Adapted by Dave Cantrell with permission from D. C. van Essen.)

not what consciousness is like), a later, second transduction into the medium that *is* consciousness must occur.

Figure 7.3 is David van Essen's famous color-coded diagram of the visual areas in the brain of the macaque monkey. The big area on the left is V1. Another famous area is the motion-serving area (MT). V4 is where color is the main topic. Other areas deal primarily with shape and location. All the different aspects of vision are parceled out to specialist areas in the brain. Looking at a map like this, one is inclined to wonder whether there is some area on this map (or in some more covert area not yet drawn, deeper in the system) where "it all comes together" for consciousness—the place in which the soundtrack is put together with the color track, which is put together with the shape track, the location track, the motion track, the smell track—everything brought together into a single, multimedia representation: the representation of one's *conscious experience*.

If you are tempted to think that way (and if you are not, I think you are a very rare individual) you are making a fundamental mistake, the mistake I

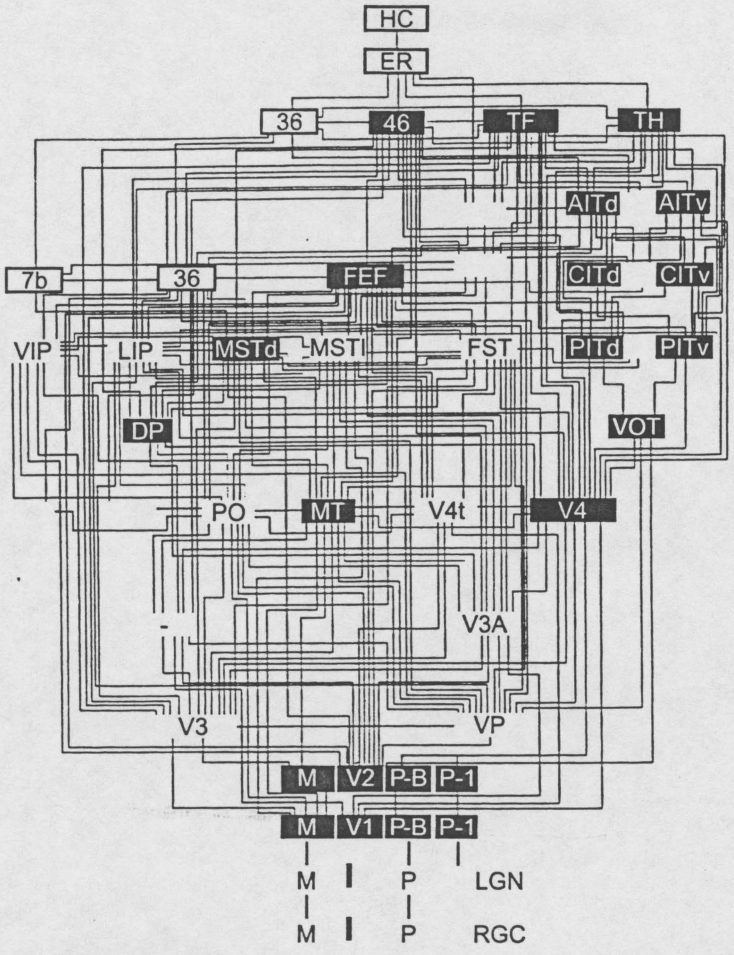

Figure 7.4 Hierarchical connections among visual areas in macaque cerebral cortex. (Adapted by Dave Cantrell with permission from D. C. van Essen.)

call *Cartesian Materialism*. This idea is that there is a second transduction somewhere in the brain (that is why this is a form of materialism). The idea is that a privileged medium exists in the brain where *and when* the consciousness happens, a place in which all the various features of a conscious experience become "bound"—and then, most importantly, are appreciated.

Figure 7.4 is another of van Essen's diagrams. It shows some of the connectivity between the areas in Figure 7.3, but there are no arrowheads on the connecting lines because arrowheads would be redundant—there would have to be arrowheads on both ends of nearly every connecting line. In the illustration at least as many pathways go down to V1 and down to V2 from "higher" regions as come up from the eyes. The very tempting idea that one can move up, up, up from the sense organs at the periphery to "the summit" or in, in, in to "the center" must be abandoned when you get to V1 because

there you are already "home"! There is no more central headquarters toward which input moves in the brain.

There might have been. People ask whether my theory is an empirical theory or an a priori theory. It is both. It *might have been* the case that there was a little guy who sat in a little house in the middle of your brain, who looked at many screens, who listened to many speakers, and who pushed many buttons. There *might have been* what I call a "Cartesian Theater" in each of our heads. It is an empirical discovery that there is no such thing. If there were, however, we'd simply have to start our theory all over with the homunculus: What happened inside *his* brain (or whatever occupied that functional role in him)? The empirical side of my theory is quite uncontroversial, though not trivial because empirical researchers often stumble over it: there is no Cartesian Theater in the human brain. The conceptual point of my theory is that at some point, one must get rid of the Cartesian Theater, get rid of the little guy sitting in the control center. As soon as that happens, one must change the assumptions of the theory in a rather dramatic way.

Figure 7.5 is the Cartesian Theater (in an old spoof that first appeared in *Smithsonian* magazine many years ago). We can all laugh at this; we all know this picture is wrong. The tough question is: With what do we replace it? Exactly what is wrong? It is not that the little observers are wearing white coats or that they have arms and legs or that there are two of them. What is wrong is actually much subtler. We know what is wrong, but we have not yet come to grips with all the implications. Here is what is wrong:

The work done by the homunculus in the Cartesian Theater must be distributed in both space and time within the brain.

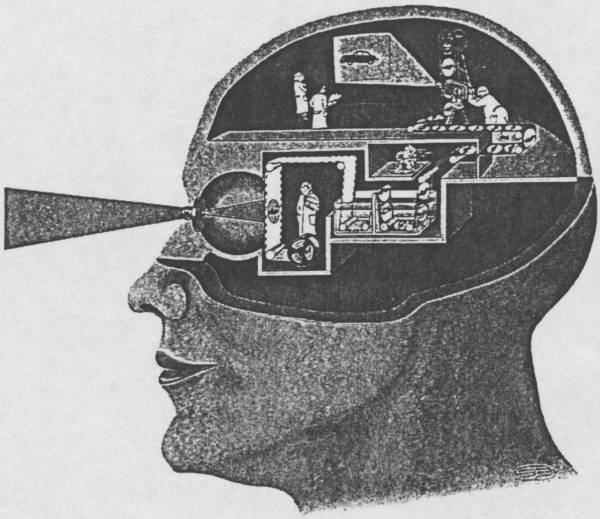

Figure 7.5 The "Cartesian Theater." (Adapted by Dave Cantrell with permission from Sally Bensussen. Originally published in *Smithsonian* 16, no. 1 (April 1985): 97.

There is widescale agreement on this view. But there is also a lot of backsliding; understanding this is no easy task.

I want to comment at this point on the word "work." David Chalmers might say:

> Yes, the *work* done in the Cartesian Theater—all that *functional* work—that's all distributed around in various parts of the brain just as you insist. Getting clear about all that *work*—those are the Easy Problems. But after we've figured out how all that work is done (and distributed it around in the brain) we won't have touched the interesting, Hard Problem—which is the *play*, the fun, the qualia, the subjective experience itself (which of course we want to distinguish from all that *work*).

However, the work that must be distributed is not just pattern recognition, discrimination, locomotion, and the like. It is also the generation of feelings of disgust or delight; it is the appreciation of the scene. Getting the appropriate emotional reactions is just as much part of the work that has to be distributed within the brain as controlling the stroke of a tennis racket. If we recognize that the work done in the Cartesian Theater has to be distributed in space and time, we have an answer to Jeffrey Gray's problem about the tennis player. It is simply a mistake to suppose that first all the work is done in the brain and later comes the consciousness (ohmigosh, too late to matter!—as if consciousness itself were something that happened too late to control the tennis stroke.)

Consciousness itself is distributed in time, not just in space, and there is no good reason not to say that part of one's *conscious* reaction to the tennis stroke starts within 10, 15, 20, 50, or 100 milliseconds of the arrival of the visual information at your cortex. It is a mistake to suppose that one must wait 100, 200, or 500 milliseconds for a *finished product* to be created in your consciousness, a building process that takes too much time, so much that it will be too late for *you* to do anything about it. We fall into this error if we suppose that there is still a task to be performed in the Cartesian Theater—a task of conscious appreciation and decision in addition to the processes that have already begun to take effect thanks to the distributed activity in the brain.

The task of deciding when and how to hit the tennis ball has to be distributed in time. One should not make the mistake of withholding the honorific of consciousness until all the work is done. (This point has often come up in discussion of Benjamin Libet's results. All his evidence for how long it takes to become conscious of something—what he calls "rising time to consciousness"—could just as well be taken as evidence for what we might call "curing time in memory." There can be only arbitrary grounds for taking some point in that interval of several hundred milliseconds and declaring it to be the onset of consciousness. Consciousness does not *have to have* an onset measurable to the millisecond; it is much better to think of consciousness as distributed in both space and time.

Figure 7.6 is one of Vesalius' wonderful anatomical drawings of the brain. In the middle, marked L, is the pineal gland, the epiphysis. But for the

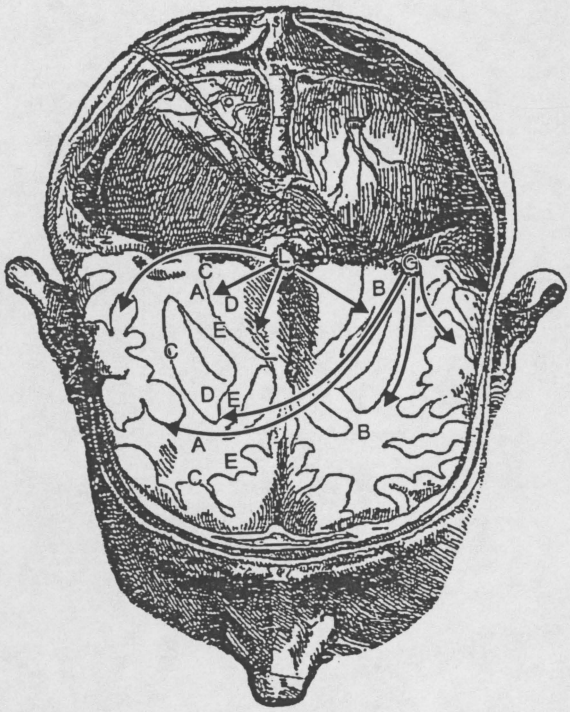

Figure 7.6 Anatomical drawing of the brain. (By Vesalius; adapted by Dave Cantrell.)

moment consider Figure 7.6 a picture of the globe, and let me use it to tell a rather different story—the story of the war of 1812 and the battle of New Orleans. The war officially ended on Christmas Eve, 1814, at Ghent (marked G in Vesalius' diagram!), when the British and American ambassadors signed a treaty. The news that the war was over traveled out from Ghent in all directions, at various rather slow speeds. The news arrived in London in hours or maybe even a day or two after the signing. It took weeks for the news to get across the Atlantic to New Orleans, where, in the meantime, British and American troops fought a battle in which several thousand British troops were killed. The battle was needless in one sense because the war was officially over; the truce had already been signed. But neither army knew that. It may have been even later when the people in Calcutta or Bombay or Moscow or Hong Kong learned about the signing of the truce.

Ask yourself a weird question: When did the British Empire learn about the signing of the truce? You may protest that the British Empire was not the sort of thing that could be said to *learn* anything. But that is not really true. In a certain important sense, the British Empire was a person, a legally responsible agent, a framer and executer of intentional actions, a promise maker, a contract signer. The Battle of New Orleans was a legitimate, intended activity of an arm (almost literally) of the British Empire, fighting

under the British flag. It was not some sort of renegade action. The signing of the truce was another official action of the British Empire, executed by one of its proper parts.

Suppose historians could tell us the exact day or hour or minute when every official of the British Empire (every proper part of the Empire, one might say) learned of the signing of the truce. It is hard to imagine that any *further* facts could be relevant to the question of when the British Empire "itself" learned. But then we can readily see that we cannot date the onset of this recognition, the onset of consciousness *by the British Empire*, any more finely than by specifying a period of about 3 weeks. When did the British Empire learn? It learned during the period between Christmas Eve 1814 and several weeks into January, 1815. One might be tempted to disagree, saying that what mattered was when the news reached London (or Whitehall, or the ears of George III) (*L'Empire, c'est moi?*). After all, was not London the headquarters of the Empire? In this instance, that's particularly implausible because when George III learned anything was hardly decisive! (Not so many years ago many of us were going around chanting "What did he know and when did he know it?"—but that was Ronald Reagan, and in that case, too, the question was of merely ceremonial interest. It did not really matter much when either of these "rulers" learned things.)

The point is this: Unless one is prepared to identify some quite specific subregion of the brain as *headquarters*, the place where "moi" is, so that entry into that charmed precinct is what *counts*, one cannot, logically, make precise determinations about *when* consciousness of one item or another happens. (In a discussion of Jeffrey Gray's presentation, a number of questions brought out that his "subicular comparator," although perhaps a very important crossroads in the brain, would have to function by sending out appeals here and there. The system under discussion was not just the hippocampus and the subiculum. It is only in concert with other large and widespread parts of the brain, as Jeffrey eventually acknowledged, that the effects that matter could be produced. It follows from this that one can only arbitrarily point to *any* one stage in this process and say *that* is where the consciousness happens—or more important for present purposes—that is *when* the consciousness happens. (*When* various vehicles of information arrive at the hippocampus is an interesting fact, but it does not settle anything about the timing of consciousness.)

Here is the moral of my story so far:

Since *you* are nothing beyond the various subagencies and processes in your nervous system that compose you, the following question is always a trap: "Exactly when did *I* (as opposed to various parts of my brain) become informed, aware, conscious, of some event?"

It is easy to see reasons why people are tempted by the hypothesis of a second transduction. In some regards, the computer revolution has made this mistake easier. We now have many systems that are *media neutral*. Consider the steering system of a modern ship, in which the helm is attached to the

distant rudder not by pulleys and wires or cables or chains but by a couple of little electric wires. Information is all that is carried by these wires—information about the position of the helm and feedback from the rudder. The wires could readily be replaced by any medium that carried the same information (glass fibers, radio waves, etc.). Such familiar examples give us a clear understanding of how one can have information transmission in, essentially, a media-neutral way between the "Boss"—the steerer, the governor, the cyberneticker—and the rudder. As long as the information is preserved, what the medium is does not matter.

However, when we think of information transmission in neurons, mere nerve impulses, it seems, just cannot be the medium of consciousness! It doesn't seem to be enough. Somehow, it seems, we have to put the helmsman back in there. It seems that we have to have a boss or an appreciator, some more central agent to be the audience for all that information. Otherwise, the nervous system seems to be a phone system with no subscribers. There is nobody home to answer the phone. There is a television cable network without any viewers. It certainly seems as if we need to posit an audience somewhere to appreciate all that information—to appreciate it in a second "transduction."

The alternative to this bad idea takes some getting used to. The alternative is the idea that the network *itself*—by virtue of its structure and the powers of transformation that it has, and hence its capacity for controlling the body—could assume all the roles of the inner Boss and thus harbor consciousness. That idea at first seems preposterous to many people. Both David Chalmers and Michael Lockwood remarked in their sessions in Tucson that although they acknowledge that there are people who maintain this view, they think it is simply a nonstarter. That "the subjective point of view" can somehow be captured in the third-person resources of the structure of this functional network strikes them as inconceivable. To me it is not. When people declare to me that they cannot conceive of consciousness as simply the activity of such a functional network, I tell them to try harder.

A common reaction to this suggestion is frank bewilderment, expressed more or less as follows:

OK. Suppose all these strange competitive processes are going on in my brain, and suppose that, as you say, the conscious processes are simply those that win the competitions. How does that make them conscious? What happens next to them that makes it true that *I* know about them? For after all, it is *my* consciousness, as I know it from the first-person point of view, that needs explaining!

That question, I think, betrays a deep confusion; it presupposes that what *you* are is something *else*, some Cartesian *res cogitans* in addition to all the brain-and-body activity. What you are, however, *is* this organization of all the competitive activity between a host of competences that your body has developed. You "automatically" know about these things going on in your

body because, if you did not, it would not be your body! The acts and events you can tell us about, and the reasons for them, are yours because you made them—and they made you. What you *are* is that agent whose life you can tell about. For me the "hard problem" is getting people to see that once you have solved the "easy problems," you have solved them all—except for *my* hard problem, which I am continuing to work on as you can see!

REFERENCES

Frisby, J. P. 1979. *Seeing: Illusion, Brain and Mind.* Oxford: Oxford University Press.