

WHY CONSCIOUSNESS IS NECESSARY BUT SCIENCE CAN'T EXPLAIN IT

© 2017 Dan Bruiger
dbruiger [at] telus [dot] net

Why consciousness is necessary is a question that only the modern mind could ask. We must clarify at the outset what is meant by 'consciousness' and in doing so clarify what is meant by 'science'. Then we will attempt to show why consciousness exists at all and why science still cannot explain it.

I can remember as a young child noticing certain things about perception: for example, that my two eyes saw color slightly differently. Or that food can taste very different depending on how hungry you are. Such experiences led me then, as they probably did you, to conclude that perception is not a straightforward matter of exposure to the world, but depends also on oneself. There is a subjective component to how we experience what may appear to be objective reality; and there are experiences, like dreams, that seem unbound to the real world. Yet, one does not come to a full-blown concept of *subjective consciousness* until adolescence, when the ability to think abstractly comes into play in the individual's development.

Similarly, it may be that ideas about consciousness or mind could only enter human culture at a certain stage of development. Society has long benefitted practically from the psychological ability to distinguish subjective from objective. To recognize that people can see things differently, and that one's own perceptions have no absolute claim on truth, enables cooperation within groups and tolerance of other groups. Yet this recognition of subjectivity is not the same as a theory of how the mind works that can explain the relationship of subjective experience to objective reality. Not every culture or age has thought it necessary to *explain* consciousness at all, nor made such a clear distinction between subject and object. The modern scientific culture, however, is founded upon it.

A further complication is that many words pertaining to our psychological life have ambiguous or double meanings. In particular, they can refer either to a first-person or to a third-person perspective. This distinction itself is reflected in the grammar of language, and is a product of our peculiar situation as beings aware of being aware. In English, the very term 'experience' can mean either the sensory awareness of here and now, or else a state of knowledge or memory about something that has happened (as in 'work experience' and 'traumatic experience'). We can think of experience either as a present inner state of consciousness or as outer things that have happened to us. Since we have material bodies, we can think of ourselves and of others as subjects or as objects. We can imagine ourselves or others both from the inside and the outside, so to speak. And this raises the question of how inside relates to outside. Is one somehow the product of the other?

The first challenge in talking about the so-called Mind-Body Problem is to identify just what the problem is. While most people make regular practical use of the distinction between subjective and objective, not everyone is concerned to understand their relationship in a theoretical way. To live our lives, we do not need to know how the

brain creates experience. Day-to-day experience is mainly about negotiating a path through the world as we take it for granted. Only under unusual circumstances do we dwell on the nature of experience itself—or mind. But this is the very job that falls to philosophers—and, more recently, to scientists. The MBP is exactly this question of the relationship between inside and outside, subject and object, first-person and third-person perspectives, conscious experience and brain.

Another thing that makes consciousness hard to talk about is that language provides only figurative ways of speaking. Like other fundamental questions, we have only metaphors through which to understand consciousness. Except physiologically, there is no “inside” and “outside.” We speak of the external world, but external to what? Certainly, there is a world outside the skull, but the brain is part of that world. In some sense, therefore, even the brain (and certainly the rest of the body) is external to the perceiving subject. The relationship between subject and object is not between two things in the world, but between a point of view and the view from that point. ‘Point of view’ is both a visual and a linguistic metaphor. First-person and third-person are elements of language, which tells us how much our views of anything, even consciousness itself, depend on language.

All this to underline that ‘consciousness’ is an ambiguous term. It can mean the actual here-and-now fact of being aware of something—whether sensation, dream, emotion, thought, imagination, or hallucination. For clarity I will call this awareness *phenomenality*, in preference to other possible locutions that seem ambiguous (such as ‘experience’) or else redundant (such as ‘conscious experience’, ‘subjective experience’, ‘conscious awareness’, ‘phenomenal awareness’, etc.) But, secondly, ‘consciousness’ can mean a specific *function* within the human organism, which is different from non-conscious cognitive activity that can take place without it. I intend to explain what this function does and why it is necessary, at least for the human organism, and how it requires the subjective state of awareness we call consciousness, or conscious experience, or phenomenality.

The *need* for any such explanation is modern and peculiar to Western culture—especially the attempt to explain phenomenality as a state or product of the brain. Since the scientific conception of matter is of something essentially inert, mechanical, and objective, it then becomes particularly troublesome to grasp how a blob of such matter can produce the panoramic subjective “show” of phenomenality.¹ From the get-go, the quest for a scientific explanation of consciousness rather shoots itself in the foot. This is because science has exiled the observer’s subjectivity in order to better study the objective aspects of nature; but then it is stymied when it tries to understand subjectivity itself in objective terms. For a long while, science was only willing to consider mental *functions* performed consciously, without inquiring how they require corresponding *experience*. Functions could be described in terms of observed *behavior*, but there was no comparable way to understand the associated *phenomenality*, which seemed to be causally superfluous. It seemed possible to understand the world, but impossible (and perhaps irrelevant) to understand the appearance of the world in consciousness. To do science, the observer stands outside the system observed. Yet, the observer is also a part

¹ In the expression ‘mind-body problem’, ‘mind’ means that show and ‘body’ means that matter.

of the nature observed, an embodied organism. The problem is that the observer is both subject and object.

The ancient Greeks understood the distinction between subjective and objective, but lacked a clear basis on which to found a scientific theory of mind.² (As proto-scientists, they were more concerned with a theory of matter.) In the Christian culture of medieval Europe, there was no need for a theory of mind, because both reality itself and its appearance in our consciousness were provided directly by God. Similarly, in the East, Atman and Brahman were distinguished, yet entailed each other without separating mind from matter. Buddhism developed an elaborate theory of perception, but with no relation to the brain as an organ responsible for it. In pre-scientific times generally, where an *agent* was deemed responsible for perception and thought, or to account for the presence of the subject, that agent was considered the *soul*, an indestructible unit of consciousness. Most religions and eastern philosophies are inherently idealist, in the sense that *idea* (that is, consciousness) is considered to be what fundamentally exists and so requires no explanation. The western secular tradition that evolved from the Greeks parted ways with these traditions, to consider the *material world* fundamental. We have inherited that materialism and the science based on it, and therefore also the problem of explaining consciousness in material terms. Here I embrace that problem. Part of what I propose is to explain why consciousness is so challenging within the scientific worldview, and how the latter can be revised to accommodate it without compromising its positive features.

How should we go about inquiring what consciousness is? Indeed, what makes us think there is any such thing, as distinct from the world that appears *in* it? And if it is a *thing*, what is it made of and where does it exist? If it is not material, what is it, then, and how does it fit within the material world?

As children and even as adults, we mainly relate to experience as though it were an open window on the world. That is, most of the time the world simply *is*, without any thought of being aware. This is normal, because the very nature of consciousness is to efface its own tracks, so that as subjects we simply dwell and navigate in the world of objects without any reason to focus on *being* conscious. That is how nature designed us, a feature that becomes part of the task of explaining consciousness. We could call this state of affairs *naïve realism*.

Yet, there are specific clues that there is more going on than the simple presence of the world. There are cues within the content of experience that someone is doing the experiencing: there is the world, on the one hand, and there is our consciousness of it, on the other. There are objects, but there is also a subject. We look out upon the world as through open portals, and we see our own bodies within that visual field. The visual field actually has an outline—a literal frame—formed by the bone structure in which the eyes are embedded. Normally one focuses on the world beyond the eye sockets, but upon occasion the realization that *someone is looking* intrudes literally as plain as the nose on your face! While that is a visual cue, the *subjective frame* is any experiential cue to one's own existence as a subject. Such experiences as hallucinations, dreams, perceptual illusions, after images, sensory adaptation, imagination, memory, and emotion tell us not

² Some early Greek philosophers identified sensation and thought with the brain, others did not. With his metaphor of the cave, the closest to come to a representational theory of mind was Plato, whose intuition was furthered centuries later by Kant.

only that we are subjects, and that there are objects of experience that are not part of the physical world. They tell us that phenomenality exists in its own right, and is more inclusive than an open window on the world. Once one grasps the fact of being a conscious subject, it is natural enough to found a concept of mind upon that realization and a worldview based on subjectivity, even to think that the mind is what primarily exists and that the material world is no more than an appearance within it. The belief that nothing “really” exists but such appearances could be called *naïve idealism*.

If it is all in the mind, like a dream, why is there a communal delusion of a real material world? On the other hand, if the world is purely material, what place is there in it for feelings, thoughts, imagination, dreams, hallucinations? Is everything mind or is everything matter? (For some reason, humans prefer that things be all one way or another, and not confusingly both.) Accordingly, philosophers tend toward either a worldview in which matter is the primary reality or one in which mind or consciousness is the primary reality. Unfortunately, either extreme leads to a restricted perspective that cannot reconcile apparently conflicting aspects of human experience. Our very nature as self-conscious beings renders dualism of some sort inevitable. We are aware of an external world but *also* of an internal one. Hence, no one-sided approach can explain either consciousness or the world to complete satisfaction. Phenomenality, thought, and behavior are all *co-determined* by the organism and its environment, which means co-determined by subject and object together. One or the other aspect may be favored for various reasons, but neither can be ignored.

In the materialist tradition, mind is explored indirectly through observing the *behavior* of organisms, which are considered material systems, just as the physicist observes the behavior of non-living matter. Until the last quarter of the twentieth century, science generally refused to consider phenomenality or consciousness as a topic for discussion. For one thing, it was awkward enough to consider the subjective experience of other human beings, let alone of other creatures. It is now recognized that consciousness must serve some purpose, at least for the human organism. But, for a long while it was considered literally a superfluous afterthought; the real causal processes were material, taking place in the nervous system, of which phenomenality was a useless byproduct. Two modern developments in technology made it feasible to consider consciousness fit for scientific study: brain scanning and stimulation, on the one hand, and computer modeling on the other. The computer is not only a mathematical tool but serves also as a powerful metaphor for understanding how the mind works.

If we admit co-determination of subject and object, then we must wonder what sort of agent a subject is, if it is not merely another part of the physical causal system. What is the nature of this “agency” if it is to account for phenomenality? The mind-body problem is that phenomenality does not seem to be material, yet we are trying to account for how it is caused within a material system. In desperation, some deny that phenomenality even exists, and that there are only material causal processes at work. Others deny that it arises in the brain. I will take a different tack: phenomenality is produced within a *virtual* system (which happens also to be a material system), and the type of causality involved is not that conventionally of physics (which has no place for any agency but that of the physicist, nor any causality but “efficient” causation). I will consider *intention* to be the sort of cause that agents initiate and which underlies consciousness.

Many human experiences do not fit well within the materialist worldview. All sorts of “paranormal phenomena” (out-of-body and near-death experiences, memories of past lives or of alien abductions, ESP, etc.) resist scientific explanation and may seem more compatible with an idealist worldview. However, I do not believe that consciousness *per se* falls into this category. In the idealist worldview, there may seem no *need* to explain consciousness at all, since it is considered the fundamental reality from which all else is derived or composed. In such a view, some principle of consciousness does our perceiving for us. But such an arrangement reminds one of the character in Moliere’s play who attributes the power of a sedative to a *dormative principle* it contains. Of course, this does nothing by way of explanation. Similarly, it does nothing to attribute consciousness to the activity of some conscious entity within the person, whether a soul or a “spiritual body.” This simply defers the problem. Moreover, idealism is not free from the need to explain the apparent materiality of the world, the existence of nature as something that pre-exists human beings, and obvious experiences of being embodied—such as pain, injury, mortality, and having a limited perspective in space and time. If the nature of consciousness is not material, then why does it dwell in material bodies? Why is there a material world at all? Many answers have been proposed to such questions and entire systems of thought constructed around them. Since there is no accounting for taste, I will not argue against idealism, but will simply adopt the materialist worldview as an exercise and embrace the challenge of explaining consciousness within its framework.

To succeed at that challenge, one must explain consciousness in terms of material processes that are at least potentially observable. But let us distinguish the *actions* of an agent from the passive behavior usually attributed to matter. Even so, we fall into the trap alluded to above if we consider this agent to be already a subject capable of perceiving. For, if it is this inner agent that does the perceiving for us, then we must explain *its* perceptions in turn. Does that agent then defer to another agent within it? If so, where does it end? A view of perception in which the “soul” or “inner agent” looks out through the portals of the senses as through an open window may be a step up from naïve realism. But then we must explain how this soul or inner agent perceives, with a potentially infinite regression of agents within agents.³

I will show that there is no need for a regression of observers within observers, or agents within agents, or for an inner theater in which the contents of consciousness are literally displayed. For, there is but *a single* inner observer: the conscious self. The dilemma of regression results from passing the buck from one agent to another, thereby failing to explain how the subject perceives at all. The question is *how* does an agent manage to perceive, and what role does phenomenality play in that process? If we believe that phenomenality arises in the brain, then we must know what the brain *does* that constitutes perception and results in phenomenality.

We can note a few obvious facts. First, regardless of what may or may not be going on in the brain, we do indeed experience the appearance of a world outside the skull—a world that includes the skull and the body. In other words, experience is

³ Perhaps this dilemma comes of considering the perceiving subject to be as passive as matter is assumed to be: merely a witness to objective reality rather than an active participant in creating experience. Science has simply institutionalized these cultural assumptions.

projected outside the brain. Second, the image on the retina is upside down with respect to the body, but we nevertheless experience the world as right-side up. So, neural processing makes an appropriate *adaptation* that presents the world as it “actually” is in relation to the body. Third, we note that the brain is indeed sealed within the skull in such a way that its only connection to the external world is through electrochemical signals transmitted through nerves (and perhaps through chemicals in the blood supply). Perception must involve actively *interpreting* such signals. In particular, it cannot be a matter of passively witnessing an image, whether on the retina or somewhere in the brain. There is therefore no need or place for an inner theater. If some kind of inner representation turns out to be useful, it is not a literal image transmitted by the senses and projected in the brain for some ready-made observer to passively look at. Rather, it is a device by which the brain accomplishes looking. An internal representation or *model* of the external world is not a literal copy but a schematic map. There is no need (or possibility) of a one-to-one relation between the model and the reality.

The isolation of the brain in the skull suggested to Descartes the possibility that inputs to the brain could be faked, giving it the delusion of being in a body in the world. This situation has been popularized in science fiction as the “brain in a vat.” A human brain has been removed from its donor and kept alive in a chemical bath. Electrodes are implanted in it to provide the same sort of signals it would normally receive from the senses; but it is a computer that supplies these to simulate experience in a real world. Descartes’ concern was that we have no way to prove we are not being deceived by the input of our senses. His way out of the dilemma was to argue that God would not sanction such deception. However, a more modern solution is to realize that there are not only inputs to the brain but also outputs from the brain to the body, which interacts with the world. While the brain is sealed within the skull, the skull is part of the body, which is part of the world.⁴ The brain is engaged in a sensory-motor cycle of self-regulation, such that motor output changes the world, which changes the sensory input, which changes the motor output, which changes the world... etc. The challenge to the brain (its job) is to regulate inputs and outputs in such a way that the body (with the brain in it) survives. If conscious experience is in some sense illusory, it is not a lethal deception, which by definition *nature* would not sanction.

While the brain is a black box to an outside observer, the outside world is a black box to the brain. The observer can open the brain surgically to observe its structures (as well as observing its inputs and outputs) in order to try to deduce what goes on inside it. The only way that the brain can open the world, however, is indirectly through the actions of the body that intervene in the world and cause new inputs. This remains a task of interpreting inputs, whether from the senses or from instruments. (Scientific experiments are such interventions.) The points of view of the brain and of the outside observer are essentially different. This becomes especially tricky when the brain of the observer happens to be the object of study and also the subject who is doing the studying.

From the observer’s point of view, there is clearly a world outside the brain. The brain also may believe that there is a world outside it, but this “belief” is strictly an

⁴ An analogy would be if the brain in the vat had electrodes connected to a robot, so that the brain not only received signals from the robot’s sensors but also issued motor signals to control its behavior. This is the situation the normal brain is actually in.

incidental (and functional) aspect of its primary obligation, which is merely to auto-regulate a balance of inputs and outputs to maintain stasis within a zone of viable conditions. From the brain's point of view, the outside world is a pragmatic fiction, a theory. Since the observer has a brain, she may consider that the world that seems so obviously real, and which contains brains such as her own, is likewise her own brain's pragmatic fiction. I call this paradox the "problem of cognitive domains." It arises when the output of a cognitive system is recycled as its input. For example, the scientific observer tries to explain the appearance of the world in consciousness as an output of the brain. However, the input to the brain can only be conceived in terms of that output (that is, in the terms her brain presents to her as the appearance of the world). Precisely because this is an inescapable dilemma, which uniformly attends all forms of cognition and inquiry, it need not deter us from attempting to understand how the brain produces phenomenality in the course of its auto-regulation.

The brain's task can be likened to that of a submarine navigator or an aircraft pilot flying by instrument. In both cases, the task is to coordinate instrument readings with appropriate control outputs. (For example, to compensate changes in altimeter reading by moving the control stick.) The pilot or navigator has no direct view outside, only the inputs and outputs of electronic instruments and controls. While the human pilot or navigator has previously set foot outside the cockpit or submarine, this is not the situation for the brain. The human pilot can visualize air space, landing fields, possible obstructions, etc, because of prior experience. She has a mental picture of the external world based on such experience in terms of which to interpret instrument readings. The brain, however, must invent its picture from scratch. The brain has a theory, so to speak, about the external world, based on feedback in the auto-regulation cycle. The brain can test this model, within limits, by varying its outputs in relation to inputs and learn accordingly. In the overall scenario, however, the "theory" is tested by natural selection. The model is "true" (or at least adequate) if the body survives to reproduce.⁵

What relation must the model bear to the external reality? From the point of view of an observer, it might seem optimal that the brain's model should be an accurate copy of the external world. A perfect literal copy, of course, would be no more feasible than building an "exact" scale model of the universe! But such a replica is unnecessary, for in truth the only requirement is that the model should not hinder the survival and reproduction of the organism.⁶ In any case, the model will be some highly selective *mapping* between two domains. From the observer's point of view, this is a correspondence between structures in the external world and structures in the brain. From the organism's point of view, however, it is a correspondence between inputs and outputs. The observer might judge the correspondence in terms of accuracy and completeness. But this is really a comparison with the observer's model, not with reality itself. Natural selection judges it simply in terms of survival. The theory does not need to be true; it only needs to work. Accordingly, we must therefore redefine "accuracy" to

⁵ Other analogies might include self-driving cars and autonomous robots. The inputs and outputs of these systems are regulated by a computer instead of a biological brain. But there is still a sensory-motor loop and the challenge is essentially the same: to guide the vehicle safely.

⁶ In the self-driving car analogy, the requirement might be that the model should not put the company out of business!

mean functionality for survival. Because the human brain, even of the scientist, is in this very same situation, she faces the same dilemma: while one can imagine that there must be some way the world truly is—in its own right, apart from our cognitive modeling—one has no access to that reality except through that modeling. Like the pilot flying blind, the brain has access only to instrument readings and controls, to patterns of inputs and outputs summarized in its model, which helps us to see what we *need* to see to navigate successfully. The question then remains, how does the brain's cognitive activity result in "seeing" at all? How is it that we experience the brain's model *as* an external world? How does the firing of nerves result in the luminous experience of phenomenality?

The computer metaphor puts the programmer in the shoes of the brain, so to speak. If you were in the brain's position, how would you program yourself to coordinate inputs and outputs? An obvious first approach would be trial and error: try various (motor) outputs and see what new (sensory) inputs result. Does the action take you farther from equilibrium? If so, try the opposite, etc. Various rules of thumb might emerge from this exercise. These could form a simple program or even be hardwired into the system. Such an arrangement would amount to *implicit* knowledge of the environment. At some point in the exploration of increasingly complex situations and responses, it might be advantageous to codify that information explicitly in a model—a simulation of the outside world. The proposal here is that consciousness *is* this simulation, as presented to an executive agent, part of whose job is to monitor and oversee the organism's relationship to the outside. We will see shortly why this inner agent does not imply a regression and also why it *does* imply consciousness.

The brain's simulation is a virtual reality that is successful insofar as it is realistic—in other words, insofar as it facilitates the organism's survival. The concept of virtual reality, as it has been developed for recreational or training purposes, expands the computer metaphor: the computer can be used to produce artificial sensory inputs for a human user that simulate a real world. (This is like the brain in a vat scenario, except that the brain is not directly wired to the computer but retains its normal sensory and motor interfaces.) The metaphor is imperfect because the brain has no "goggles" to put on and the VR user has her own brain and senses with which to view the simulation. But we can extend the metaphor to the situation of the brain in the vat—that is, to the situation of the brain sealed within the skull, which has only the interface of electro-chemical inputs and outputs. Unlike the human user who witnesses a simulation created by someone else, say for commercial purposes, the brain creates a simulation for its own purposes. The metaphor aims to explain experience *as* simulation, not experience *of* a simulation. The challenge is to show how consciousness *is* the active operation of this simulation and why it is necessary at all.

It might seem that the VR metaphor implies the solipsism of the brain in a vat; but this is not so, because the experience that the brain's simulation provides is largely communal for the species and because the brain is not a passive subject but interacts with the world. Nor does it imply idealism, since the "dream" or "illusion" is directly guided by this interaction. Nor should one infer that the putative real world we interact with is nothing more than a simulation itself. For, by definition, a simulation is produced in and by a real computer (or brain). Even if a vastly superior alien civilization (or God) had created our universe as a simulation (for their entertainment?), there would still have to

be a real computer somewhere in which the simulation is produced. Logically, somewhere there has to be a real reality.

To show how a causal system such as the brain (or a computer) could produce the “guided illusion” of a real external world, we must first show that the causal system in question is a *virtual* system. The traditional scientific notion of cause is like the domino effect. Something *outside* a system disturbs it and the disturbance is mechanically transmitted through the system. In contrast, a system could be an *agent* that initiates a disturbance within itself, which may be transmitted beyond its boundaries. This sort of agent causation could be called *intention*. While a virtual system may be physical (as it must be if not merely conceptual), its processes and connections *as a virtual system* are intentional and logical, rather than causal. The wiring diagram for an electrical circuit illustrates the difference. It shows the *logical* or *functional* connections of the circuit rather than the points of soldered physical connections between wires. Yet, for the system to work as a real object, the circuit diagram must be realized in physical connections that correspond to the logical connections. Similarly, the brain is a physical system, but its operation as a virtual system must be considered in terms of its logical connections. The physical wires, transistors, and connections of a computer’s central processor are the hardware of a physical system; the computer’s operating system is the virtual system enabled by the hardware. Similarly, it is the virtual system of the brain that produces consciousness, even though the virtual system consists of physical nerve cells.

One difference between a brain and a computer, however, is that the computer was designed and programmed by human agents, according to their purposes. An organism has its own purposes (or those imparted by natural selection); its brain is, so to speak, self-programming. It has its own intentions, which (religious dogma aside) are not those of an external designer. We might seek to understand the brain as a *physical* system in terms of its physical connections (its neurophysiology). But to understand its phenomenality, and any but its simplest behavior, we must regard it as an agent in its own right.⁷ A physical system operates on causes; an intentional or virtual system operates on *reasons*. These are two different perspectives on the same system, corresponding to the difference between third-person and first-person description. To understand how the brain creates phenomenality, we must approach it as a virtual system from its own point of view.

An observer *discovers* causal connections; an agent *makes* intentional connections. Intentional connections exist by creative *fiat*, as in the divine declaration: “Let there be light!” Or by *hypothesis*, as in the mathematician’s: “Let *x* stand for... (such and such).” Just as a computer program results from original acts of the programmer, so the brain’s “assertions” are not *caused* by something outside it, but are its original acts, which constitute an internal language. While this “language” may seem like gibberish to the observer, it is precisely how the brain creates images and meaning for itself.

Another analogy is paint-by-numbers. Each “cell” to be colored in is filled according to neural code with what the brain “asserts” to be there. The degree of detail

⁷ In contrast to the notion of “efficient cause” in physics, an agent is an original or “first” cause. Efficient causation is one domino hitting the next in a chain reaction. But the chain reaction requires something to start it—an uncaused first cause, which in the literal illustration of the domino effect is the finger of a human being who *intends* to set the chain reaction into motion.

perceived is therefore relative and impressionistic. *Finer* details are simply smaller designated areas to be painted in with what the brain decides is there. The point is that the brain provides itself not with a realistic, indefinitely detailed portrait of reality, but simply with coded information displayed in a form useful to itself at the moment. It would be challenging and pointless to try to imagine the colorful scene implicitly represented in the numbered outline *before* it has been actually painted. The very point of the “painting” is to represent such coded information conveniently—for the benefit, we shall see, of a unique inner agent.

I propose that the basic dynamic behind phenomenality is this “painting in,” which is also demonstrated in the laboratory in various perceptual completion effects, habituation, perceptual adaptations, and phenomena of spatial projection. Just as the visual blind spot is filled over (in the sense that it goes unnoticed), the detail in the peripheral field of vision is *assumed* to exist to the same degree as it appears in the fovea. In effect, the brain *asserts* that it exists and sees what it asserts. But the same argument can be made all over again regarding the detail within the fovea itself: the brain sees what it believes is there. This “belief” is not arbitrarily made up; it is guided by real-time sensory input. But the way it represents that sensory input to itself is its own convention.

Why does it bother doing this “representing” at all? If phenomenality is a virtual reality show, what is the use of it? I propose that consciousness serves as an interface between an executive agent (the conscious self) and other parts of the brain’s virtual system. Its purpose is to allow this executive agent to readily monitor the state of the organism and its world and to enable *considered* action as distinguished from automatic reflex.

Many experiences are obviously associated with a cognitive judgment and a behavior. We tend to be attracted to things that are good for the body, which often feel, taste, smell, and even look inviting. Conversely, we tend to avoid things that are bad for the body and perceive them as unpleasant. The judgment of bad or good is associated with behavior of aversion or attraction, and both are associated with feelings or appearances that involve valuation. Hence, injury *must* hurt if pain is to serve any purpose. Sugar *must* taste good and rotten foods smell bad. It is clear that these conscious sensations are functional. But, what of visual sensations of color or auditory sensations of tone? What is the valuation or behavior involved in these experiences? Why do ripe apples appear red rather than some other color or no color at all? Why does an acoustic vibration sound as a tone?

Such questions lead us to first distinguish the distance senses from the proximal senses. Immediate contact with the skin has a very different implication than perception of a possible danger or opportunity in the far distance. Contact may elicit a reflex response in addition to a sensation (or even instead of it), while the sighting or hearing (or even smelling) of a distant predator or prey may elicit conscious alertness.

Let us consider the pain response in particular. There are two levels of response, traveling in the nervous system through different pathways and at different rates. With regard to an external stimulus, the first in time is a quick unconscious reflex of withdrawal, designed to avoid damage caused by the stimulus—for example, a hot surface. This is followed by a lingering conscious sensation of pain if there actually has been tissue damage. This serves an entirely different purpose from the reflex: namely, to avoid *further* or renewed damage and to insure that the injured part is protected and

avored during a prolonged period of healing. While the unconscious reflex is an instantaneous reaction produced by the external stimulus, the conscious sensation of pain is the result of ongoing nervous activity *initiated by the organism itself*. It therefore reflects an intention and an internal message to enforce a certain protective behavior toward the injured part.

But for whom is this “message” intended? The metaphor I prefer is the corporate boardroom. Various departments of the company meet together to report on their activities. But they must cooperate and act in concert to produce the decisive action that is often called for. There must be someone in charge, to finalize planning and decisions, and to take ultimate responsibility for monitoring the “market” and the internal and external relations of the company. This is the job of the executive agent: the conscious self.

The conscious experience of pain, as opposed to the pre-conscious reflex, is a separate response after the fact, which regards future rather than immediate behavior. We may generalize to say that conscious experience (phenomenality) serves a different purpose than unconscious processing in general. Much perception and behavior can take place “automatically,” as we commonly experience when driving a car or performing a rote task. Another way to put this is that existing algorithms can handle familiar or unthreatening situations. But let there be an emergency, and something “emerges” in consciousness. Following our metaphor, much of company business does not require the attention of the CEO—but some does, particularly novel business threats or opportunities, or internal disruptions within the company. This is why consciousness exists and is necessary. A corporation that only ran by the book and by precedent would tend to lack a competitive edge in a complex and changing economy.

One might yet wonder why this executive function could not operate with more sophistication than automatism but still unconsciously. Why could there not exist a fully functional human being who was simply “dark” inside and did not experience phenomenality? (In philosophical jargon, such a creature is called a ‘zombie’.) The answer has to do with the decisiveness required for considered action and the interface required for the monitoring that leads up to it. Sensory information always begs for interpretation, which must be packaged unambiguously for decisive action. One way or another, the situation must be perceived as definite and with clear implications. Hence, the recommendations of the board must emerge for the CEO in a definite perception of the situation, upon which action can be taken. Phenomenality is always definite. (This is illustrated by such ambivalent figures as the Necker cube; the brain’s interpretation switches from one gestalt to the other and back, but is definitely one way or the other at a given moment.) Phenomenality thus facilitates decisiveness. But it is also necessary for the monitoring and planning involved in considered action, which are functions of the distance senses and/or processes requiring time. Conscious sensations are how the world (including the body) is monitored spatially and temporally at the interface between the executive agent and the input from subsystems.

The *qualities* of sensations (the hurtfulness of pains, the hue of colors, the tone of sounds) constitute the internal language in which information—generally derived from the senses—is summarized for the executive agent on an ongoing basis. Contact with the surface of the organism (touch) can elicit an immediate unconscious reflex response. But,

like the *felt* pain, ongoing monitoring of the situation and planning for the future require the organism to initiate its own internal activity carried out over time, which is a form of internal communication. A single wave front of sound might physically affect a living cell as a form of direct physical contact. A photon of light might similarly affect it. But a complex multi-celled organism as a whole would not be significantly impacted by such miniscule events, which accordingly would elicit no reflex. A continuous *series* of such events, however, *would* bear significance and potentially merit response. A series of sonic wave fronts is registered as a *tone* and a series of photons (or light waves) is registered as sensations of *color* and/or *brightness*. The individual impacts of impinging sound waves or photons in the series, which on a cellular level might trigger a reflex, are instead integrated over time into a *sensation*. The experience of the sensation is how the executive agent is presented with the pertinent information it bears. Phenomenality is, so to speak, the language of the senses, in which each distinguishable sensation has a different meaning, bearing different information. As we have seen, the brain as a virtual system establishes these meanings by its own *fiat*, just as words are created or selected to represent information. Sensation is how the brain's executive agent presents to itself the sensory information it must monitor. Since the perceiving subject *is* that executive agent, conscious experience is the virtual reality through which we (primarily) present to ourselves the world and our bodily relationship to it.

Like any language, this internal language can be used to depict reality or to create fiction. Hence, the inner virtual reality display can be more or less veridical. It can display dream, imagination, and hallucination along with the external world—all with even the conviction of representing something real. Just as grammatical language enables inventiveness, abstraction, and expression unbound by literal truth, so the inner language enables us to have experience that is not literally of the external world but convincing nevertheless. It is this very possibility that necessitates the divide between object and subject, reality and imagination, matter and mind.

Let us return, finally, to the question of whether science can explain consciousness. The account I have given is admittedly not a scientific theory of consciousness. For one thing, it is unclear how to falsify it.⁸ For another, it purposefully deals in hypothetical agencies that are virtual rather than physical systems, intentional rather than causal. The rationale for these notions is that science has limited itself historically and unnecessarily to a certain type of causation and to a certain vision of matter as intrinsically passive or inert. In that vision, matter is a *patient* and only the observing scientist is an *agent*. Science defines itself as a third-person perspective on the world. It studies phenomena but not (first-person) phenomenality. This is dualism writ large, excluding consciousness and agency within the scientific purview, which allows only the mechanical transmission of efficient causes. How to solve a problem that your mental framework will not allow? Well, evidently by altering the framework!

⁸ Inductive generalizations can never be proven, but can be *disproven* simply by finding a counter-example. I am unsure what would be a counter-example to the hypothesis that consciousness results from action-by-fiat within a virtual system. Various “paranormal” experiences might indicate that consciousness is somehow independent of the brain.

So, science *as it stands* cannot explain consciousness. This has not prevented vain attempts to account for phenomenality in strictly physical terms—for example, to locate it in “microtubules” within the brain, which somehow make use of quantum effects to produce consciousness. To my thinking, that simply trades one mystery for another and is no more satisfying than Descartes’ location of consciousness in the pineal gland. All such explanations come up against the perennial objection that they do not address the explanatory gap between mind and matter—the real issue, the so-called “hard” problem. To the extent the gap is an artifact of the scientific approach and its metaphysical assumptions, the only way forward is to modify that approach and those assumptions. I believe that would have benefits as well in other areas of science that involve the troublesome presence of the observer, such as occurs at both extreme ends of the size scale. Both the quantum realm and the realm of cosmology are not only at the limits of observation, but confront paradoxes of the observer’s presence in a framework that is supposed to be purely objective. The co-determination of subject and object ultimately requires an expanded framework.