4
The Boundary Problem for Experiencing Subjects

4.1 First Steps

If physicalism is false, we must look for an alternative way to place conscious experience in the universe. The alternative I explore is not a Cartesian dualism but a version of Liberal Naturalism. Liberal Naturalism is the view that nature is built on a single fundamental kind, and, if so, that some aspects or properties of this fundamental natural kind are not physical. Liberal Naturalists cast the problem of consciousness differently than do some who claim that it is the problem of reconciling consciousness with materialism. William Lycan (1996) expresses this view of the problem:

> It has to do with the internal or subjective character of experience, paradigmatically sensory experience, and how such a thing can be accommodated in, or even tolerated by, a materialist theory of the mind. (p. 1)

> Lycan's statement of the problem makes a pretheoretical commitment to the salvation of a certain metaphysic, physicalism. To the extent that Lycan is representative, one could say that the physicalist's overriding priority is to be ontologically conservative, and that to honor this commitment, he or she has to pay the price of being methodologically radical. As chapter 3 discussed, their prior commitment to physicalism forces physicalists to take such measures as blaming theoretical failures on cognitive deficits of the theory makers rather than on the quality of the theories; approving of appeals to unique and not clearly meaningful kinds of necessity; postulating primitive identities; or arguing for the elimination of self-evident observables.

In contrast, Liberal Naturalism primarily wants to explain consciousness clearly, without appealing to anomalous standards of explanation. The Liberal Naturalist point of view is that the scientific enterprise accepts the discovery of natural ontology as its purpose. The thing keeping Liberal Naturalism honest is

not its commitment to a metaphysic but its rigorous standards for rational explanation. For Liberal Naturalism, setting aside these standards to save an ontological viewpoint is an unwise perversion of science.

Liberal Naturalism believes just this: The problem of consciousness is to understand why it exists; what its relations are to the other things we know exist; and what difference it makes, if any, to the natural order of things. Liberal Naturalism has weaker metaphysical commitments than physicalism because its primary allegiance is to the empirical project of explanation. One might suggest that Liberal Naturalism is metaphysics in the service of explanation, whereas physicalism is explanation in service to metaphysics. Accordingly, the Liberal Naturalist is methodologically conservative, and this conservatism will lead its adherents to be ontologically radical.

My form of Liberal Naturalism is a variant on a kind of view put forward before by authors such as Whitehead (1929), Russell (1927), Maxwell (1979), Lockwood (1989), Griffin (1993), and Sprigge (1994), tentatively endorsed by Chalmers (1995), and recently suggested again by Galen Strawson (1999) and Thomas Nagel (1998):

> It may be that the physical description of the brain states associated with consciousness is an incomplete account of their essence - that it is merely the outside view of what we recognize from within as conscious experience. If anything like that is true, then our present conceptions of mind and body are radically inadequate to the reality, and do not provide us with adequate tools for a priori reasoning about them. (Nagel, 1998, p.)

So this suggestion arose here and there in the twentieth century. Russell suggested that the problem stems from science portraying matter structurally, focusing on its form and not its content. Restating Russell, Lockwood adopts some physicalist terminology from J. J. C. Smart and explains that our concepts of the physical are topic neutral, meaning that they say nothing about the basis of physical being. That is, physical theory says nothing about what physical things are like "in themselves." Whitehead called belief in the adequacy of such descriptions the "fallacy of misplaced concreteness" and argued against any such notion of vacuous actuality.

> One can see a commitment to traditional materialism in the names that proponents of the view often give it. Maxwell called it Nonmaterialist Physicalism. Lockwood distinguishes between Physicalism, which he thinks is false, and Materialism, which he thinks is true. David Ray Griffin writes of Panexperientialist Physicalism. Galen Strawson calls it Realistic Materialism.

Many of these authors seem unwilling to go beyond physicalism in any way more radical than the hypothesis that the physical has an "inner aspect" tied somehow to experience. Seeing no more significance than this in the view will block efforts to fully develop it and make it viable. As I argue in the rest of the book, making this sort of Liberal Naturalism work requires undertaking more thorough revisions to our view of nature than these authors entertain.

Although Liberal Naturalism might feel liberating, we have too much freedom. To find a place for consciousness, we need tests for the minimal adequacy of proposed explanations and also a class of problems able to provide clues that help us triangulate to the point of fundamental incompleteness in our knowledge. As a beginning for the effort, I wish to step back to examine assumptions and to try to identify the deepest problems and clues in the vicinity.

Because we are searching for new facts about nature that are fundamental, the most helpful kinds of puzzles to focus on may not be specifically cognitive puzzles. By *cognitive puzzles* I have in mind the sorts of questions raised by facts such as: conscious states tend to be reliably reportable, conscious states are representational, conscious states contain information that is globally available in the control of behavior, and the fact that the structure of consciousness mirrors the structure of cognitive processing. All of those facts will be very important *eventually*, and any theory must allow us to understand why consciousness has those features. In the context of a foundational search, however, they are not likely to be the best pointers to follow. The next few chapters discuss ways that consciousness raises problems for our general view of nature, not just for our view of the mind or of traditional cognitive science or neuroscience.

For example, the links between conscious experience, voluntary action, and functional awareness lead to very interesting puzzles when considering multiple personality cases (Braude 1991) or commissurotomy patients (Marks 1981) or blindsight patients (Weiskrantz, 1986, 1988). These puzzle cases can be very seductive, philosophically, but if Liberal Naturalism is correct they are likely more intriguing than they are fundamental. Were we to focus exclusively on overtly cognitive features of consciousness such as these, we would run the danger of confusing the inessential with the essential and of overlooking promising paths in our search.

The history of discovery should lead us to expect the deepest insights to come from reflection on the places of paradox, so, ideally, the features we focus on will yield a paradoxical view of the world when combined with its physical image. The task of removing the paradox-driven tension can provide constraints for our search. Each acts as an explanatory target for a Liberal Naturalist view of nature. They might also provide further clues about the location of our missing knowledge.

In this chapter I make the case that there is a puzzle about how consciousness can exist at the middle level of nature, where it does. In subsequent chapters I discuss whether Liberal Naturalism can plausibly restrain itself to a conservative view that only cognitive entities have experiences. Finally, I examine a set of paradoxes involving such things as the unity of consciousness and the causal relevance of consciousness. I ultimately treat each investigation as producing not just a puzzle or a paradox but also potential clues and explanatory targets. Part II of the volume takes up the task of making sense of these clues and accepts the challenge of meeting the explanatory targets.

4.2 Overview of the Boundary Problem

Bertrand Russell once said that the aim of philosophy is to start with something so obvious as to not be worth mentioning and to end up with something so absurd that no one will believe it. My development of the boundary problem for experiencing subjects is in the spirit of Russell. As with many issues surrounding conscious experience, it takes a bit of hard work to bring the depth of the problem into focus.

I start with the observation that consciousness has inherent boundaries. Only some experiences are part of my consciousness; most experiences in the world are not. Arguably, these boundaries are what individuate me as an experiencing subject in the world. I argue that this poses a problem that any theory of consciousness must answer. How can consciousness have boundaries? What element of the natural world dictates the way these boundaries are drawn? This is the boundary problem for experiencing subjects: We must find something in nature to ground the natural possibility of an experiencing subject bounded in just the way human consciousness is bounded.

4.3 The Foundations of the Problem: Obvious Observations

There are obvious observations that help define for us what it is to be an experiencing subject. First, reflect on the fact that experiencing subjects come in discrete tokens. Without too much strain, we can think of each subject of experience as being a kind of quantum. I am one such quantum, so is Trey Kirven, and so are you.[i] These quanta, the individuated phenomenal fields of experiencing subjects, contain coevolving elements. In some vague but compelling sense of *unified*, these coevolving elements are naturally unified into a subject of experience.

Second, the phenomenal field has boundaries. Not every feeling is part of my phenomenal field because I do not feel the pains produced by damage to your body. The unity and boundedness of the phenomenal field stand together at the core of the concept of an experiencing subject. The driving intuition is that experiencing subjects are inherent individuals in a sense of *inherent* that we must try to make clear. If these boundaries could not exist, then nothing like human consciousness would be possible.

Third, our human consciousness is only a species of experiencing subject. An experiencing subject is a manifold of qualitative entities teeming with variety. We only roughly name these entities in our own case, with such words as *feeling, sensation,* and *appearance.* Other kinds of experience may exist in other kinds of beings.

Fourth, the human subject belongs to a human body and its cognitive processing. Humans, and the activity of human cognitive systems, are individuated at a middle level of the physical world. Typically, our individuation of objects at this middle level of nature is fluid, context sensitive, and interest relative. It is highly conceptual and hinges on facts about the abstract organization and causal cohesion of physical activity. As a consequence, events or objects may form parts of

many individuals simultaneously, depending on how one organizes the world and draws the individuating boundaries. For example, a cell may be an individual; also, at the same time, it may be part of an organ; at the same time, it may be part of an individuated bodily system such as the reproductive system; at the same time, it may be part of the organism as a whole and part of that organism's society; at the same time, it might be part of an ecosystem. For each of these different individuals, a different kind of causal organization exists in the world.

Finally, levels of abstract organization and causal cohesion exist between microphysics and human cognition and between human cognition and the universe as a whole. Lycan (1990) especially emphasizes the importance and continuity of the levels of nature, whereas Scott (1995) has emphasized important differences between them.

All those things should be more or less obvious. The part that is not obvious requires putting these observations together in a way that makes it clear that the experiencing subjects could have been different individuals than they are and different in ways that would prohibit human consciousness. Therefore, we need to explain why experiencing subjects are associated with the specific, middle-level patterns of interaction and organization with which they happen, in fact, to be associated.

4.4 Defining the Problem

The very obviousness of our own existence as middle-level experiencing subjects is an obstacle to appreciating the boundary problem. Because the boundaries of consciousness are something that is always with us, it may not be easy to realize how remarkable it is that things are this way. I now want to bring out the stark brutishness of the fact that experiencing subjects like us could even exist, individuals localized at a middle level, with middle-level boundaries to what we feel. The main points are: (1) if it were not possible to draw these boundaries to the phenomenal field, humanlike experiencing subjects could not exist; and (2) the fact that such boundaries exist where they do is surprising, and their basis is not obvious. To bring out the problem more vividly, I use several thought experiments designed to loosen our sense that there is a natural inevitability to the boundaries that actually exist.

Abnormal forms of consciousness, such as multiple personality disorder (MPD), open the door to the possibility that, in some circumstances, multiple experiencing subjects may coexist within a single brain. Braude (1991) describes cases of MPD in which different personalities may be copresent, each claiming to be a distinct center of awareness. Among many peculiarities, these centers of awareness (which Braude describes as apperceptive centers) make claims to sharing a variety of relations among their experiences. Sometimes they claim distinct experiences altogether. In these cases, the experiences of each personality are "screened off" from the others, so the different personalities achieve, apparently, privacy of experience. In other situations, their experiences partially overlap,

some belonging to multiple centers of awareness and others only to one. In still other cases, experience may be completely shared, although particular experiences may sometimes claim to be owned only by one or another center of awareness. In such cases, John and Mary may both claim to have an experience, but only John claims it as his experience.

These cases raise very puzzling issues about the facts of the matter. What are the number and boundaries of the experiencing subjects that exist in these cases? Is there really one for each personality, or are the claims issuing from confabulation? Whatever the truth, it seems to me that different hypotheses are at least coherent: there could be one experiencing subject, there could be many, and perhaps there are even overlapping subjects of experience. Accepting that there is a legitimate scientific question here, that nature could deliver any one of several possible answers, is a first step in beginning to see the boundary problem.

Most likely, the boundaries of consciousness correspond to the boundaries of certain distinctive activity in our brains. Some evidence suggests that specially synchronized activity in and around the cortex, modulated chiefly by the thalamus, constitutes the boundary maker for human experiencing subjects. For example, Crick (1995) hypothesizes that different regions of the thalamus coordinate each level of visual processing. Llinas (1994, 1996) reports the existence of a wave of coherent oscillatory activity that sweeps the cortex every 12-13 msec that is perhaps generated by thalamic activity and postulates that it binds separate sensory content together into a unified representation. Newman (1996) and Newman et al. (1997) argue that this activity constitutes the binding of sensory contents into a global workspace whose contents are neurally broadcast to specialist subsystems. According to the picture that is emerging, this activity as a whole corresponds to an experiencing subject.

But this raises the questions, What counts as a "whole" for nature, as far as it is concerned with experience? And why? I begin clarifying the importance of these questions by asking, Might any of the subsystems oscillating within this magnificent whole also constitute an experiencing subject? Consider the patterns of synchronized activity that carry and organize auditory information from our ears. Is there an experiencing subject, existing at a different level of nature, that is associated with this activity alone? The picture I am proposing is something like this: Within ourselves as fully human experiencing subjects, there would be other experiencing subjects, themselves perfectly complete subjects of experience, although simpler. Like Russian dolls, there would be individuals within individuals within individuals, all of them subjects of phenomenal experience. The hierarchy of nature might then contain a hierarchy of experiencing subjects, each more or less complex.

The very definiteness of what it means to be an experiencing subject seems to require an answer. Given that we understood the obvious observations about experiencing subjects, extension by analogy should allow us to make sense of this question. Bring to mind the physical image of the world: We have a multitude of

interacting microphysical entities at places and times, this multitude congealing into a macroscopic whirlwind of finely layered patterns of organization. Simply imagine looking at the patterns of physical activity in the world from the perspective of a third-person observer. Note the coherence of causal and abstract organization at the many levels and the many ways it exists. We know that a set of these patterns supports boundaries that allow for the existence of us, where we are one kind of experiencing subject.

The question here is, Does nature support other kinds of feeling subjects, other kinds of experiencing beings? Analogous to us as experiencing subjects, might there exist simpler experiencing subjects whose boundaries are given by subsets of the activity that determines our experience as a whole? After all, the relevant subsets of activity are like the more complete set in many ways. They share common biology with the larger set of events; they carry information and are processing it; they process it in a very similar way; and within themselves, they are internally synchronized and coherent. Do any of these subsets of activity support experiencing subjects, also?

We do not need to assume that an experiencing subject corresponding to auditory activity experiences sound. Instead, it might experience qualia uniquely appropriate to its own level of reality and be responsive to its own finer grained causal organization, just as we are. The case of multiple personality disorder (now called dissociative identity disorder) suggests that a single brain might be able to support several experiencing subjects at the same level of organization. The next step is to wonder whether the normal way of things in the brain might be to support several experiencing subjects but at different levels of organization. The one with which we identify might just be the one at the highest level of organization. Is this the way our world is?

Because such brain activity, taken as a whole, corresponds to the existence of experiencing subjects, *us,* the point seems to generalize to a relatively mild claim. Our intuitive concept of what it is to be an experiencing subject allows for the possibility of simpler experiencing subjects, individuals whose manifold of experience consists of much less rich and less cognitive experience. Given this, the physical activity in fact corresponding to the existence of an experiencing subject might also support other, simpler experiencing subjects via the simpler patterns of organization it contains. Of course, it might not.

I am not claiming anything about the plausibility of this view. I am only claiming that it is an epistemically possible view. It is not a question that one can answer through a priori reflection on the nature of experience and the nature of the physical. It is an empirical question that arises only after one is aware of the physical facts and is suggested by them, as after reflecting on the physical situation, both yes and no seem possible. Why couldn't there be experiencing subjects at many levels of processing, some associated with subsets of the cognitive activity corresponding to our own experiences? On the other hand, why would there be? I now take this openness in the concept and, in steps, parlay it into the full-blown boundary problem for experiencing subjects.

4.5 Sailing toward Scylla and Charybdis

The next step toward the boundary problem is built on a variant of Ned Block's well-known fiction of the Chinese nation simulating the functionality of a human brain (1980). To make it a little less fantastic, we can imagine the simulation of some other, simpler kind of organism's brain, maybe a fish. Very likely, a fish is an experiencing subject.

Imagine building a robot fish. Imagine also that we have designed its nervous system to be functionally isomorphic to the nervous system of a naturally occurring fish (assuming that's possible). The processing has been made remote in the usual way, with inputs and outputs to the fish's central nervous system employing relays. These relays send signals to remote stations manned by human beings. The humans monitor the signals as they come in and relay an output to other destinations. Some signals are sent between the remote stations, and some are relayed back to the fish as motor outputs. In this way, we imagine that the relay system is functionally isomorphic to an actual fish's brain.

> The question Block raises is whether or not a system like this system would be conscious. Block uses the example in an attempt to show that our concept of phenomenal consciousness is not a concept of a purely functional entity and that it supports the view that consciousness does not conceptually supervene upon functional organization. However, Block's argument fails to show that the system will not contingently support the existence of consciousness. As is often pointed out, it seems surprising that our brains would support consciousness, but we know firsthand that they do.

Because the relay system is functionally like a fish's brain, it is certainly conceivable that this system actually supports an experiencing subject. The system has parts that are phenomenal homunculi, and I previously argued for the consistency of the idea of experiencing subjects whose physical organization supported the existence of other experiencing subjects. In fact, both of these seem to be at least epistemic possibilities:

(1) Each homunculus is an experiencing subject, but the whole system is not.

(2) Each homunculus is an experiencing subject, and so is the whole system.

These possibilities are eye-opening because we can redirect the principles that make them plausible back to a local system for the fish. The homunculi system is functionally isomorphic to the fish's cognitive system. Each homunculus maps onto some important part of the organizational structure of a naturally evolved fish. Imagine the mapping being made with one of the homunculi, call her Edna, mapped onto some functional part of the fish, call it the E-system. There is no principled reason to restrict possibility (1) to the robot fish alone. By analogy, (1) would seem to ground the possibility that the natural fish's E-system, the part corresponding to Edna, could be an experiencing subject, even though the fish as a whole would not be.

How does the analogy go? By admitting possibility (1), we are admitting the coherence of the idea that the robot system as a whole may not be an experience

ing subject. In doing so, we are admitting the coherence of a world in which (a) a system may contain experiencing subjects, (b) that system may be functionally isomorphic to the fish's system, and yet (c) that system is not an experiencing subject. In the previous section I gave reasons why it seemed coherent that ordinary cognitive subsystems could themselves be experiencing subjects. To imagine the E-system as an experiencing subject, but not the fish, we have to combine the two points.

To combine them, conceptually shift the boundaries that make experiencing subjects. Shift one's view of nature so that the phenomenal boundaries stretch through the E-system, encompassing all the activity within it but not overflowing the boundaries of the E-system. The larger individual is abolished. In its place is a collection of simpler experiencing subjects in a system of competitive and cooperative interaction. Of course, the experience of the E-system would be vastly different from the experiences of Edna. Accommodating the difference between Edna and the E-system requires postulating alien experiences for some simpler beings. Our ordinary concept of experience is tolerantly open-ended in this way, so this requirement does not stand in the way of our being able to change the intuitively assigned phenomenal boundaries. Lacking clear criteria for natural boundaries, conceptually we can rearrange the boundaries, forcing the individuality down to the E-system level. By doing it, we rob the natural fish of its status as an experiencing subject.

Once we have seen the essential analogy between Edna and the E-system, we can begin to engage in other conceptual shifts. Obviously, the E-system might not be an experiencing subject. It is perfectly coherent to suppose that the only experiencing subject associated with the fish's brain is the one existing at the global fish level. The coherence of the idea that natural fish do not have phenomenal E-systems seems to support a third possibility:

> (3) The homunculi system would be an experiencing subject, but none of the homunculi would be.

The possibility that the E-system is not an experiencing subject means that some systems that have experiences in some (epistemically) possible worlds do not have experiences in others. We can apply this principle to Edna. Although everyone knows that in fact Edna would be an experiencing subject, we cannot overlook the failure of consciousness to conceptually supervene on the physical. This failure raises the logical possibility of phenomenal Zombies. Phenomenal Zombies are physical systems organizationally just like human beings but without consciousness. With the specter of Zombies looming, we need to explain why (3) is not true even of our world. For instance, some people think that experiencing subjects emerge at a certain level of complexity, and this seems like an empirical possibility. If it is possible, then imagining (3) merely requires imagining such a world, dictating an appropriate kind of complexity, and then moving the starting point upward, past Edna. The complexity point at which experiencing subjects emerge would be higher than that possessed by the homunculi, but not

by the homunculi fish. The result is that Edna could be a zombie and a component in a system supporting an experiencing subject. What keeps conscious experience right there, between Edna's ears?

This reconception of boundaries is just the flip side of the earlier suggestion. Earlier the reconception was a movement of the phenomenal boundaries to lower levels of organization, robbing wholes of their experiences. Here, the reconception is to a higher level of organization, robbing parts of their experiences. Such a world would be one in which eddies of coherent causation that are human bodies would not support experiences, and human phenomenal consciousness would not exist. We might say that in this imagined world there are humanlike *bodies* but no human *beings*. Instead, the experiencing subjects exist at a higher level. As human bodies act, exchanging signals with one another, as well as interacting with other causal eddies, the phenomenal individual arises only for the supersystem. The possibility is analogous to the way we (or many of us) normally imagine that our cognitive subsystems contribute to our conscious lives without themselves being experiencing subjects.

I intend these science fiction tales to make vivid how an intuitive understanding of experiencing subjects supports a great deal of possible variation in their boundaries. Once the basic point has been appreciated, we can make the point without using philosophers' thought experiments. Even actual systems, such as economies or political systems or nation states, bring it out. An economy is an extremely complex dynamical system of self-organizing components. As a physical system, it stores information and seems to have a kind of distributed memory, a high degree of synchrony between its parts, massive parallel communication, global broadcasting and dominance of certain information, feedback loops, and so forth. The spatial and temporal scales at which this all takes place are much larger than in an individual brain, and much different in detail, but the same basic kinds of activity exist. It would be nothing new to suggest that an economy might represent some kind of group mind. And, really, it is not just a philosopher's question. It is also a deep scientific question about the nature of mind, as well as a legitimate question of fact about something that actually exists in the real world.

What is the main reason for rejecting the idea? The economy certainly has representational properties, and it is a representation consumer: Money, its lifeblood, is a representational vehicle through and through. Mostly, the problem is the bizarreness of believing that the U.S. economy is conscious, and most people consider consciousness essential to mind. Now, I do not know if the economy in fact supports the existence of an experiencing subject and actually tend to doubt it myself. Still, it seems as a priori coherent to me that an economy could (on a much slower time scale) support such existence as it does that my brain would, and I am quite sure my brain does. The economy would have to possess a very different kind of phenomenology, but there does not seem to be good reason for thinking that human-type experiencing subjects are the only kinds that could exist. There's no escaping that economies share many of our mind's most salient characteristics, stretched out vastly in scale over space and time.

Even the scale differences do not amount to much once one considers that our experiences arise from collections of atoms and molecules. The time scale they operate on is far faster than the time scale on which brains produce consciousness. If bunches of neurons (or molecules) stand to us as we do to the economy, and if their organization supports us as experiencing subjects, why couldn't we support the economy similarly? Once we see the possibility that both the economy and our bodies might support experiencing subjects, we are only a short step away from seeing another possibility. It might have been that we are not phenomenally conscious but that the economy nevertheless would be. After all, it seems coherent that our neurons are not experiencing subjects, even though we are. The boundaries of experience, once loosened, can begin to shift radically. Again, why are our bodies not simply local, nonphenomenal causal eddies within a larger phenomenal individual? What grounds the brutishness of these boundaries?

4.6 Scylla and Charybdis: The Boundary Problem

The boundary shifting that occurs in these thought experiments is enabled by the fact that information about physical pattern and organization alone does not fix the boundaries of experience. We individuate most objects at higher levels of organization by extracting some significant pattern from the flux of microphysical interaction. Consistent with a given pattern of microphysical causation, innumerable ways exist of conceiving and reconceiving the abstract organizations that supervene.

Just adding these facts about pattern, or abstract organization, to the causation between the microphysical entities does not seem to go far enough in determining the proper sense of *inherent* in the idea that an experiencing subject enjoys a kind of inherent individuality. One can coherently hypothesize almost as many ways of determining boundaries for experiencing subjects as there are of abstractly organizing and reorganizing the patterns of microphysical interaction in the world. The resulting scenarios are intuitively bizarre, but bizarreness is not inconsistency. The fact that nature's boundaries yield human consciousness stands out as a brute fact.

We are faced with the need to understand more deeply what it is to be an inherent individual in the natural world. We need a natural criterion for individuation, one that illuminates the specialness of some patterns over others as supporters of experience. The fields of the most primitive particles (or strings or whatnot) make one good set of candidates. Each of these has a natural dynamic unity, one that seems inherent. An experiencing subject might be associated with each of these.

This suggestion threatens human consciousness. If the fields of the primitive individuals of physics are the only natural individuals, the rest of us are mere abstractions off the pattern of their interaction. Each primitive physical individual may be a simple experiencing subject, supporting firefly flickers of feeling briefly buzzing at the lowest levels of spacetime, but above them the world is dark. This world would be the panpsychist's world painted surrealistically. There is nothing

that can bootstrap us to human consciousness: feeling, feeling everywhere, but not a drop can think.

Perhaps, by flowing along the lines of interaction, the experiencing subjects could outrun the boundaries of the primitive individuals of physics. Here the trap concerns stopping the flow of interaction. It can seem that the flow of interaction in the universe is inherently unbounded, and no merely abstract pattern presents a natural condition for containing it. Those patterns merely direct it from one watershed to another, orchestrating it, moving it along through the continuity of the universe. According to this view, experience must follow the boundaries to their limits along these lines of interaction. This makes for the possibility of a universal subject of experience, perhaps some kind of cosmic consciousness. Unfortunately, no room exists for the more mundane, middle-level boundaries necessary for human consciousness to exist. Like the first view, this view banishes middle-level individuals from existence.

These two views are a Scylla and Charybdis for Liberal Naturalist theories of consciousness.[ii] One view pushes us inward, past the point of middle-level individuation, and into the realm of the subatomic. There, and only there, do we find our natural, inherent individuals. Another pushes us outward, past the boundaries of the subatomic individuals, ever outward along the lines of interaction between them, racing past the middle level to the continuous unfolding of the cosmos. Only there, at the level of the universe, do we find our inherent individual. Neither view allows for conscious human beings. To navigate the middle ground, we must find a principle that allows us to push those boundaries outward from the microphysical but only *just so*. We must be able to go only so far past the microphysical level and no farther. That is the boundary problem for experiencing subjects.

4.7 The Teeth of the Problem: Two Examples

By considering two examples of dual-aspect theories that falter on this problem, we can get a better sense of its importance. The two proposals I briefly critique are the materialism of Michael Lockwood (1989, 1993) and the information theory of David Chalmers (1996). I do not believe that either successfully navigates the way between Scylla and Charybdis.

Lockwood Michael Lockwood's materialism is a resurrection of Bertrand Russell's neutral monism in the context of quantum mechanics. In an argument similar to my argument from *Life* in chapter 2, Lockwood suggests that phenomenal consciousness fails to logically supervene on the physical because physical concepts are content-neutral, merely specifying the structure of the causal flux. Phenomenal qualities and consciousness, on the other hand, are defined precisely by their content. Lockwood suggests that a nice solution to the problem is to simply draft phenomenal properties into duty as the content of the causal flux whose structure is described by physics. The result is a kind of dual-aspect theory. Physical concepts are about the structural aspects of the causal flux, and our phenomenal concepts are about the intrinsic content that is in flux.

This is an interesting proposal, but something needs to be added before it can hope to account for the individuation of human consciousness at the midlevel. After all, if we are to believe physics, the individuals who are the natural candidates for this basic phenomenal content are the fundamental fields. Lockwood's theory needs to take us from this simple phenomenal content of simple individuals to the complex, middle-level experiencing subject necessary for human consciousness. According to one horn of the dilemma, he is stuck at the microphysical level.

Lockwood (1989) appeals to the other horn for help, postulating that phenomenology flows along the lines of interaction in the world. Unfortunately, he has no principle to allow him to resist being hung on this horn, as it urges that the boundaries be pushed further and further outward. Once the bootstrapping process has begun, Lockwood's theory gives us no explicit way to stop it. Actually, the problem is a little worse for Lockwood, because he is sympathetic to the Everett interpretation of quantum mechanics. Interaction, although structured, is seamless in Schroedinger's world, and Charybdis demands a reason for stopping it here, where there are human cognitive systems in one eigenstate. There does not seem to be a compelling reason to think Lockwood's proposal would result in anything less than a many-worlds-sized individual. Lockwood (1993) discusses this problem and makes this appealing observation:

> In quantum mechanics there is a sense in which all observables, and in particular observables corresponding to every level of structure, are to be regarded as equal in the sight of God, as are different frames of reference, relativistically conceived. As I intimated earlier, quantum mechanics seems to be telling us that it is a classical prejudice to suppose that the world is not intrinsically structured at anything but the level of elementary particles, and their actions and interactions.

This sets out the problem and a possibility for solution. Alas, Lockwood concludes:

> For our own awareness, so I have been urging, embodies a preferred set of observables, which in turn amounts to saying that its contents, at any given time, embody the answers to a set of questions about the state (the intrinsic state) of the underlying brain system. Sadly, however, we here find ourselves in a predicament....We know the answers to those questions, in a way that a scientist, merely by examining our brains from without, never could. But unfortunately, we have, as yet, no idea what the questions are!

In other words, for Lockwood's view to work, we need to find a basis for the existence of an intrinsically preferred set of quantum mechanical observables at precisely the level at which awareness emerges. This, however, is just the boundary problem rearing its head.

Chalmers David Chalmers (1996) proposes that phenomenal properties and physical properties might be two aspects of information spaces. If we take his suggestion as being unrestricted, it is immediately confronted with the problem of individuating information spaces. On Shannon's view, which Chalmers appeals to, information is a difference that makes a difference along some causal pathway. But a difference that makes a difference to *what* along the pathway?

Falling on the first horn, we can recognize informational differences to the basic individuals, but that banishes human consciousness.

Falling on the second horn, we can recognize the universe as a whole as an information space. Its structured state changes as interactions occur within it, and one state of the universe makes a difference to subsequent states, but this also banishes human consciousness. On Chalmers's proposal we should be able to save middle-level individuals by allowing for all covarying subportions of spacetime to be information spaces, but then we are left with panpsychism run wild. Even within one brain, we will have astronomically large numbers of experiencing subjects, separately experiencing, each corresponding to different ways of carving up the activity of the brain and its causal pathways. An explanation that promiscuous is not illuminating.

4.9 What to do?

The Liberal Naturalist should take the boundary problem seriously and think hard about what might be missing from our current view of individuation in the world. The suggestion that we allow inherent individuality to flow along the paths of interaction between individuals sounds promising. After all, we are looking for something more than abstract organization to ground judgments of natural individuality, and causation seems to be an inherent, natural connection par excellence. Also, an interaction divides the world by its very nature, partitioning it into different spaces that mutually condition one another. Lockwood's observation suggests that the second horn of the dilemma gains its conclusion by taking advantage of a naive view of interaction, one that capitalizes on a rough classical understanding of causation.

One good strategy to follow would be to think harder and more carefully about interactions in the world. We should think in more detail about the way they might condition nature into individual, mutually influencing regions, and do so at many levels simultaneously. We very well might discover that interactions have certain important aspects we can use to mark off candidates for natural individuation. These individuals would then be candidates for supporting experiencing subjects. The job would then be to look for a physical reflection of this special feature of interactions.

Like Lockwood, I think that we must understand the causal structure of our world better. Causal connections seem the best candidates for helping to understand more deeply the naturally individuated, middle-level structure exhibited in our phenomenal existence. This is the first example of a conclusion that pops up again and again in this section of the book. By its end, it will seem that wherever we turn in trying to understand consciousness, we end up spun around and facing questions about causation. This section of the book thus serves not only as a discussion of problems and challenges facing Liberal Naturalism but also as a runway to the eventual topic of causation.

<FN>[i]I am using "quanta" with its classical meaning of "discrete units." I do not mean to suggest ties to quantum physics.</FN>

<FN>[ii]I think it presents the same kind of challenge even for nonreductive physicalists who were not persuaded by the arguments in chapters 2 and 3. It is an explanatory problem, and the explanatory problem arises just as strongly even if the possibilities used to generate it are just epistemic possibilities.