Prologue: The Most Fundamental Question

Why is there something rather than nothing? I expect that many share my bafflement, a true sense of wonder, in the face of this ultimate question. But what, precisely, is it about? The conventional interpretation is that it is about the origin of the physical universe. Everyone acknowledges that had the big bang been called off, there would have been nothing rather than something.

There seems to be, though, an alternative way in which there would have been nothing rather than something. Consider the question, Why is there something rather than nothing, *for me*?

That seems to be a mystery at least equally fundamental. It is, at bottom, the problem of consciousness. When your consciousness first emerged from subjective nothingness, what happened was a kind of psychological big bang. Your consciousness had been absent from the physical universe for a respectable while—about 14 billion years or so—but then suddenly there was light *for you*. As a matter of fact, until recently we all had been dead since the dawn of the physical world. Our subjective existence is a kind of life after death; indeed, life after billions of years of being entirely devoid of subjective consciousness. Sadly, it looks like each and every one of us will once again return to that state of impenetrable darkness, probably far sooner than any of us would care to.

Without your consciousness ever coming into being, there would have been nothing rather than something (*for you*). Without the phenomenon of consciousness emerging in some place at some point in the history of the physical universe, there would have been nothing rather than something *for everyone*, and consequently, nothing rather than something, across the board. The physical universe, if deprived of us conscious subjects, has absolutely no way of sensing or knowing its own existence; it simply couldn't care less. The physical universe in itself, then, leads a pretty boring life, at least compared to us sentient creatures. If you don't even know whether or not you exist, and you cannot care, is there any way in which things could get more indifferent for you?

So maybe the importance of the most fundamental question, if taken as the mystery of the physical universe, has been somewhat overestimated. Why is there something

xvi Prologue

rather than nothing? A few bits of dead physical matter and energy hovering about in an otherwise empty space-time would be enough to bring up the mystery of the physical universe in its entirety, but even if that mystery could be solved somehow, would the solution make much difference anyway?

The mystery of consciousness is a horse of another color. Consciousness matters. It makes a difference whose significance it is difficult to overstate. Just think about it: What if you, that is, your *consciousness*, had never existed? What if there had never been anything for you? Surely that is not a minor difference. For whom? Well, the universe doesn't care anyway, so it is pointless to envisage things from its point of view. Of course it would not have noticed if you had not existed, but, then, it would not have noticed even if half of itself were missing. So why ask the universe in the first place? Let's ask you instead. Had you not existed, the world-for-you would not have existed either. Your subjective psychological reality, the world-for-you, would never have been. It is not a world of space-time, force fields, or subatomic particles we are talking about now, but a world of joys and sorrows, colors and sounds, thoughts and memories, embodiment and movement: a reality of lived experience and sensation. It is a world that, at the very least, senses its own existence, though it may not understand it or have the ability to conceptualize it as sensed existence. Nevertheless, existence it is, a subjective existence that matters in a way that the objective existence of an entire universe of physical substance, if devoid of consciousness, could never matter.

In the light of the above considerations, it is no wonder that the problem of consciousness has been regarded as perhaps the most significant question that still remains unsolved by science. Consciousness is literally a matter of life and death: you exist only insofar as your subjective psychological reality exists. When it is wiped out for good, the world-for-you will be gone, and so will you.

Thus, it would be really nice, wouldn't it, to have a clear scientific theory about consciousness that explains, once and for all, what these subjective worlds-for-someone are really all about; what the place of the subjective psychological reality is in the physical universe? Although the final answers will be certainly hard to come by anytime soon, it seems to me quite clear what the fruitful scientific approach to this problem should consist in. It should take very seriously both the subjective psychological reality of consciousness and the objective neurobiological reality, found in specific regions, called "brains," of the physical universe. Somehow, these two realities, although seemingly worlds apart, must be intimately connected.

If we can make any sense of that connection at all, it would seem that, one way or another, those two realities simply *must* lie along a single integrated continuum in the seamless order of nature. Perhaps such a continuum could be discovered by following—at least for the time being—the working hypothesis that consciousness is just another biological phenomenon in the brain. Not, I gather, an entirely ordinary biological phenomenon; instead, a very special one with several puzzling and amazing features,

but a natural biological phenomenon all the same. Therefore, treating consciousness as a real biological phenomenon in the brain seems to me a reasonable approach to the scientific study of consciousness.

I call this approach *biological realism* about consciousness. Biological realism treats subjective consciousness as a *real* phenomenon instead of denying or eliminating it as some sort of grand illusion or intellectual error. Biological realism regards consciousness as a *natural* phenomenon, one that requires a fully naturalistic explanation, instead of treating it as some sort of supernatural or mystical phenomenon entirely beyond the reach of human science. It takes consciousness as a *biological* phenomenon, which means that consciousness is realized at spatial and temporal scales, and levels of complexity in nature, similar to other biological phenomena. That is to deny that consciousness is basically a phenomenon found at the microscopic level of subatomic particles or quantum effects, for example.

Biological realism assumes that the biological phenomenon of consciousness is tightly anchored to, indeed literally located within, the confines of the biological reality found in the brain. Therefore, it denies the philosophical idea according to which consciousness, in some form at least, is present all over the physical universe, even in such unlikely places as electrons, stones, or thermostats (panpsychism). Biological realism also denies the position that consciousness has *no* definable spatial location or spatial properties whatsoever. It further denies the currently fashionable view of "embodied cognition" according to which consciousness is not confined to the brain, but in some unfathomable manner pervades the whole human body (or even extends beyond the body to the surrounding environment).

Biological realism treats the empirical biological sciences and the explanatory strategies applied to them as the best tools we have when trying to understand the place of consciousness in nature. Consciousness should be seen as a biological phenomenon in the brain, located at a particular level of organization in the overall hierarchical structure of nature. A biological research program on consciousness should build its philosophy, methodology and empirical approaches on these kinds of basic assumptions.

Cognitive neuroscientists may have taken for granted that something like biological realism must be the mainstream philosophical approach to the study of consciousness. After all, there is nothing terribly radical or inherently implausible about it—as a matter of fact, it is just our conventional scientific worldview extended to the problem of consciousness. Perhaps something like biological realism is tacitly assumed by most of the neuroscientists who try to uncover the neural basis of consciousness in the brain. But, strangely enough, most of the current philosophical approaches to consciousness seem to be in serious conflict with biological realism. The reason for the dominance of antibiological approaches in philosophy would be somewhat less of a mystery if the biological approach to consciousness had been already thoroughly examined by philosophers and found to be wanting. But, unfortunately, that does not seem to be the case.

xviii Prologue

Apart from the philosophical contributions of John Searle (1992, 2000), the theoretical and empirical implications of biological realism remain largely unexplored.

Therefore, my main task in this book is the following: To depict a *biological research program on consciousness* that accepts biological realism as its basic principle. A research program should incorporate a coherent global view of its subject domain. In the following chapters, I attempt to bring the study of consciousness under the unifying theoretical framework of biological realism.

This undertaking involves the following tasks. First, to get an idea of what the explanation of consciousness requires, consciousness must be placed in the same boat with other complex biological phenomena: the framework of multilevel explanation. Second, the currently available theories and empirical methods must be evaluated by asking what they can reveal about the levels of biological organization in the brain, where consciousness resides. Third, the wealth of results from different experimental approaches to consciousness should be placed in the context of the biological research program. Only in such a way will it be possible to understand how the different types of experiments and findings contribute to the overall understanding of consciousness.

The explanation of consciousness, in this view, is a slowly progressing scientific effort, an empirical and theoretical puzzle to be solved piece by piece; a biological mechanism to be elucidated gradually and piecemeal by empirical science. I fear that the problem of consciousness will never be resolved all of a sudden by constructing a single philosophical argument, no matter how witty, nor by running a single brain-imaging experiment, no matter how ingenious. Instead of such isolated attempts, an overall scientific research program is needed that provides a general philosophical and empirical framework, the big picture, to which the problem of consciousness can be firmly anchored, and, consequently, taken apart bit by bit and thereby solved, step by step. After all, that is how biological science has succeeded in solving other great mysteries of nature which once seemed eternal and irreconcilable.

So far the field of consciousness research has been quite a mess just because it lacks the big picture or an overall unifying theoretical framework. In the course of this book, I explore the major theoretical and empirical issues in consciousness research and relate them to each other under the framework of biological realism. I shall proceed in this task as follows.

The first, introductory part of the book opens with the observation that, despite frequent references to a "science" of consciousness, the field of consciousness research is still a long way from constituting a true science in the sense of a unified research program. In chapter 1 I delineate what would be required of a unified scientific research program on consciousness. I argue that any such research program must be committed to some shared fundamental assumptions concerning the nature of the *explanandum*, and I propose that biological realism provides us with a theoretical starting point from which to launch the hard core of a promising research program. Other types of funda-

mental assumptions (e.g., panpsychism, functionalism, and representationalism) lead to research programs dissimilar from biological realism in important respects. Many find the alternative approaches more plausible or acceptable than biological realism. In the course of this book I explain why I do not find the alternatives all that inviting.

Biological realism necessarily anchors the explanation of consciousness in the context of other biological sciences, especially cognitive neuroscience. Recent work in the philosophy of science has shown that the proper explanatory strategy in the context of the biological sciences is not the traditional idea of "theory reduction" or "deductive-nomological explanation" but rather an approach that may be called "multilevel explanation." I will demonstrate how consciousness as a biological phenomenon can be fruitfully placed in the framework of multilevel explanation.

This explanatory framework of the biological sciences greatly clarifies what it means to "explain" consciousness. Consciousness can be placed in the multilevel framework if it is treated simply as one particular level of biological organization: the phenomenal level in the brain. In the multilevel framework the overall explanation of consciousness then breaks down into several smaller tasks: the downward-looking, backwardlooking, and upward-looking explanations. Different sorts of questions and evidence are relevant for illuminating each of these explanatory dimensions. The downwardlooking explanation should answer the question, What are the neural microlevel mechanisms of consciousness? The backward-looking explanation illuminates the causal history of consciousness, How are the contents of consciousness modulated by external stimuli?, When does consciousness emerge and how does it develop in the brain of a newborn baby?, and Why was consciousness selected for during evolutionary history? The upward-looking explanation tells us what the functional role of consciousness is in the brain and in guiding external behavior: How does consciousness contribute to adaptive behavior? In the rest of the book these different explanatory dimensions are systematically charted.

Chapter 2 begins with the observation that before an *explanation* of any biological phenomenon is feasible, the explanandum must be clearly and unambiguously *identified*. This requires a systematic description of the phenomenon. One way to this end is to clarify the conceptual mess surrounding consciousness. The biological research program needs a detailed *conceptual map* of the territory. Only then is there a real possibility for conceptual integration so that most people would use the fundamental concepts of the field in the same way. I present an initial map of the conceptual territory by introducing the most central concepts and conceptual distinctions that I believe a research program on consciousness should include.

Part II of the book describes the theoretical and philosophical foundations of the research program. As evidenced by the history of biology, a useful strategy to make initial theoretical progress in a relatively new area of empirical research is to utilize *metaphors* and *model systems*. In chapters 3, 4, and 5 I explore the model systems and

xx Prologue

metaphors of consciousness. Visual awareness has been proposed by many to be the most promising model system for consciousness research. I suggest that in addition to visual awareness, the dreaming brain also constitutes an excellent model system for the study of consciousness. During dreaming, the full range of subjective experiences is brought about by a brain that is effectively isolated from stimulus input and motor output. The dreaming brain thus creates a kind of "pure" consciousness, which has wideranging theoretical and philosophical implications. By exploring the phenomenology of dreaming we end up with a particular metaphor of consciousness that captures the phenomenon — the explanandum—purely from the first-person point of view: the world-simulation metaphor of consciousness.

In chapter 6 the theoretical foundations of the world-simulation metaphor are explored. The basic feature of phenomenal reality is *the sense of presence*: The contents of consciousness form *the world now present for me*, a kind of virtual reality in the brain. Thus, the theoretical explanation of consciousness requires a philosophical account of *inner presence* and *presentation* (rather than *re*presentation).

The epistemological and metaphysical implications of the world-simulation metaphor are analyzed in chapter 7. I argue that the biological research program requires a philosophical account of "presence." This concept has a respectable philosophical background in old discussions of "acquaintance." A few current philosophers have reintroduced the concept of "presence" in some form into modern discussions on consciousness, but these writings seem to have appeared in isolation from each other, not accumulating as a serious alternative in the philosophy of consciousness. One reason for this may be that many philosophers still seem to think that this type of view will necessarily lead to a hopeless metaphysical and epistemological cul-de-sac where an internal homunculus must be postulated who infallibly perceives private mental objects or sense-data. I will show that it is not all that difficult to avoid those outdated philosophical traps and to construct a coherent and plausible way of thinking about consciousness as inner presence.

There are several lines of philosophical thought that resist biological realism or the internalism inherent in the world-simulation metaphor. In chapter 8 I try to expose the reasons why many reject the view that consciousness is a biological phenomenon residing within the confines of the brain, or the view that inner phenomenal presence is the most fundamental aspect of consciousness. I suggest that alternative ways of characterizing and localizing consciousness, advocated by such approaches as direct realism, representationalism and content externalism, embodied cognition, neurophenomenology, projectivism, and panpsychism, are not entirely convincing as the philosophical basis of a progressive scientific research program on consciousness.

Part III concentrates on the task of "empirical phenomenology": the description of consciousness on its own terms. In chapter 9 I describe the basic principles of such phenomenology. I end up arguing for the view that the most fundamental ingredient of the

phenomenal level is some sort of a unifying spatial coordinate system—which I call "virtual space"— in which the rest of the phenomenal contents must be embedded to be directly present in subjective experience. Virtual space is the invisible but unifying framework that underlies the global unity of consciousness. It explains why there is at any one time exactly *one* world-*for-me* in which all my phenomenal experiences occur in a fundamentally interrelated fashion. In chapter 10 I apply empirical phenomenology to illustrate how the multifarious contents of consciousness are constructed out of dissociable phenomenal parts and features. This exploration leads to insights about the nature of "virtual objects" and the "virtual self," fundamental components of phenomenal consciousness during both dreaming and wakefulness.

The phenomenological observations thus suggest that, on the one hand, consciousness is fundamentally unified—it forms a single phenomenal world-for-me—but on the other hand, the diverse phenomenal features and contents of consciousness must be based on a number of different underlying mechanisms. In part IV I concentrate on the unity of consciousness and on the problem of its mechanisms: the binding problem. In chapter 11 I show that the binding problem is not a single problem, but a set of related problems that need to be disentangled from each other. "Binding" refers to three distinct problems in three different disciplines: neuroscience, cognitive science, and consciousness research. Once these distinctions are made clear, we can see how the binding of phenomenal consciousness, or the unity of consciousness, relates to research at the cognitive and neural levels of description. Chapters 12 and 13 then apply this analysis to several neuropsychological disorders and other related phenomena in which the normal phenomenal unity of consciousness breaks up. The complexity and independency of several different binding mechanisms that must be at work when our normal unified phenomenal consciousness is constructed in the brain is thus dramatically exposed. Chapter 14 describes how the mechanisms of binding fail in the dreaming brain, producing bizarre combinations of phenomenal content in dream images. However, the failures to bind the dream world together coherently seem to follow interesting regularities which will be revealed with the help of a systematic exploration of dream phenomenology.

Chapter 15 explores the binding problem as it appears in cognitive science and cognitive psychology. Although the discussions in those fields hardly ever mention consciousness, it is possible to find interesting theoretical relations between phenomenal unity and cognitive unity. In fact, I argue that some of the data structures or representations identified in cognitive theories are none other than the subjective phenomenal contents of consciousness in disguise. In cognitive theories, selective attention and binding are closely coupled. Therefore, in cognitive science consciousness is often confused with (or even equated with) attention. I believe this is a mistake. Therefore, I try to clarify how exactly visual consciousness is related to the cognitive concept of visual attention.

xxii Prologue

Chapter 16 explores the neural mechanisms of binding. A large body of empirical evidence from neuroscience shows that the unity of phenomenal consciousness is closely related to neural synchronicity and coherence. These results have important implications also for the search of the neural correlates of consciousness (NCC), a topic to be discussed in greater detail later in the book.

Part V moves downward from the phenomenal level to explore the microlevel mechanisms of consciousness. To define those mechanisms is to give a *downward-looking* or *constitutive explanation* for consciousness. The fundamental problem related to this project is known as the *explanatory gap*. I suggest that the biological research program will be able to build bridges that cross the gap from neuroscience to phenomenal consciousness. In chapter 17 I first point out that the explanatory gap and the possibilities for a reduction between the neural and the phenomenal realities have been discussed traditionally in the context of the deductive-nomological framework of explanation. I argue that it is an inappropriate explanatory framework for the biological sciences and therefore the explanatory gap should be placed in the context of multilevel explanation. The history of science shows that the existence of an explanatory gap in some domain most likely is a sign not of an irremovable mystery nor of the sudden, inexplicable emergence of nonphysical entities, but that there are levels of natural organization involved that have not yet been discovered by empirical science.

The task of the biological research program is to uncover all the levels of organization in the brain that will be crucial for the explanation of consciousness. This task is primarily a challenge for cognitive neuroscience and functional brain imaging. Chapters 18, 19, and 20 deal with the NCC or the experimental search for consciousness in the brain by using the methods and theories of neuroscience. The core question here is, If we find the neural correlates of consciousness in the brain, will we then be able to explain consciousness in neural terms? The view recently defended by several philosophers is that even perfect empirical neuroscience will never be able to explain consciousness, while the view among many cognitive neuroscientists appears to be that the currently available brain-imaging methods are already revealing the neural mechanisms of consciousness in colorful images produced by positron emission tomography (PET) and functional magnetic resonance imaging (fMRI), and that those methods will eventually provide us with theoretical explanations of consciousness. In contrast to these views I argue that, if biological realism is taken seriously, it follows that the currently available methods of functional brain imaging, though useful up to a point, are not adequate to truly discover the phenomenal level of organization in the brain. However, more sophisticated methods might take us to the point where empirical neuroscience is capable of putting forward genuine explanations of consciousness. Since we do not yet have such methods available, we need to stretch our imagination a little to catch a glimpse of the possible future of cognitive neuroscience.

In chapter 21 I speculate how the phenomenal level and its constitutive mechanisms might eventually be discovered in the brain, thereby transforming the cognitive neuroscience of consciousness into "phenomenal" neuroscience. Such a development is not going to be a simple matter because of severe methodological limitations in the measurement, observation, and modeling of higher levels of neural organization in the brain. Although the speculative scenario I delineate may sound like a science-fiction story, we should remember that for the neuroscientists of past generations, our current neuroimaging technology (such as PET, fMRI, magnetoencephalography [MEG], and transcranial magnetic stimulation [TMS]) would have been nothing but far-fetched science fiction.

In the final part VI, (chapters 22, 23, and 24) I explore the *causal powers* of consciousness and the threat from *epiphenomenalism*. The causal role of consciousness is defined at a higher level, where consciousness is itself a part or subsystem of the whole organism behaving in its environment. Thus, in the multilevel framework, defining the causal powers and higher-level roles of consciousness constitutes the *upward-looking* or *contextual explanation* of consciousness.

The dissociation between explicit and implicit—or conscious and nonconscious—information in the brain is considered first (chapter 22): What exactly do we learn about the causal powers of the phenomenal level from the studies on blindsight, implicit perception, nonconscious visually guided actions, and similar phenomena? The data on implicit processing reveal what the brain can do with purely nonconscious stimulus representations when phenomenal content—the subjective experience—normally elicited by the stimulus is selectively carved off, leaving the brain with only nonconscious information to live on. The evidence suggests that these nonconscious input-output transformations of sensory information in one way or another bypass the sphere of consciousness entirely, thereby revealing complex information-processing mechanisms in the brain that in themselves realize no phenomenal level of organization. These nonconscious zombie systems seem to have only limited causal powers in guiding organism-environment interaction, whereas the contribution of the phenomenal level seems to be decisive for meaningful interactions with our environment.

In chapter 23 I explore a variety of fascinating disorders, such as epileptic automatisms and sleepwalking, which temporarily seem to carve off the entire state of being conscious. Thus, these disorders turn the whole person into a nonconscious zombie that continues to roam about in the world, totally oblivious of the fact that within, there is nobody home. A careful examination of such zombies, however, reveals the crucial role of the internal world simulation in guiding external behavior. If the causal contribution of consciousness is wiped out, the remaining nonconscious organism-environment interaction, even if sometimes remarkably complex, is rather pointless all the same. Conversely, other types of disorders show that the simulated phenomenal world in the brain

xxiv Prologue

has unique causal powers in determining the behavioral trajectories of our physical bodies. In the light of evidence from these disorders, consciousness surfaces as a causally potent biological system with unique causal powers. Therefore, we need not worry about epiphenomenalism any longer.

In chapter 24 I explore *the biological function of dreaming:* Why does the brain bother to construct a subjective world-*for-me* during sleep, when the internal world simulation can in no way influence behavior? I propose that the dream reality in the brain was originally a nocturnal simulation of adaptively important events that our ancestors repeatedly were exposed to during the course of evolutionary history. This theory explains most of the conspicuous statistical properties of phenomenal dream content; among other things, why nightmares, bad dreams, and other unpleasant recurring dream themes exist and why they are a universal feature of human dream life; why we all easily fall prey to posttraumatic nightmares after going through life-threatening events in the real world; and why we all frequently dream about being late, failing in studies or work, or about our other everyday worries. The threat-simulation theory postulates that our dream-production system is originally an *evolved biological defense mechanism*. Once we know what to look for in the content of dreams, we find that the dreaming brain still shows clear and abundant traces of its evolutionary function.

All in all, the biological research program on consciousness, based on the core assumptions of biological realism and the world-simulation metaphor, integrates the philosophy, methodology, and empirical study of consciousness into one unified big picture which guides empirical research in new and fruitful directions. I have tried to present the big picture in this book as clearly and in as much detail as possible. The treatment covers philosophy, psychology, and cognitive neuroscience in an attempt to integrate their independent contributions to the study of consciousness and to incorporate the emerging overall view of consciousness into the biological research program.

Still, I am afraid that in many places the treatment remains rather sketchy or superficial, at least for specialists working within the particular disciplines involved. I have had to trespass into the territories of at least the following disciplines: philosophy of mind, philosophy of science, biology, neuroscience, dream research, evolutionary psychology, neuropsychology, and cognitive science. The issues I have included in the book I have judged as being absolutely crucial for understanding consciousness, no matter which disciplines they fall into. Undoubtedly, most of the issues deserve a book of their own for a more detailed treatment. My goal, however, was not to write another purely philosophical book about consciousness for professional philosophers (and most likely almost incomprehensible to empirical scientists), or another purely scientific review of the neural basis of consciousness for neuropsychologists and cognitive neuroscientists (without any discussion of the philosophical issues involved). By contrast, my goal is to cross the traditional disciplinary borderlines to paint with broad strokes a general multidisciplinary picture of consciousness. I hope that the broadness of the overall

treatment is sufficient compensation for the unavoidable lack of detail and depth at some points.

The problem of consciousness, the most fundamental question, if ever solved, will be solved in the context of some sort of systematic empirical research program that takes phenomenal subjectivity very seriously. Biological realism looks to me like a respectable candidate for establishing exactly such a research program. The big picture that I present in this book is my idea as to how we might reach a more unified science of consciousness. I hope that some of my readers will be able to share this idea with me or, if not, at least point out where I have gone wrong.