

## I Coming-to Explained



Chances are it is less than a day since you regained consciousness. It probably happened soon after the sunlight returned this morning. What was it like for you, as you *came to*? Remember? The chink of a milk bottle, the touch of sheets, the sight of a patch of blue sky. You rubbed your eyes, stretched your limbs, and before you knew it, waves of *sensation* refilled the lake of your being. *You* re-emerged into the *subjective present*. Once more you *felt* yourself alive.

You were not alone. Something like this happened today to countless other individuals here on Planet Earth. Our planet, we are told, is merely a condensate of stardust, not so different from all the other minor cosmic bodies that litter the universe. But this one planet has become home to an extraordinary phenomenon. Here is where *sentience* evolved. Here is where *conscious selves* have come into their own. Here live *souls*.

■ In this book I will address the questions of what *sentience*, *selfhood*, and *soulfulness* amount to. In the course of it I will propose a solution to the “hard problem of consciousness.” The hard problem is to explain how an entity made entirely of physical matter—such as a human being—can experience conscious feelings. The problem is *hard* because such feelings appear to us, who are the subjects of them, to have properties that could not possibly be conjured out of matter alone. We say—because we do not know what else to say—that “it’s like something” to be conscious. Yet, the problem with this inadequate phrase, “it’s like something,” is that *what it is like* seems to us—no, *is* to us—unlike anything else out there in the material world.

There are philosophers who think the problem is simply too hard to admit of a solution. For Colin McGinn, trying to explain phenomenal consciousness as a product of the brain is like trying to explain how you can get “numbers from biscuits, or ethics from rhubarb.”<sup>1</sup> For Jerry Fodor, “We can’t, as things stand now, so much as imagine the solution of the hard problem. The revisions of our concepts and theories that imagining a solution will eventually require are likely to be very deep and very unsettling. . . . There is hardly anything that we may not have to cut loose from before the hard problem is through with us.”<sup>2</sup>

I disagree. I acknowledge, of course, that theorists have not been doing too well in imagining the solution. I am as impressed as anyone by what *seem* to be the insuperable difficulties. But I suggest we attend to the word “seem.” The fact that something *seems* to have mysterious and inexplicable features does not necessarily mean it really has them.

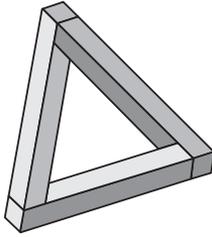


Figure 1.  
The Penrose Triangle.

Let me illustrate the difference between *seeming impossible* and *being impossible* with the help of a well-known example. Suppose you were to come across a solid wooden object that looked just like the object shown in figure 1, Penrose's "impossible triangle." Certainly, it would seem to be a physical impossibility. Yet no one would say that just because of what the object *looks like* you should throw away your physics books and cut loose from everything you know. You would soon realize, of course, that it must be an illusion. And sure enough, if you could only *change your viewpoint*, you would discover that what you are actually looking at is the curious object shown on the next page in figure 2. This object was cunningly constructed by the psychologist Richard Gregory, precisely so that, when it is seen from a certain position, *it creates the impression* of an impossible triangle. This object deserves a name. With Gregory's permission, I call it the "Gregundrum."<sup>3</sup>

If you were to come across the Gregundrum lying on a laboratory bench, without knowing its "function," I am sure you would never guess that it holds the key to anything interesting. It is certainly not a pretty thing in its own right. Who would have thought that such a perfect thing as the Penrose triangle could have such an ugly explanation? Yet, as Sherlock Holmes said to Dr. Watson, "When you have eliminated the

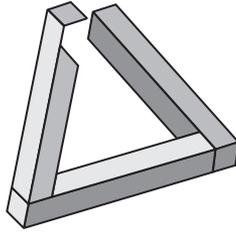


Figure 2.  
The Gregundrum.

impossible, whatever remains, however improbable, must be the truth.”<sup>4</sup>

I will argue that the truth about consciousness—if and when we see it from the right perspective—is that it is indeed the product of a highly improbable bit of biological engineering: a wonderful artwork of nature that gives rise to all sorts of mysterious impressions in our minds, yet something that has a relatively straightforward physical explanation. As Holmes went on, “We know that he did not come through the door, the window, or the chimney. We also know that he could not have been concealed in the room, as there is no concealment possible. Whence, then, did he come?” “He came through the hole in the roof,” Watson cried. Our job as consciousness researchers is to find the hole in the roof.

I do not say it will be easy. To start with, in an area where theorists continually talk past each other, there will be issues about the use of *words*. To forestall at least some potential verbal misunderstandings, I have set out in the box a rough guide to the conceptual territory as I see it. (You should not get hung up on anything in this list at this stage—I will justify and explain these definitions further as we go on.)

But it is not just words that may come between us and the truth; it may be the deep-seated biases that we bring to the table as subjects of consciousness ourselves. We cannot of course

- In general, when I talk about consciousness I mean “phenomenal consciousness.”
- A subject is “phenomenally conscious” (or plain “conscious”) when and if there is something *it’s like to be him* at this moment.
- There is “something it’s like to be him” when he experiences *feelings*, or what philosophers call *qualia*.
- Qualia—for example, the felt redness of fire, the sweetness of honey, the pain of a bee sting—are features of *sensations*.
- The subject is “phenomenally conscious” just when he experiences sensations *as having* these peculiar features.
- To experience sensations “as having” these features is to form *a mental representation* to that effect (with the meaning of “represent” still to be decided).
- Thus “consciousness” (or “being conscious”), as a state of mind, is the *cognitive* state of entertaining such mental representations.
- Consciousness can change the subject’s life just to the extent that these representations *feed forward to influence what he thinks and does*.

opt out of our privileged position, but we can at least try to imagine where we would be without it. To that end, I want to begin our investigation of the problem by handing it over to someone else, someone who should have a remoter and more objective view of what consciousness is doing for us than we ourselves have.

■ Let us return, then, to this morning. Only now imagine that a few hundred miles out in space, a visiting scientist from an advanced civilization in the Andromeda galaxy is orbiting our planet, on her first trip to investigate life on Earth. (I call her “her” because I assume the Andromedans long ago dispensed with the male sex.)

Situating her craft so as to get a good view of the boundary as night turns to day on the Earth’s surface, she observes how, all along this boundary, living creatures are emerging from their nighttime coma. Birds are breaking into song, butterflies are taking to the wing, monkeys are leaving their beds in the trees, and human beings are going downstairs to brew their morning coffee.

She observes this great awakening, and she nods knowingly. She recognizes, of course, that the central processors that run these earthlings’ onboard software have been in sleep mode overnight, so as to save energy and perform system maintenance. And now, with the sun’s rays bringing light and warmth, it is time for them to resume their life tasks. As a scientist, she has much to look forward to. Once she gets down among these creatures, how interesting it will be for her to study their brains and behavior and figure out how it all works. Indeed, she fancies herself as a bit of a philosopher: one day she will write a book called *Coming-to Explained*.

Our visitor has every reason to trust the scientific method. Wherever else in the universe she and her colleagues have applied it, natural phenomena have given up their secrets. No doubt, she reckons, there can be nothing so different or difficult about those living organisms down there on Earth.

But is she right? What about *consciousness*? Will it ever dawn on the Andromedan visitor that there is a dimension to

the lives of at least some of the creatures she is studying that needs special treatment, that when they “come to,” it is as if a light is coming on *inside* their heads? Given that she can see things only from outside, is it possible that she will miss this altogether, that she will never even suspect that consciousness exists?<sup>5</sup>

■ I think we should assume the Andromedan does not have the circuits in her own brain that would make her phenomenally conscious herself. Otherwise we will not know how to assess any claims she may make to have *discovered* the existence of consciousness in other creatures. (She might just be arguing from analogy with her own case, in the way you or I might argue, for example, that it is obvious that a dog feels pain the way we do.)

The absence of phenomenal consciousness may or may not affect the way she thinks about philosophical and scientific issues (this is something we should be better placed to judge by the end of the book). But I see no reason, as of now, why it should place any limit on her intelligence (“artificial intelligence,” as we might want to call it) or her skills at scientific research. Let us suppose, indeed, that she does have an exceptionally brilliant analytic mind. And let us allow her every other scientific advantage anyone might ask for. She can undertake meticulous behavioral studies of how earthlings behave in the wild, and then follow up this fieldwork with whatever laboratory investigations are suggested. She has all the research instruments she could possibly need: scanners and imagers and calculators of a sophistication yet undreamed of here on Earth. She can prod and probe and listen in and cross-question. She can, if she wants, take the earthlings apart

and examine their machinery (the Andromedan ethics committees have no objection to alien vivisection). Then, back home, she will be able to run theoretical simulations on her computer and build working models in the robot shop.

Then, what *will* she discover, and what will she not? Let us consider some possibilities.

- She will find, to her surprise, that in order to explain the *behavior* of certain species of earthlings, she needs to postulate the existence of an extraspecial mental state—a state with peculiar qualitative properties, unlike anything else, which *just because of what it is like* is changing how these creatures live their lives.
- Though perhaps she will be unable to deduce the existence of any such special inner state from what she observes of public behavior, she will nevertheless realize that such a state exists when she examines in detail the flow of information in the earthlings' brains and figures out what kind of *private mental representations* are being generated.
- She will do better still. Beyond simply discovering the existence of conscious states, she will be able—either from behavioral observations or from brain scans—to arrive at a complete description of what it is like to be the subject of a particular state. Perhaps she will even get to the point where she can compare one individual's state with another's—so that she can tell, for example, whether different subjects are experiencing the sensation of red in the same way.
- Or then again, perhaps she will be able to do none of the above.

Now, as it happens, there are a good many students of consciousness here on Earth—they may even be in the majority—who believe the answer can be only the last of these. In their view our visitor will fail to discover *anything* about consciousness by any of the scientific means at her disposal because of an awkward but undeniable truth: consciousness, for all its subjective importance, *is physically featureless; it does not show.*

The psychologist Jeffrey Gray has written, for example, “Nothing that we so far know about behaviour, physiology, the evolution of either behaviour or physiology, or the possibilities of constructing automata to carry out complex forms of behaviour, is such that the hypothesis of consciousness would arise, if it did not occur in addition as a datum in our own experience; nor, having arisen, does it provide a useful explanation of the phenomena observed in those domains.”<sup>6</sup>

Others have gone further still, arguing for what the philosopher Owen Flanagan has called “consciousness inessentialism”—“the view that for any intelligent activity *I*, performed in any cognitive domain *d*, even if *we* do *I* with conscious accompaniments, *I* can in principle be done without these conscious accompaniments.”<sup>7</sup> Thus, according to John Searle, “We could have *identical behavior* in two different systems, one of which is conscious and the other totally unconscious.”<sup>8</sup> There could even exist a “philosophical zombie human,” David Chalmers has suggested, who is physically identical to a normal human being and who looks and acts in every respect just like one, yet who is not phenomenally conscious—“all is dark inside.”<sup>9</sup> Then, if you or I were to meet such a philosophical zombie in the street, we would not—and could not—know the difference.

True, each of us is presumably convinced that consciousness exists in our own case, and therefore we may want to give the benefit of the doubt to others who so obviously resemble us. But the Andromedan scientist does *not* know about consciousness from her own case. Therefore, if and when *she* notes resemblances between herself and any of the earthling creatures she is studying (those naked bipeds who seem to have taken over the planet are certainly technologically ingenious!), she is likely to assume they resemble her in this respect as well. And if consciousness inessentialism is right, she will not discover anything in the course of her research to make her revise her opinion. At the end of the day, she will not think she has missed anything. So she will return to Andromeda—and write her book—with a satisfied sense of mission accomplished: “Coming-to Explained Away.”

I said I wanted to hand over the investigation of the hard problem to this visitor, because we might expect her to have “a remoter and more objective view of what consciousness is doing for us than we ourselves have.” But if this is really how things stand, it seems the problem will not even cross her horizon. Fodor wrote, “There is hardly anything that we may not have to cut loose from before the hard problem is through with us.” He cannot have meant this interpretation, but is the lesson that if we want to keep up with the best science in the universe, we ought to cut loose from the concept of consciousness itself?

■ You will realize—if for no other reason than because my own book does not end here—that I do not think so. My starting point is that consciousness, however elusive and enigmatic from a scientific perspective, is a fact of nature. And if

it is not *evidently* a fact of nature, that can be only because scientists and philosophers have been looking for evidence in the wrong places. I believe this because I think the idea that consciousness has no observable effects is daft (and the notion of a “philosophical zombie”—a physical duplicate of a conscious human who completely lacks consciousness—is dafter still). However, I have to say I do not think it is daft to suppose that certain aspects of conscious experience could have no observable effects. So, before we go further, I want to consider just to what extent conscious experience will—and will not—be observable to an outsider.

We know, of course, that not everything that goes on in the mind of a person or an animal has to show up *in behavior*. There can obviously be purely private mental states. Indeed, most ordinary mental states are private, insofar as they occur without anyone’s—except the subject—knowing about them. No one but you knows what your thoughts are right now (why else would anyone give you a penny for them?). No one but me knows about my dreams last night (and, as it happens, even I do not know any longer).

Still, we might want to argue that states such as these are only contingently private. If you were given the penny, you *could* tell me what your thoughts are. If I had kept a dream diary, I *could* have shared my dream with you. And even without language, there would probably be ways of communicating much of the content of these mental states.

But that is *thoughts*. And with *feelings* it would seem to be a different matter. How about basic sensory experiences? They undoubtedly seem to be more absolutely private. You would be hard put to it, however much you tried, to reveal the full content of what it is like to experience the smell of a rose or

the coldness of a snowball. Though you could surely communicate some part of it, you would not know how to capture the subjective quality of the sensations, the qualia.

It is by no mean obvious exactly what the problem is. Is it that there is something about the logical status of qualia, as intrinsically subjective properties, that makes them incommunicable in principle? Or is it simply that in practice we do not have the requisite communication skills? Could it even be that our minds have been designed to have some kind of fire wall around sensory experience which puts adaptive limits on what others can discover about us?

There could be some truth in all these possibilities. But whatever is causing the problem, we must surely accept that there *is* a problem; we must concede that in practice, even if not in principle, conscious sensations are private in crucial respects, so that nothing the subject can say or do can reveal everything about them.

However, I would say this is *all* we need concede. We need not—and should not—accept either of two stronger propositions, namely, (1) while an outside observer is restricted to studying behavior, she will not even be able to detect that phenomenal consciousness *is present*, and (2) even if the observer were allowed complete access to the subject's brain, she would not be able to discover *the full content*.

■ Let us look at these two issues. First, why do I believe that consciousness must reveal its *presence*, if nothing more, at the level of behavior?

The reason is the ultimate one, the hand of natural selection. Since consciousness, as we know it, is a feature of life on earth, we can take it for granted that—like every other spe-

cialized feature of living organisms—it has evolved because it confers selective advantage. In one way or another, it must be helping the organism to survive and reproduce. And of course this can happen only if somehow it is changing the way the organism *relates to the outside world*.

Now, how could this be happening? Conscious creatures do not smell different or look prettier. Consciousness does not provide extra strength or better health. Instead, consciousness can have its effects on survival only by changing what we may loosely call the creature’s “psychology.” In other words, being phenomenally conscious must be influencing how the creature *thinks* or what it *wants* or what it *believes*, in just such a way that it now acts in the world in adaptive ways it would not have done otherwise.

In later chapters I will explore in detail just how this may be working: how the effects may be present on several levels, and how they may be more or less important for different kinds of animals, pushing the evolution of consciousness along species-specific lines. As we will see, human beings, with their developed sense of “conscious self,” are most likely in a class of their own. But the important point for now is that if natural selection can “see” the effects—whatever they are—of the changed psychology on behavior, presumably so too can other outside observers (if only they knew where to look). What is more, if these observers can see what natural selection sees, they should also be able to see what it is about it that is beneficial—and hence why natural selection has favored it. Thus they should be well on their way to constructing a story about *why* consciousness evolved.

Still, do not get me wrong about this. I am not suggesting that because consciousness has been designed by natural

selection, this means that every one of the features of the design must be able to be seen from the outside. Rather, what it means is that every one of these features must be contributing to the beneficial effects that natural selection *does* see.

It would be easy to misunderstand this, so let me tell a parable to make it clearer. Imagine that in a certain country the government has a Department of Happiness, whose minister has the job of maximizing the general happiness of the population. The minister must therefore be on the lookout for things that put people into a good humor. One day he comes across a group of people who are looking at a cartoon and smiling broadly. From where he stands, the minister cannot actually see the picture they are looking at, and so he does not get the joke. Still, he can see the positive feelings the cartoon is eliciting. And that is enough to persuade him to take departmental action to “breed” this cartoon by ordering additional cartoons in the same style. So he does this, and the next day he sees more people laughing at the new drawings. He repeats his order, and soon enough cartoons in this special style are everywhere. The style has become, as it were, a ministerial design feature. But note that at no point has the minister himself needed to know what the cartoons look like. All he has needed to see is evidence that the cartoons exist and that they are funny.

My point is that, likewise, natural selection need never have known what conscious experience is actually like for the subject. All it must have seen is evidence that conscious experience exists and that it is in some way life enhancing. This being the case, it is possible and even quite likely that the detailed phenomenal content of sensations will not ever have been evident in behavior. And so today our visiting scientist, while she

relies on outside observations, will be able to get only halfway to discovering the facts of consciousness. She should certainly be able to detect that the special inner state exists in some creatures and that, in whatever way their behavior suggests, it adds to their success in life. However, this may well be as far as she can go.<sup>10</sup>

■ Yet, what if she were able to *search inside their heads*? Why do I believe that an observer who can go beyond behavior down to the level of brain activity should be able to discover *all* there is to know?

My reason is simply the guiding principle, which underlies all science, that *nothing interesting occurs without a material cause*. In short, *miracles do not happen*. When conscious experience arises in a person's mind, it is the outcome of events in the brain. Moreover, if and when these events (in their totality) occur, the outcome *has to be* that the person is conscious (which is why the idea of a philosophical zombie makes no sense). Thus, if a scientist can go inside and observe these crucial events, she should be able, in principle, to deduce what the outcome is—provided only that she has a *theory* linking brain states to experience, a theory that enables her to move from one level of description to the other.

What kind of theory would this be? Philosopher Dan Lloyd has written: “What we need is a transparent theory. One that, once you get it, you see that anything built like *this* will have *this* particular conscious experience.”<sup>11</sup> We can draw an analogy with explaining the properties of water. Scientists are able to deduce that a pail of molecules, whose chemical composition is H<sub>2</sub>O, at room temperature will have the physical properties of the substance we know as water (fluidity,

wetness, and so on) because, with their understanding of the laws of physical chemistry, they have a theory of why *water under its chemical description* must amount to *water under its physical description*.

Then, so too, we may reasonably hope that if and when scientists have a comparable understanding of the laws of what we may call neurophenomenology, so that they have a theory of why *brain activity under its neuroscientific description* must amount to *mental activity under its experiential description*, they will be able to deduce that, for example, a man whose brain is in a particular state is a man who is thinking such and such thoughts.<sup>12</sup>

It is already widely agreed by those who study mind-brain relationships that it is the pattern of information flow in the brain that determines mental states. I would say we can assume therefore that the neurophenomenological laws will essentially be laws about *how experience is computed*. Admittedly, apart from having this one insight, our scientists here on Earth are nowhere near to discovering what the laws actually are. Still, we need not doubt that the laws exist and will eventually be found out. So, to continue with our story of the Andromedan scientist, let us imagine that the theorists on Andromeda are far more advanced than ours are, and—in anticipation of their sister’s mission (or perhaps just for the fun of it)—they have worked out ahead of time the relevant laws as they apply to alien brains.

Thus, let us suppose the Andromedan scientist has arrived among us prearmed with the theoretical tools she needs for interpreting earthlings’ brain activity in experiential terms. Where will this take her? Given what was said above, we may assume that, on the basis of her purely behavioral observations,

she will already have concluded that in some of the earthlings under study (notably, human beings) there does exist a special inner state that is influencing their outlook on life—though a state of which the detailed content is so far a mystery to her. But now that her brain research is under way, she will, with the help of the theory, be able to deduce that these particular subjects are having experiences with exactly the weird and wonderful phenomenal content that you and I know so well firsthand.

“Well, blow me!” she may say. “Who’d have guessed it?” For she will indeed have deduced the existence of qualia. She will, as it were, have arrived at a complete description of the private joke that lies behind the public smile.

■ Are you with me still? Or do you think I have tried to pull a fast one on you (in fact, did I not try to pull it a few pages back)? Can it be true that the Andromedan—who is not conscious herself, remember—has discovered what consciousness is *really* like? Or has she merely discovered its pale shadow?

The big question, you may insist, is whether the scientist, when she examines the brain of someone who is having a conscious sensation, can deduce what that person’s experience *actually is*, and not merely deduce a *description of what that experience is* (and calling it a “phenomenological description” simply begs the question).

But, no, I have not pulled a fast one. Rather, if you make this objection, I would say you have just pulled a fast one on yourself. You have fallen for the tempting idea that there is something conscious experience *actually is* that is separate from what the subject *thinks* it is—that is, the *mental representation* that he makes of it. But it is not so. If you do not see

this now, I hope to persuade you of it as we go on. To give a foretaste of what is coming, in the very next chapter I will argue that what I called at the start of this book the inadequate phrase “it’s like something” is not such a bad phrase after all. Because, when it comes to it, for a subject to have a sensory experience that *is like something* really is for him to represent the object of experience *as if it is something* with some very peculiar features. In short, for the subject to have a sensory experience that is like something is just for him to experience it *as what it is like*.

The philosopher John Searle (with whom, on the question of consciousness, I agree about very little) put his finger on this point precisely when he wrote: “If it seems to me exactly as if I am having conscious experiences, then I am having conscious experiences.”<sup>13</sup> Just so. “Seems to Searle exactly as if” can only mean “is represented mentally by Searle exactly as being.”

What follows from this? Since mental representations can, in principle, always be described or re-represented in some public medium—they would not count as representations otherwise—it surely follows that, despite what was said above about the *de facto* incommunicability of private experience, it must be possible in principle to *describe* what it is like to be conscious.

It is undeniably true that, as of now, we humans do not know how to do this satisfactorily. We lack both the theory and the language for the job. But these, we should assume, are contingent limitations—already overcome in Andromeda and soon enough to be overcome back here on Earth.

I would say we should acknowledge that the phenomenological descriptions of conscious experience that will feature in the final theory will probably require a new vocabulary,

even a new grammar.<sup>14</sup> But we should not be too alarmed by this, let alone see it as a reason for giving up. It has happened before in the history of science that scientists required a new conceptual language before they could move on—and yet, after initial awkwardness and even disbelief, soon enough everyone gets used to it. Think, for example, of how mathematics has had to come to terms with “complex numbers” involving the square root of minus one, or of “transfinite numbers” that are bigger than infinity. Think of how physics has had to come to terms with relativity.

Future descriptions of conscious experience will almost certainly require concepts that sit oddly with our standard ways of thinking today. I already remarked at the opening of this chapter that the problem with saying “it is like something” to be conscious is that *what it is like* seems to us—no, *is* to us—unlike anything else out there in the material world. The phenomenal experience of the “subjective present” as existing in “thick time”—as I have attempted to describe it elsewhere<sup>15</sup> and as I will revisit shortly—is perhaps just such an apparently *essential yet nonsensical* concept.

Yet, let us stick with our story. We have assumed that scientists on Andromeda are well ahead of us in recognizing the neurophenomenological laws. Contained within this assumption must be the assumption that they have already developed a suitably esoteric language for describing conscious experience (even if the development of this language must have been, as it were, “on spec,” since the Andromedans, having never encountered creatures such as human beings before, cannot yet have had occasion to apply it).<sup>16</sup> So we are assuming that our visitor will have the tools for describing what it is like for us, even if we humans at present do not.

However, I do not want to make our own inadequacy an absolute sticking point. To claim—as many philosophers would—that consciousness is *essentially ineffable* is to underestimate human ingenuity and creativity. As we will see later in this book, humans may be more capable of expressing publicly what it is like to be conscious than the philosophical and scientific skeptics would have us believe—though when they do so, they “cheat” by using the language of *art* rather than that of *science*. Well, we will see.<sup>17</sup>

■ This introductory chapter, which started off so breezily, is getting heavy. It is time to sum up—and lighten up, if possible.

I wanted the Andromedan scientist’s help with the project of understanding the hard problem—the nature of consciousness—because I hoped that to see the problem from her perspective might provide us with some useful guidelines for our own inquiry. Whatever the differences between us and her are, I take it that *science is science* wherever in the universe it is being done. What counts as evidence and conclusions for this researcher from a far-off galaxy should be what counts as such for us on Earth. That is why I asked above what will the Andromedan find out about consciousness, and what will she not. I assume this is what, at the limits of our human abilities, we can expect to find out too.

Here is the score.

We have established that the Andromedan scientist will be able to discover at a behavioral level crucial hints that consciousness is present in some creatures. At the very least, she will discover that consciousness is having certain beneficial effects—these are the effects on which natural selection has been acting in the course of evolution. She will discover that

consciousness *exists* and—in the larger picture—what consciousness *is for*.

Nevertheless, while she stays on the outside, she will probably be unable to reach a deep understanding of the contents of consciousness. This is because the crucial features of what it is like for the subject will, in normal circumstances, probably be hidden from public view—even though these features are ultimately responsible for the beneficial effects.

To find out more of the details, she will have to go inside. When she does so, using all the neuroscientific techniques at her disposal, she should indeed be able to discover *everything* about what it is like to be conscious, provided she has a *theory*. But this neurophenomenological theory will have to be a new and remarkable theory: not a theory that we human beings can never get to understand (as some philosophers, notably Colin McGinn, have suggested),<sup>18</sup> but certainly a theory we will not understand until we have put in some more work.

So now, let me set out my agenda for this work and my book. What I plan to do is to emulate, in my own way, the Andromedan scientist's investigation. Yet, because, first, I am not as clever as she is, and, second, I am a living example of the phenomenon under investigation, my strategic goals will be a little different.

On Andromeda, I have suggested, scientists have already developed the theoretical tools for solving the hard problem of how matter could in principle give rise to consciousness, even if they have never yet come across a case of consciousness in fact. By contrast, we humans know consciousness exists in fact, but we do not at the moment have the theory of it. The first task for the book, then, must be to come up with at least the beginnings of a plausible theory of what consciousness *is*

and how it relates to the brain. To do this I will, in the next few chapters, argue for a radically new account of what we mean when we say that “it is like something” to experience sensations. I will make a proposal as to what *the thing in the brain* that the subject represents as “being like something” really is, and I will suggest what its biological origins in nonconscious animals may have been.

The Andromedan scientist, I have assumed, being completely new to the world of conscious creatures, will, at the start, have no idea what difference consciousness is making at either the private or the public level, let alone what good, if any, comes of it. By contrast, we humans know rather a lot about the difference that consciousness is making to our private lives, though we are far from understanding how this translates into public benefits. The second task for the book then will be to figure out—knowing what we already do—how being conscious changes people’s psychology (and perhaps that of other conscious animals as well) in ways that ultimately increase their chances of survival.

Having read this far, you may be nervous that the book is going to be unduly scientific. Do not worry. There is indeed work to be done. We need to get the science right if we can. But my book is called *Soul Dust*, and it will live up to that title. The book will continue with some hard-going philosophical analysis, but it will end with a fairy tale—a scientifically based fairy tale—about how consciousness lights up the world.