time in evolution, information acquires a new potential—the possibility of subjectivity. It is information "for somebody"; in short, it becomes consciousness itself.

We have claimed that the sophisticated forms of information exchange developed by humans would not be conceivable in the absence of consciousness. After Homo sapiens and higher-order consciousness appeared, it became possible to create syntactically rich symbol systems, to create codes, and even to create logic. Eventually, methods of scientific analysis were invented, resulting in the formulation of natural laws. To us, those laws are information. To nature, outside of us, is it energy or coded information that is being exchanged? Does it come from bit or bit from it?[3] What came first: biology or logic?

SELECTIONISM AND LOGIC

We live in a world dominated by computers and computing. It has become commonplace these days to think about the brain as a computer, a device based on logic. Although we do not believe that this is a tenable point of view, it does pose an interesting epistemological question: At the most fundamental level, how many modes of thought are there: Is logic the only one?

In a purely formal sense, some philosophers have defined logic as the study of all sentential relations invariant with respect to lexical substitution. "All A are B. X is an A; therefore, X is a B," no matter what A, B, or X refer to. In a broader, more psychological framework, logic may be considered the study of the relationships between such formal issues and intuition (or pattern recognition). This latter framework is much less tidy, but it does prompt one to ask: What is the connection between the ability to recognize and match patterns, the ability to think, and the ability to carry out logical operations? This last ability can be mechanized, but as Gödel's theorem implies, there are certain patterned mathematical relationships whose truth cannot be proved or disproved within a consistent axiom system. Nevertheless, the argument has often been made that the brain is some kind of computer and therefore is describable as a Turing machine.

A universal Turing machine can carry out any sequence of logical operations and, according to the Church-Post thesis, can carry out any sequence of effective procedures or precisely specified algorithms (see figure 17.1). This is a powerful capability, and we suppose that it is this capability that has prompted the suggestion that the brain is a Turing machine. We have considered elsewhere why this cannot be the case. Briefly, the argument goes as

follows: Each brain is formed in such a way that its wiring and dynamics are enormously variable at the level of its synapses. It is a selectional system, and each brain is therefore unique. That uniqueness and unpredictability can be significant in carrying out certain brain operations and must be taken into account in considering any particular brain operation. Moreover, brain function is degenerate: When presented with an unpredictable context, nonisomorphic brain structures at more than one level of construction and operation can lead to the same output or function. Furthermore, many brain operations in perception and memory are nonrepresentational, constructive, and context dependent and are not necessarily guided by an effective proce-

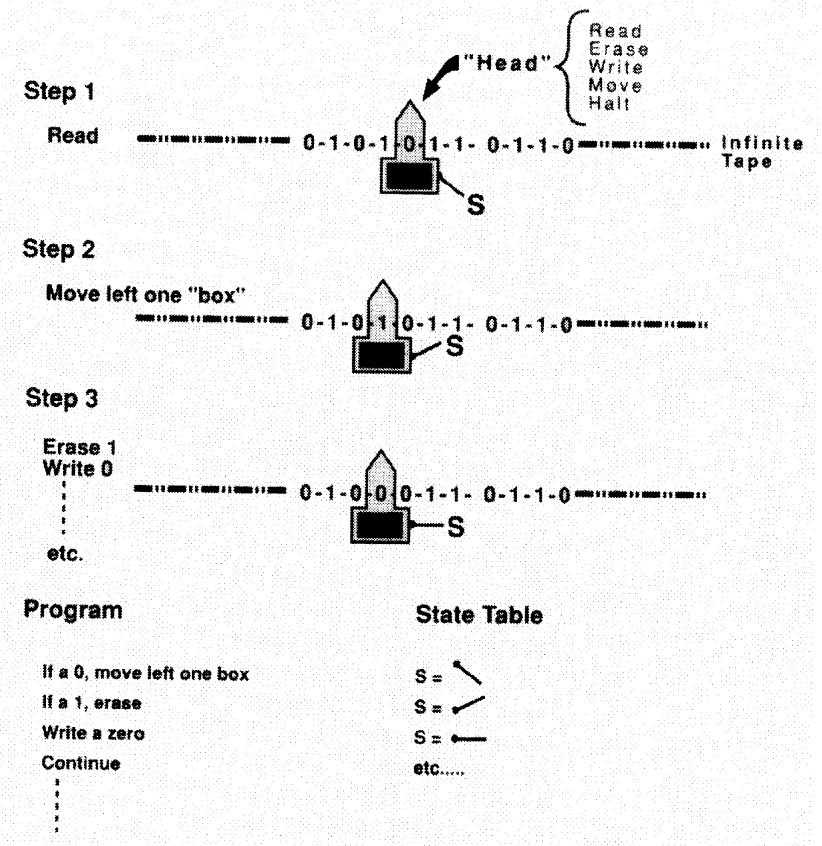

FIGURE 17.1 DIAGRAM OF A TURING MACHINE. When a program designates an action, as shown in the figure, the machine state *S* changes, and this change alters the next step prescribed by the set of precise instructions. Brains do not work this way.

dure. This is so because their key operations involve selection, not instruction, and there is no evidence for rigorous or preestablished neural codes resembling those of computers. Finally, the environmental input and the context of signals to the brain are not uniquely specified or single valued; that is, the world, although it obeys physical laws, does not behave like a computer tape.

If the brain is not a Turing machine, we need another explanation for its workings. That explanation is provided by the theory of neuronal group selection (TNGS). As we have seen, a series of simulations based on the TNGS can actually carry out pattern recognition and perceptual categorization. Moreover, a large body of disparate experimental evidence not only indicates that selectional events occur in the brain, but suggests that much of that apparently disparate evidence can actually be reconciled by an analysis of such events. We have taken it as established that after the brain arose in evolution by natural selection, which set up value constraints and major structures, each individual brain operates by a process of somatic selection. Instead of being guided mainly by a set of *effective procedures*, it is governed by a degenerate set of *effective structures*, the dynamics of which allow its correlated activities to arise by selection, rather than by the rules of logic.

Clearly, if the brain evolved in such a fashion, and this evolution provided the biological basis for the eventual discovery and refinement of logical systems in human culture, then we may conclude that, in the generative sense, selection is more powerful than logic. It is selection—natural and somatic—that gave rise to language and to metaphor, and it is selection, not logic, that underlies pattern recognition and thinking in metaphorical terms. Thought is thus ultimately based on our bodily interactions and structure, and its powers are therefore limited to some degree. Our capacity for pattern recognition may nevertheless exceed the power to prove propositions by logical means. Indeed, conscious human thought can create new axioms, which a computer cannot do. This realization does not, of course, imply that selection can take the place of logic, nor does it deny the enormous power of logical operations.

In the realm of either organisms or of the synthetic artifacts that we may someday build, we conjecture that there are only two fundamental kinds—Turing machines and selectional systems. Inasmuch as the latter preceded the emergence of the former in evolution, we conclude that selection is biologically the more fundamental process. In any case, the interesting conjecture is that there appear to be only two deeply fundamental ways of patterning thought: selectionism and logic. It would be a momentous occasion in the history of philosophy if a third way were found or demonstrated.

PHILOSOPHICAL CLAIMS

Whatever the powers and limits of a thought process that is embodied and is the result of natural and neuronal group selection, that thought process has given rise to a remarkable propensity for asking questions about the ultimate significance of things—questions of philosophy. Two grand areas of philosophy are metaphysics, which is concerned with the ultimate nature of reality, and epistemology, which is concerned with the basis and justification of knowledge and belief. These two areas are linked to some extent, and they touch on others of human concern that are connected to values, for example, ethics and esthetics. Does a theory of consciousness have implications for our approach to such questions? We believe it does.

Perhaps the best way to support our belief is to consider certain matters of epistemology and metaphysics and their connection to scientific matters. We believe that there is a real world, the one described by the laws of physics, which so far appear to apply everywhere. As humans, we must follow such laws without exception because we have evolved in that world from ancient animal origins. As living systems, we are also subject to evolutionary constraints not considered by the laws of physics. Consciousness, while special, arose as a result of evolutionary innovations in the morphology of the brain and body. The mind arises from the body and its development; it is embodied and therefore part of nature. All these statements follow from what we called, in chapter 2, the physics assumption and the evolutionary assumption, both of which underlie our theory.

Once we agree that these assumptions are justified by a large body of evidence, a number of consequences can be identified. First, we must reject the Cartesian assumption of dualism and any form of idealism. Thus, we cannot accept the position taken by those who embrace a materialist metaphysics combined with a dualist, rationalist or idealist, epistemology. On the other hand, we must be skeptical about extreme reductionist accounts that attempt to explain consciousness on the basis of quantum mechanics but ignore the facts of evolution and neurology. The same skepticism holds for attempts to imbue the world at large with conscious properties—the view of panpsychism.

The metaphysical and epistemological positions that we do espouse we have called, respectively, qualified realism and biologically based epistemology.[4] A key idea on which qualified realism and biologically based epistemology both depend is the notion that concepts are not, in the first instance, sentential. That is, concepts are not propositions in a language (the common usage

of this term); rather, they are constructs the brain develops by mapping its own responses prior to language. Just as signals from the world are not organized as information before interaction with the brain, neither is language specified in terms of a genetically inherited universal grammar. Concepts, in our view, precede language, which develops by epigenetic means to further enhance our conceptual and emotional exchanges.

Consciousness is a dynamic property of a special kind of morphology—the reentrant meshwork of the thalamocortical system—as it interacts with the environment. Our knowledge of the real world comes as a result of the physical, psychological, and social interactions of our minds and bodies with that world. Those interactions do not involve a direct transfer of information, however, and we must therefore reject naive realism, the position that the perception of objects is direct and that the qualities we perceive are, in fact, those of the objects that are perceived. Realism based on perception must be qualified by the bodily means we have available for perception. Most of these means, while enormously powerful, are nevertheless indirect and of limited range. They constrain how our brains develop their conceptual systems, and we therefore conclude that our realism must be qualified, at least to some extent.

Of course, as conscious persons capable of language develop and communicate in a culture, their conceptual abilities become enormously enriched. The products of that enrichment, for example, logic and mathematics, can transcend some of the phenotypic constraints on the embodied mind and remove some of the qualifications imposed on us by the limitations of our phenotype. Indeed, perceptual categories based on neural structures, buttressed by conceptual categorization and aided by conscious planning and choice in a linguistic milieu, have led to the extraordinary systems of thought that underlie the scientific investigation of the world. The important point that emerges from the work reviewed in this book is that a scientific investigation of consciousness is also consistent with the facts of human individuality and subjectivity.

Before the advent of modern experimental psychology and neuroscience, epistemology was based solely on normative issues and thinking about thinking. But from the time of Charles Darwin on, and in light of recent scientific advances, it has been proposed that epistemology should be "naturalized" and take its grounds from a behavioristic psychology.[5] Naturalized epistemology necessarily stops, however, at the stimulation of receptor sheets—the retina, the skin, the taste buds—and while it includes the analysis of language, it leaves the inner workings of the body and brain

untouched. We claim that this position is insufficient: Epistemology should be grounded in biology, notably neuroscience and a theory of consciousness, which, of course, includes psychology. We therefore reject the effort by philosophical behaviorists to naturalize epistemology via behavioral psychology alone.[6] A move to a biologically based epistemology would not only recast certain arguments about, for example, the possibility of a "synthetic a priori," but would create a much broader base for thinking about thinking and feeling. Moreover, it would not limit our descriptions to the boundary between our skin and the rest of our world. Most important, it would open our inquiry to include feelings and emotions in terms of bodily mechanisms that go far beyond computation.

While these attempts give due scientific recognition to the subjective domain, subjectivism itself is no basis for a sound scientific understanding of the mind. Consequently, we reject phenomenology and introspectionism, along with philosophical behaviorism. We believe that the inner mechanisms of consciousness can be explored scientifically without exclusive resort to either simple behaviorism or introspection. The first puts the observer outside the key phenomena, while the second assumes falsely that by taking thought alone, one can analyze the underlying bases of conscious experience.

Taking the position of biologically based epistemology changes how we look upon that fictive but useful agent, the scientific observer. As Erwin Schrödinger once noted, a physical scientist does not introduce sensation or perception into his theories. Having removed the mind from nature, he cannot therefore expect to find it there. Taking a God's-eye view to observe an individual person from the outside leaves the scientific observer with an impoverished picture of mind. It can lead to the paradoxical conclusion that consciousness is merely a bottleneck in an information-processing cascade, a bottleneck that at any given time can contain "just a few chunks" of information. We have insisted, instead, that the observer must consider consciousness by viewing the brain from within, in terms of what makes a difference to its underlying neural processes. From that vantage point, the observer will discover a unified physical process of the utmost complexity—a process that, unlike anything we have built so far, can rapidly integrate immense amounts of information.

Taking the position of a biologically based epistemology offers a major opportunity to extend our scientific view of animal behavior and human nature. This position accepts physics and evolution as two main pillars for philosophical reflection. It then proposes that the efficacious role of con-

sciousness is to construct an informative scene ("the remembered present") that connects present reality to the past value-ridden history of each individual, conscious animal. The efficacy of consciousness for the rapid integration of information and planning yields significant evolutionary advantages. The translation of such planning into unconscious learned routines is also essential for survival, and such routines constitute a vast proportion of the basic mechanisms of behavior. Indeed, consciousness plays upon such routines to give the possibility of enhanced planning and of ever more complex learned acts.

While the efficacy of nonconscious mechanisms cannot be denied, biologically based epistemology considers consciousness to be a sine qua non of mental acts. Without entering into definitional controversies, we take the position that thought is a conscious process underlaid by a deep structure of necessary nonconscious mechanisms, including nonrepresentational memory, value constraints, and the action of cortical appendages like the basal ganglia, hippocampus, and cerebellum.

The incorporation of value systems as necessary constraints on the workings of the brain as a selectional system ties the view of biologically based epistemology to the view that emotions are fundamental both to the origins of and the appetite for conscious thought. As described by Spinoza, emotions may represent human bondage, but despite the apparent paradox, we think it likely that it was mainly emotions that impelled him to create his magnificent edifice of thought. Value systems and emotions are essential to the selectional workings of the brain that underlie consciousness. Further neuroscientific research on these systems and their modification by learning should shed light on an important issue: the place of value in a world of facts.[7]

A final word on the causal efficacy of consciousness: If anyone doubts the efficacy of consciousness, let him compare the workings of eusocial insects to the abstract constructions of poets, composers, mathematicians, and scientists. Without life, the intricate behavioral webs of wasps and the structures of termite colonies certainly are not likely to arise spontaneously. But as impressive as these colonies are, they cannot be compared to the grand view of the universe that has emerged from the workings of higher-order consciousness in human beings. We continue to describe our place in the universe by scientific means and, at the same time, give ourselves comfort and significance in that place by artistic means. In the realization of both ends, it is consciousness that provides the freedom and the warrant.

CONSCIOUSNESS AS A PHYSICAL PROCESS

We have argued throughout this book that consciousness arises from certain arrangements in the material order of the brain. There is a common prejudice that to call something material is somehow to refuse its entry into the realm of exalted things—mind, spirit, pure thought. The word *material* can be used to refer to many things or states. As it is used in these pages, it applies to what we commonly call the real world of sensible or measurable things, the world that scientists study. That world is considerably more subtle than it first appears. A chair is material (shaped by us, of course), a star is material, atoms and fundamental particles are material—they are made of matter-energy. The thought, "thinking about Vienna," however, while couched in material terms, is, as Willard Van Orman Quine pointed out, a materially based process but is, itself, not material.

What is the difference? It is that conscious thought is a set of relations with a meaning that goes beyond just energy or matter (although it involves both). And what of the mind that gave rise to that thought? The answer is, it is both material and meaningful. There is a material basis for the mind as a set of relations: The action of your brain and all its mechanisms, bottom to top, atoms to behavior, results in a mind that can be concerned with processes of meaning. While generating such immaterial relationships that are recognized by it and other minds, this mind is completely based in and dependent on the physical processes that occur in its own workings, in those of other minds, and in the events involved in communication. There are no completely separate domains of matter and mind and no grounds for dualism. But obviously, there is a realm created by the physical order of the brain, the body, and the social world in which meaning is consciously made. That meaning is essential both to our description of the world and to our scientific understanding of it. It is the amazingly complex material structures of the nervous system and body that give rise to dynamic mental processes and to meaning. Nothing else need be assumed—neither other worlds, or spirits, or remarkable forces as yet unplumbed, such as quantum gravity.[8]

There is a web to untangle here: Humans were capable of meaning and of thought before they had a scientific description of the world. Any such scientific description, even when clarified, cannot be fully tested or sustained by just one person for an indefinite period of time. It needs social interactions or, at least, two persons to make an ongoing experimental science. Yet a single person can have both private thoughts, not fully capturable by a scientific description, at the same time that he or she has a quite correct scientific

understanding. So, what happens when we turn scientific inquiry in the direction of the individual human brain and mind? What are the limits? What can we expect to capture and understand by such a scientific adventure?

Our claim is that we may capture the material bases of mind even to the extent of having a satisfactory understanding of the origins of exalted things, such as the mental. To do so, we may have to invent further ways of looking at brains and their activities. We may even have to synthesize artifacts resembling brains connected to bodily functions in order fully to understand those processes. Although the day when we shall be able to create such conscious artifacts is far off, we may have to make them—that is, use synthetic means—before we deeply understand the processes of thought itself. However far off the date of their construction, such artifacts shall be made. After all, it has been done at least once by evolution. The history of science, particularly of biological science, has shown repeatedly that apparently mysterious or impassable barriers to our understanding were based on false views or technical limitations. The material bases of mind are no exception.

This position does not contradict the conclusion that each mind is unique, not fully exhaustible by scientific means, and not a machine. Do not search for the mystical here. Our statements about the material order and immaterial meaning are not only mutually consistent within a scientific framework, but live in a useful symbiosis.

PRISONERS OF DESCRIPTION OR MASTERS OF MEANING?

Our analysis has been predicated on the notion that while we can construct a sensible scientific theory of consciousness that explains how matter becomes imagination, that theory cannot replace experience: Being is not describing. A scientific description can have predictive and explanatory power, but it cannot directly convey the phenomenal experience that depends on having an individual brain and body. In our theory of brain complexity, we have removed the paradoxes that arise by assuming only the God's-eye view of the external observer and, by adhering to selectionism, we have removed the homunculus. Nevertheless, because of the nature of embodiment, we still remain, to some extent, prisoners of description, only somewhat better off than the occupants of Plato's cave. Can we get around this limitation—this qualification of our realism? Not completely, but we return to the extravagant thought that we may transcend our analytic limits by synthetic means. Even if, some long time into the future, we can eventually construct a conscious artifact that, *mirabile dictu*, has linguistic capability, we will, even then, not directly know the actual phenomenal experience of that artifactual individual; the qualia we experience, each of us, artifact or person, rests in our own embodiment, our own phenotype.

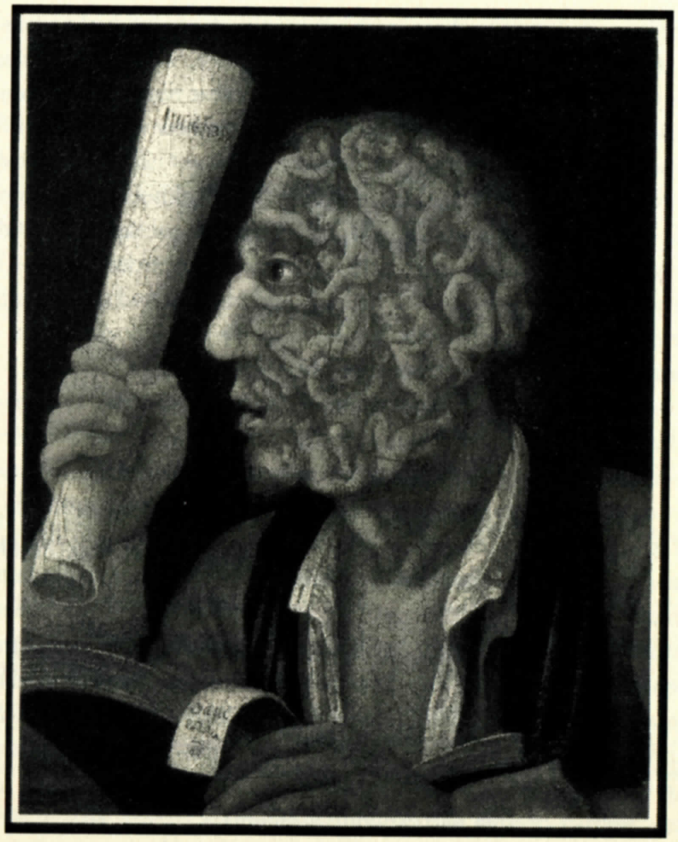

FIGURE 17.2 The world knot can be expressed in many ways, as, for example, in this figure *Counterpart* by Giuseppe Arcimboldo (1527–93) and in this fragment of a poem by Emily Dickinson (1830–86):

The Brain—is wider than the Sky—
For—put them side by side—
The one the other will contain
With ease—and You—beside.

There is no real mystery here; embodiment in the individual corpus is the price of admission for any such qualitative experience. There is, however, one new expanse of knowledge that would open up at that remarkable moment in our intellectual journey. It is the opportunity to see how a radically different phenotype capable of higher-order consciousness actually categorized the same world, the world we would share with it. The likelihood that this phenotype will be like ours or even like that of a complex animal appears diminishingly small. But could such an artifact, in its descriptions, reach a lawful generality identical to ours, even if its body and mind carved up the signals of this world in a radically different fashion than we do? If it did, some of the qualifications constraining our position of realism might be removed.

The limits and ranges of exploration of the material order as it gives rise to mind may or may not fall out according to our speculations here. But there is one fascinating point that, here and now, bears on the exhaustiveness of the scientific pursuit. It concerns whether *all* meaningful relations at the level of consciousness constitute objects for scientific study. Think, for example, of meaningful sentences in ordinary language or, even better, of poetic exchanges as they are enacted by sentient humans. Our conjecture is that they are, here and now, not fit objects for scientific study except in some trivial sense. Their meaning and description rest on too many unique historical patterns; on multiple ambiguous references; and, in the case of a unique poetic utterance (see figure 17.2), on an incomparable sample. To grasp their meaning requires both the unique phenomenal experience and the historically based culture of each participating individual.

We do not wish to be misunderstood: Just as consciousness itself may be explained, the *bases* of such objects and utterances can be fully explained by scientific inquiry as arising in the material order. But even though they may be more directly accessible as objects for study than the origin of the cosmos, they are unfit (except in the most trivial sense) to be scientific subjects, and they will not yield up their significance through scientific inquiry alone. They do yield significance, however, as a result of our individual embodiment and the mutual grammatical exchanges that allow us to experience higher-order consciousness.

If we consider that most of our lives take on meaning in the rich soup of such exchanges, we need not fear exhaustion by scientific reduction. But neither need we call upon mystical explanations to account for such richness. It is enough to recognize that some scientifically founded objects are not appropriate scientific subjects. Rejoice in it. While we remain prisoners of description, our freedom is in the grammar.