CHAPTER 6

THE MIND

BELIEF IN THE intrinsic unity of knowledge—the reality of the labyrinth—rides ultimately on the hypothesis that every mental process has a physical grounding and is consistent with the natural sciences. The mind is supremely important to the consilience program for a reason both elementary and disturbingly profound: Everything that we know and can ever know about existence is created there.

The loftier forms of such reflection and belief may seem at first to be the proper domain of philosophy, not science. But history shows that logic launched from introspection alone lacks thrust, can travel only so far, and usually heads in the wrong direction. Much of the history of modern philosophy, from Descartes and Kant forward, consists of failed models of the brain. The shortcoming is not the fault of the philosophers, who have doggedly pushed their methods to the limit, but a straightforward consequence of the biological evolution of the brain. All that has been learned empirically about evolution in general and mental process in particular suggests that the brain is a machine assembled not to understand itself, but to survive. Because these two ends are basically different, the mind unaided by factual knowledge from science sees the world only in little pieces. It throws a spotlight on those portions of the world it must know in order to live to the next day, and surrenders the rest to darkness. For thousands of generations people lived and reproduced with no need to know how the machinery of the brain works. Myth and self-deception, tribal identity and ritual, more than objective truth, gave them the adaptive edge.

That is why even today people know more about their automobiles than they do about their own minds—and why the fundamental explanation of mind is an empirical rather than a philosophical or religious quest. It requires a journey into the brain's interior darkness with preconceptions left behind. The ships that brought us here are to be left scuttled and burning at the shore.

THE BRAIN IS a helmet-shaped mass of gray and white tissue about the size of a grapefruit, one to two quarts in volume, and on average weighing three pounds (Einstein's brain, for example, was 2.75 pounds). Its surface is wrinkled like that of a cleaning sponge, and its consistency is custardlike, firm enough to keep from puddling on the floor of the brain case, soft enough to be scooped out with a spoon.

The brain's true meaning is hidden in its microscopic detail. Its fluffy mass is an intricately wired system of about a hundred billion nerve cells, each a few millionths of a meter wide and connected to other nerve cells by hundreds or thousands of endings. If we could shrink ourselves to the size of a bacterium and explore the brain's interior on foot, as philosophers since Leibniz in 1713 have imagined doing, we might eventually succeed in mapping all the nerve cells and tracking all the electrical circuits. But we could never thereby understand the whole. Far more information is needed. We need to know what the electric patterns mean, as well as how the circuits were put together and, most puzzling of all, for what purpose.

What we know of the heredity and development of the brain shows them to be almost unimaginably complicated. The human genome database accumulated to 1995 reveals that the brain's structure is prescribed by at least 3,195 distinctive genes, 50 percent more than for any other organ or tissue (the total number of genes in the entire human genome is estimated to be 50,000 to 100,000). The molecular processes that guide the growth of neurons to their assigned places have only begun to be deciphered. Overall, the human brain is the most complex object known in the universe—known, that is, to itself.

It rose by evolution to its present form swiftly, even by the standards of the generally hurried pace of mammalian phylogeny evident in the fossil record. Across three million years, from the ancestral man-apes of Africa to the earliest anatomically modern *Homo sapiens*, who lived about 200,000 years ago, the brain increased in volume four times over. Much of the growth occurred in the neocortex, the seat of the higher functions of mind, including, especially, language and its symbol-based product, culture.

The result was the capacity to take possession of the planet. Advanced humans, their big spherical skulls teetering precariously on fragile stems of compacted cervical vertebrae, walked, paddled, and sailed out of Africa through Europe and Asia and thence to all the remaining continents and great archipelagoes except uninhabitable Antarctica. By 1000 A.D. they reached the outermost islands of the Pacific and Indian Oceans. Only a handful of remote mid-Atlantic islands, including St. Helena and the Azores, remained pristine for a few centuries longer.

It is, I must acknowledge, unfashionable in academic circles nowadays to speak of evolutionary progress. *All the more reason to do so.* In fact, the dilemma that has consumed so much ink can be evaporated with a simple semantic distinction. If we mean by progress the advance toward a preset goal, such as that composed by intention in the human mind, then evolution by natural selection, which has no preset goals, is not progress. But if we mean the production through time of increasingly complex and controlling organisms and societies, in at least some lines of descent, with regression always a possibility, then evolutionary progress is an obvious reality. In this second sense, the human attainment of high intelligence and culture ranks as the last of the four great steps in the overall history of life. They followed one upon the other at roughly one-billion-year intervals. The first was the beginning of life itself, in the form of simple bacteriumlike organisms. Then came the origin of the complex eukaryotic cell through the assembly of the nucleus and other membrane-enclosed organelles into a tightly organized unit. With the eukaryotic building block available, the next advance was the origin of large, multicellular animals such as crustaceans and mollusks, whose movements were guided by sense organs and central nervous systems. Finally, to the grief of most preexisting life forms, came humanity.

VIRTUALLY ALL contemporary scientists and philosophers expert on the subject agree that the mind, which comprises consciousness and rational process, is the brain at work. They have rejected the mind-brain dualism of René Descartes, who in *Meditationes* (1642) concluded that "by the divine power the mind can exist without the body, and the body without the mind." According to the great philosopher, the noncorporeal mind and hence the immortal soul repose somewhere in the corporeal and mortal body. Its location, he suggested, might be the pineal gland, a tiny organ located at the base of the brain. In this early neurobiological model, the brain receives information from all over the body and feeds it into the pineal headquarters, where it is translated somehow into conscious thought. Dualism was congenial to the philosophy and science of Descartes' time, appealing as it did to the materialistic explanation of the universe while remaining safely pious. In one form or other, it has persisted into the late twentieth century.

The brain and its satellite glands have now been probed to the point where no particular site remains that can reasonably be supposed to harbor a nonphysical mind. The pineal gland, for example, is known to secrete the hormone melatonin and to assist in regulating the body's biological clock and daily rhythms. But even as mind-body dualism is being completely abandoned at long last, in the 1990s, scientists remain unsure about the precise material basis of mind. Some are convinced that conscious experience has unique physical and biological properties that remain to be discovered. A few among them, archly called the mysterians by their colleagues, believe that conscious experience is too alien, too complex, or both, ever to be comprehended.

No doubt, the transcendent difficulty of the subject inspires this kind of denial. As late as 1970 most scientists thought the concept of mind a topic best left to philosophers. Now the issue has been joined where it belongs, at the juncture of biology and psychology. With the aid of powerful new techniques, researchers have shifted the frame of discourse to a new way of thinking, expressed in the language of nerve cells, neurotransmitters, hormone surges, and recurrent neural networks.

The cutting edge of the endeavor is cognitive neuroscience, also and more popularly known as the brain sciences, an alliance formed by neurobiologists, cognitive psychologists, and a new school of empirically minded philosophers sometimes referred to as neurophilosophers. Their research reports are dispatched weekly to premier scientific journals, and their theories and impassioned disagreements fill the pages of such open-commentary periodicals as *Behavioral and Brain Sciences*. Many of the popular books and articles they write rank among the best in contemporary science exposition.

Such traits are the hallmark of the heroic period, or romantic period as it is often called, experienced by every successful scientific discipline during its youth. For a relatively brief interval, usually a decade or two, rarely more than half a century, researchers are intoxicated with a mix of the newly discovered and the imaginable unknown. For the first time the really important questions are asked in a form that can be answered, thus: *What are the cellular events that compose the mind?* Not create the mind—too vague, that expression—but compose the mind. The pioneers are paradigm hunters. They are risk takers, who compete with rival theorists for big stakes and are willing to endure painful shake-outs. They bear comparison with explorers of the sixteenth century, who, having discovered a new coastline, worked rivers up to the fall line, drew crude maps, and commuted home to beg for more expeditionary funds. And governmental and private patrons of the brain scientists, like royal geographic commissions of past centuries, are generous. They know that history can be made by a single sighting of coastland, where inland lies virgin land and the future lineaments of empire.

Call the impulse Western if you wish, call it androcentric, and by all means dismiss it as colonialist if you feel you must. I think it instead basic to human nature. Whatever its source, the impulse drives major scientific advance. During my career I have been privileged to witness close at hand the heroic periods of molecular biology, plate tectonics in geology, and the modern synthesis of evolutionary biology. Now it is the turn of the brain sciences.

THE EARLY GROUNDWORK for the revolution was laid in the nineteenth century by physicians, who noticed that injuries to certain parts of the brain result in special kinds of disability. Perhaps the most famous case was that of Phineas P. Gage, who in 1848 was a young construction foreman in charge of a crew laying railroad track across Vermont. Part of the job was to blast away outcrops of hard rock in order to straighten out turns in the advancing path. As Gage pressed powder into a newly drilled hole, a premature explosion fired the iron tamping bar like a missile toward his head. It entered his left cheek and exited the top of his skull, carrying with it a good part of the prefrontal lobe of his cerebral cortex, then arced away more than a hundred feet before coming to earth. Gage fell to the ground, miraculously still alive. To the amazement of all, he was able within minutes to sit up and even walk with assistance. He never lost consciousness. "Wonderful accident" was the later headline in the *Vermont Mercury*. In time his external injuries healed, and he retained the ability to speak and reason. But his personality had changed drastically. Where previously he had been cheerful, responsible, and well-mannered, a valued employee of the Rutland & Burlington Railroad, now he was a habitual liar, unreliable at work, and given to vagrant, self-destructive behavior. Studies on other patients with injuries to the same part of the brain over many years have confirmed the general conclusion suggested by Gage's misfortune: The prefrontal lobe houses centers important for initiative and emotional balance.

For two centuries the medical archives have filled with such anecdotes on the effects of localized brain damage. The data have made it possible for neurologists to piece together a map of functions performed by different parts of the brain. The injuries, which occur throughout the brain, include physical traumas, strokes, tumors, infections, and poisoning. They vary in extent from barely detectable pinpoints to deletions and transections of large parts of the brain. Depending on location and magnitude, they have multifarious effects on thought and behavior.

The most celebrated recent case is that of Karen Ann Quinlan. On April 14, 1975, the young New Jersey woman, while dosed with the tranquilizer Valium and painkiller drug Darvon, made the mistake of drinking gin and tonic. Although the combination does not sound dangerous, it essentially killed Karen Ann Quinlan. She fell into a coma that lasted until her death from massive infections ten years later. An autopsy revealed that her brain was largely intact, which explains why her body survived and even continued its daily cycle of waking and sleep. It lived on even when Quinlan's parents arranged, in the midst of national controversy, to have her ventilator removed. The autopsy revealed that Quinlan's brain damage was local but very severe: The thalamus had been obliterated as though burned out with a laser. Why that particular center deteriorated is unknown. A brain injured by a heavy blow or certain forms of poisoning usually responds by widespread swelling. If the reaction is intense, it presses on centers that control heartbeat and respiration, shutting down blood circulation and soon ending in death of the whole body.

The result of thalamus excision alone is brain death, or, more precisely put, mind death. The thalamus comprises twin egg-shaped masses of nerve cells near the center of the brain. It functions as a relay center through which all sensory information other than smell is transmitted to the cerebral cortex, and therefore to the conscious mind. Even dreams are triggered by impulses that pass through thalamic circuits. Quinlan's drug accident was the equivalent of blowing up a power station: All her lights downline went out, and she entered a sleep from which she had no chance of wakening. Her cerebral cortex lived on, waiting to be activated. But consciousness, even in dreams, was no longer possible.

Such research on brain damage, while enormously informative, is nevertheless dependent on chance occurrence. Over the years it has been greatly enhanced by experimental brain surgery. Neurosurgeons routinely keep patients conscious to test their response to electrical stimulation of the cortex, in order to locate healthy tissue and avoid excising it. The procedure is not uncomfortable: Brain tissue, while processing impulses from all over the body, has no receptors of its own. Instead of pain, the roving probes evoke a medley of sensations and muscular contractions. When certain sites on the surface of the cortex are stimulated, patients experience images, melodies, incoherent sounds, and a gamut of other impressions. Sometimes they involuntarily move fingers and other body parts.

Beginning with experiments in brain surgery by Wilder Penfield and other pioneers in the 1920s and 1930s, researchers have mapped sensory and motor functions over all parts of the cerebral cortex. The method is nevertheless limited in two important respects. It is not easily extended beneath the cortex into the dark nether regions of the brain, and it cannot be used to observe neural activity through time. To reach those objectives—to create motion pictures of the whole brain in action—scientists have adopted a broad range of sophisticated techniques borrowed from physics and chemistry. Since its inception in the 1970s, brain imaging, as the methods are collectively called, has followed a trajectory similar to that of microscopy, toward ever finer resolution in snapshots separated by shrinking intervals of time. The scientists hope eventually to monitor the activity of entire networks of individual nerve cells, both continuously and throughout the living brain.

GRANTED, the brain's machinery remains forbiddingly alien and scientists have traced only a minute fraction of its circuitry. Still, the major anatomical features of the brain are known, and a great deal has been learned of their various functions. Before addressing the nature of mind as a product of these operations, I wish to provide a quick look at the physical groundwork.

The surest way to grasp complexity in the brain, as in any other biological system, is to think of it as an engineering problem. What are the broad principles needed to create a brain from scratch? Whether contrived by advance planning or by blind natural selection, the key features of architecture can be expected to be very broadly predictable. Researchers in biomechanics have discovered time and again that organic structures evolved by natural selection conform to high levels of efficiency when judged by engineering criteria. And at a more microscopic level, biochemists marvel at the exactitude and power of the enzyme molecules controlling the actions of the cells. Like the mills of God, the processes of evolution grind slowly—yet, as the poet said, they grind exceeding fine.

So let us spread the specification sheets out and consider the brain as a solution to a set of physical problems. It is best to start with simple geometry. Because a huge amount of circuitry is required, and the wiring elements must be built from living cells, a relatively huge mass of new tissue needs to be manufactured and housed in the brain case. The ideal brain case will be spherical or close to it. One compelling reason is that a sphere has the smallest surface relative to volume of any geometric form and hence provides the least access to its vulnerable interior. Another reason is that a sphere allows more circuits to be placed close together. The average length of circuits can thus be minimized, raising the speed of transmission while lowering the energy cost for their construction and maintenance.

Because the basic units of the brain-machine must be made of cells, it is best to stretch these elements out into string-shaped forms that serve simultaneously as receiving stations and coaxial cables. The dual-purpose cells created by evolution are in fact the neurons, also called nerve cells or nerve fibers. It is further practical to design the neurons so that their main bodies serve as the receiving sites for impulses from other cells. The neurons can send their own signals out along axons, cablelike extensions of the cell bodies.

For speed, make the transmission an electric discharge by depolarization of the cell membrane. The neurons are then said to "fire." For accuracy during neuron firing, surround the axons with insulating sheaths. These in fact exist as white fatty myelin membranes that together give the brain its light color.

To achieve a higher level of integration, the brain must be very intricately and precisely wired. Given again that its elements are living cells, the number of neuron connections are best multiplied by growing threadlike extensions from the tips of the axons, which reach out and transmit individually to the bodies of many other cells. The discharge of the axon travels to these multiple terminal extensions all the way to their tips, which then make contact with the receptor cells. The receptor cells accept some of the tips of the terminal axon branches on the surface of their main cell bodies. They accept other tips on their dendrites, which are threadlike receptor branches growing out from the cell bodies.

Now visualize the entire nerve cell as a miniature squid. From its body sprouts a cluster of tentacles (the dendrites). One tentacle (the axon) is much longer than the others, and from its tip it sprouts more tentacles. The message is received on the body and short tentacles of the squid and travels along the long tentacle to other squids. The brain comprises the equivalent of one hundred billion squids linked together.

The cell-to-cell connections—more precisely, the points of connection and the ultramicroscopic spaces separating them—are called synapses. When an electric discharge reaches a synapse, it induces the tip of the terminal branch to release a neurotransmitter, a chemical that either excites an electric discharge in the receiving cell or prevents one from occurring. Each nerve cell sends signals to hundreds or thousands of other cells through its synapses at the end of its axon, and it receives input from a similar myriad of synapses on its main cell body and dendrites. In each instant a nerve cell either fires an impulse along its axon to other cells or falls silent. Which of the two responses it makes depends on the summation of the neurotransmissions received from all the cells that feed stimuli into it.

The activity of the brain as a whole, hence the wakefulness and moods experienced by the conscious mind, is profoundly affected by the levels of the neurotransmitters that wash its trillions of synapses. Among the most important of the neurotransmitters are acetylcholine and the amines norepinephrine, serotonin, and dopamine. Others include the amino acid GABA (gamma aminobutyric acid) and, surprisingly, the elementary gas nitrous oxide. Some neurotransmitters excite the neurons they contact, while others inhibit them. Still others can exert either effect depending on the location of the circuit within the nervous system.

During development of the nervous system in the fetus and infant, the neurons extend their axons and dendrites into the cellular environment—like growing tentacles of squids. The connections they make are precisely programmed and guided to their destinations by chemical cues. Once in place each neuron is poised to play a special role in signal transmission. Its axon may stretch only a few millionths of a meter or thousands of times longer. Its dendrites and terminal axon branches can take any of a number of forms, coming to resemble, say, the leafless crown of a tree in winter or a dense, feltlike mat. Possessing the aesthetic inherent to pure function, and riveting to behold, they invite us to imagine their powers. Concerning them, Santiago Ramón y Cajal, the great Spanish histologist, wrote of his own experience, after receiving the 1906 Nobel Prize for his research on the subject: "Like the entomologist in pursuit of brightly colored butterflies, my attention hunted, in the flower garden of the gray matter, cells with delicate and elegant forms, the mysterious butterflies of the soul, the beatings of whose wings may some day—who knows?—clarify the secret of mental life."

The meaning of the neuron shape, which so pleases the biologist, is this: Neuron systems are directed networks, receiving and broadcasting signals. They cross-talk with other complexes to form systems of systems, in places forming a circle, like a snake catching its own tail, to create reverberating circuits. Each neuron is touched by the terminal axon branches of many other neurons, established by a kind of democratic vote whether it is to be active or silent. Using a Morselike code of staccato firing, the cell sends its own messages outward to others. The number of connections made by the cell, their pattern of spread, and the code they use determine the role the cell plays in the overall activity of the brain.

Now to complete the engineering metaphor. When you're setting out to design a hominid brain, it is important to observe another optimum design principle: Information transfer is improved when neuron circuits filling specialized functions are placed together in clusters. Examples of such aggregates in the real brain are the sensory relay stations, integrative centers, memory modules, and emotional control centers identified thus far by neurobiologists. Nerve cell bodies are gathered in flat assemblages called layers and rounded ones called nuclei. Most are placed at or near the surface of the brain. They are interconnected both by their own axons and by intervening neurons that course through the deeper brain tissues. One result is the gray or light-brown color of the surface due to the massing of the cell bodies—the "gray matter" of the brain—and a white color from the myelin sheaths of axons in the interior of the brain.

Human beings may possess the most voluminous brain in proportion to body size of any large animal species that has ever lived. For a primate species the human brain is evidently at or close to its physical limit. If it were much larger in the newborn, the passage of its protecting skull through the birth canal would be dangerous to both mother and child. Even the adult brain size is mechanically risky: The head is a fragile, internally liquescent globe balanced on a delicate bone-and-muscle stem, within which the brain is vulnerable and the mind easily stunned and disabled. Human beings are innately disposed to avoid violent physical contact. Because our evolving ancestors traded brute strength for intelligence, we no longer need to seize and rip enemies with fanged jaws.

Given this intrinsic limit in brain volume, some way must be found to fit in the memory banks and higher-order integrating systems needed to generate conscious thought. The only means available is to increase surface area: Spread the cells out into a broad sheet and crumple it up into a ball. The human cerebral cortex is such a sheet about one thousand square inches in area, packed with millions of cell bodies per square inch, folded and wadded precisely like an origami into many winding ridges and fissures, neatly stuffed in turn into the quart-sized cranial cavity.

WHAT MORE CAN be said of brain structure? If a Divine Engineer designed it, unconstrained by humanity's biological history, He might have chosen mortal but angelic beings cast in His own image. They would presumably be rational, far-seeing, wise, benevolent, unrebellious, selfless, and guilt-free, and, as such, ready-made stewards of the beautiful planet bequeathed them. But we are nothing like that. We have original sin, which makes us *better* than angels. Whatever good we possess we have earned, during a long and arduous evolutionary history. The human brain bears the stamp of 400 million years of trial and error, traceable by fossils and molecular homology in nearly unbroken sequence from fish to amphibian to reptile to primitive mammal to our immediate primate forerunners. In the final step the brain was catapulted to a radically new level, equipped for language and culture. Because of its ancient pedigree, however, it could not be planted like a new computer into an empty cranial space. The old brain had been assembled there as a vehicle of instinct, and remained vital from one heartbeat to the next as new parts were added. The new brain had to be jury-rigged in steps within and around the old brain. Otherwise the organism could not have survived generation by generation. The result was human nature: genius animated with animal craftiness and emotion, combining the passion of politics and art with rationality, to create a new instrument of survival.

Brain scientists have vindicated the evolutionary view of mind. They have established that passion is inseverably linked to reason. Emotion is not just a perturbation of reason but a vital part of it. This chimeric quality of the mind is what makes it so elusive. The hardest task of brain scientists is to explain the products-tested engineering of the cortical circuits against the background of the species' deep history. Beyond the elements of gross anatomy I have just summarized, the hypothetical role of Divine Engineer is not open to them. Unable to deduce from first principles the optimum balance of instinct and reason, they must ferret out the location and function of the brain's governing circuits one by one. Progress is measured by piecemeal discoveries and cautious inferences. Here are a few of the most important made by researchers to date:

- The human brain preserves the three primitive divisions found throughout the vertebrates from fishes to mammals: hindbrain, midbrain, and forebrain. The first two together, referred to as the brain stem, form the swollen posthead on which the massively enlarged forebrain rests.

• The hindbrain comprises in turn the pons, medulla, and cerebellum. Together they regulate breathing, heartbeat, and coordination of body movements. The midbrain controls sleep and arousal. It also partly regulates auditory reflexes and perception.

• A major part of the forebrain is composed of the limbic system, the master traffic-control complex that regulates emotional response as well as the integration and transfer of sensory information. Its key centers are the amygdala (emotion), hippocampus (memory, especially short-term memory), hypothalamus (memory, temperature control, sexual drive, hunger, and thirst), and thalamus (awareness of temperature and all other senses except smell, awareness of pain, and the mediation of some processes of memory).

• The forebrain also includes the cerebral cortex, which has grown and expanded during evolution to cover the rest of the brain. As the primary seat of consciousness, it stores and collates information from the senses. It also directs voluntary motor activity and integrates higher functions, including speech and motivation.

• The key functions of the three successive divisions—hind- plus midbrain, limbic system, and cerebral cortex—can be neatly summarized in this sequence: *heartbeat, heartstrings, heartless*.

• No single part of the forebrain is the site of conscious experience. Higher levels of mental activity sweep through circuits that embrace a large part of the forebrain. When we see and speak of color, for example, visual information passes from the cones and interneurons of the retina through the thalamus to the visual cortex at the rear of the brain. After the information is codified and integrated anew at each step, through patterns of neuron firing, it then spreads forward to the speech centers of the lateral cortex. As a result, we first see red and then say "red." Thinking about the phenomenon consists of adding more and more connections of pattern and meaning, and thus activating additional areas of the brain. The more novel and complicated the connections, the greater the amount of this spreading activation. The better the connections are learned by such experience, the more they are put on autopilot. When the same stimulus is applied later, new activation is diminished and the circuits are more predictable. The procedure becomes a "habit." In one such inferred pathway of memory formation, sensory information is conveyed from the cerebral cortex to the amygdala and hippocampus, then to the thalamus, then to the prefrontal cortex (just behind the brow), and back to the original sensory regions of the cortex for storage. Along the way codes are interpreted and altered according to inputs from other parts of the brain.

• Because of the microscopic size of the nerve cells, a large amount of circuitry can be packed into a very small space. The hypothalamus, a major relay and control center at the base of the brain, is about the size of a lima bean. (The nervous systems of animals are even more impressively miniaturized. The entire brains of gnats and other extremely small insects, which carry instructions for a series of complex instinctive acts, from flight to mating, are barely visible to the naked eye.)

• Disturbance of particular circuits of the human brain often produce bizarre results. Injuries to certain sites of the undersurface of the parietal and occipital lobes, which occupy the side and rear of the cerebral cortex, cause the rare condition called prosopagnosia. The patient can no longer recognize other persons by their faces, but he can still remember them by their voices. Just as oddly, he retains the ability to recognize objects other than faces by sight alone.

• There may be centers in the brain that are especially active in the organization and perception of free will. One appears to be located within or at least close to the anterior cingulate sulcus, on the inside of a fold of the cerebral cortex. Patients who have sustained damage to the region lose initiative and concern for their own welfare. From one moment to the next they focus on nothing in particular, yet remain capable of reasoned responses when pressed.

• Other complex mental operations, while engaging regions over large parts of the brain, are vulnerable to localized perturbation. Patients with temporal lobe epilepsy often develop hyperreligiosity, the tendency to charge all events, large and small, with cosmic significance. They are also prone to hypergraphia, a compulsion to express their visions in an undisciplined stream of poems, letters, or stories.

• The neural pathways used in sensory integration are also highly specialized. When subjects name pictures of animals during PET (positron emission tomography) imaging, a method that reveals patterns of nerve-cell firing, their visual cortices light up in the same pattern seen when they sort out subtle differences in the appearance of objects. When, on the other hand, they silently name pictures of tools, neural activity shifts to parts of the cortex concerned with hand movements and action words, such as "write" for pencil.

I HAVE SPOKEN so far about the physical processes that produce the mind. Now, to come to the heart of the matter, what *is* the mind? Brain scientists understandably dance around this question. Wisely, they rarely commit themselves to a simple declarative definition. Most believe that the fundamental properties of the elements responsible for mind—neurons, neurotransmitters, and hormones—are reasonably well known. What is lacking is a sufficient grasp of the emergent, holistic properties of the neuron circuits, and of cognition, the way the circuits process information to create perception and knowledge. Although dispatches from the research front grow yearly in number and sophistication, it is hard to judge how much we know in comparison with what we need to know in order to create a powerful and enduring theory of mind production by the brain. The grand synthesis could come quickly, or it could come with painful slowness over a period of decades.

Still, the experts cannot resist speculation on the essential nature of mind. While it is very risky to speak of consensus, and while I have no great trust in my own biases as interpreter, I believe I have been able to piece together enough of their overlapping opinions to forecast a probable outline of the eventual theory, as follows.

Mind is a stream of conscious and subconscious experience. It is at root the coded representation of sensory impressions and the memory and imagination of sensory impressions. The information composing it is most likely sorted and retrieved by vector coding, which denotes direction and magnitude. For example, a particular taste might be partly classified by the combined activity of nerve cells responding to different degrees of sweetness, saltiness, and sourness. If the brain were designed to distinguish ten increments in each of these taste dimensions, the coding could discriminate 10 × 10 × 10, or 1,000 substances.

Consciousness consists of the parallel processing of vast numbers of such coding networks. Many are linked by the synchronized firing of the nerve cells at forty cycles per second, allowing the simultaneous internal mapping of multiple sensory impressions. Some of the impressions are real, fed by ongoing stimulation from outside the nervous system, while others are recalled from the memory banks of the cortex.

All together they create scenarios that flow realistically back and forth through time. The scenarios are a virtual reality. They can either closely match pieces of the external world or depart indefinitely far from it. They re-create the past and cast up alternative futures that serve as choices for future thought and bodily action. The scenarios comprise dense and finely differentiated patterns in the brain circuits. When fully open to input from the outside, they correspond well to all the parts of the environment, including activity of the body parts, monitored by the sense organs.

Who or what within the brain monitors all this activity? No one. Nothing. The scenarios are not seen by some other part of the brain. They just *are*. Consciousness is the virtual world composed by the scenarios. There is not even a Cartesian theater, to use Daniel Dennett's dismissive phrase, no single locus of the brain where the scenarios are played out in coherent form. Instead, there are interlacing patterns of neural activity within and among particular sites throughout the forebrain, from cerebral cortex to other specialized centers of cognition such as the thalamus, amygdala, and hippocampus. There is no single stream of consciousness in which all information is brought together by an executive ego. There are instead multiple streams of activity, some of which contribute momentarily to conscious thought and then phase out. Consciousness is the massive coupled aggregates of such participating circuits. The mind is a self-organizing republic of scenarios that individually germinate, grow, evolve, disappear, and occasionally linger to spawn additional thought and physical activity.

The neural circuits do not turn on and off like parts of an electrical grid. In many sectors of the forebrain at least, they are arranged in parallel relays stepping from one neuron level to the next, integrating more and more coded information with each step. The energy of light striking the retina, to expand the example I gave earlier, is transduced into patterns of neuron firing. The patterns are relayed through a sequence of intermediate neuron systems out of the retinal fields through the lateral geniculate nuclei of the thalamus back to the primary visual cortex at the rear of the brain. Cells in the visual cortex fed by the integrated stimuli sum up the information from different parts of the retina. They recognize and by their own pattern of firing specify spots or lines. Further systems of these higher-order cells integrate the information from multiple feeder cells to map the shape and movement of objects. In ways still not understood, this pattern is coupled with simultaneous input from other parts of the brain to create the full scenarios of consciousness. The biologist S. J. Singer has drily expressed the matter thus: I link, therefore I am.

Because just to generate consciousness requires an astronomically large population of cells, the brain is sharply limited in its capacity to create and hold complex moving imagery. A key measure of that capacity lies in the distinction made by psychologists between short-term and long-term memory. Short-term memory is the ready state of the conscious mind. It composes all of the current and remembered parts of the virtual scenarios. It can handle only about seven words or other symbols simultaneously. The brain takes about one second to scan these symbols fully, and it forgets most of the information within thirty seconds. Long-term memory takes much longer to acquire, but it has an almost unlimited capacity, and a large fraction of it is retained for life. By spreading activation, the conscious mind summons information from the store of long-term memory and holds it for a brief interval in short-term memory. During this time it processes the information, at a rate of about one symbol per 25 milliseconds, while scenarios arising from the information compete for dominance.

Long-term memory recalls specific events by drawing particular persons, objects, and actions into the conscious mind through a time sequence. For example, it easily re-creates an Olympic moment: the lighting of the torch, a running athlete, the cheering of the crowd. It also re-creates not just moving images and sound but *meaning* in the form of linked concepts simultaneously experienced. Fire is connected to hot, red, dangerous, cooked, the passion of sex, and the creative act, and on out through multitudinous hypertext pathways selected by context, sometimes building new associations in memory for future recall. The concepts are the nodes or reference points in long-term memory. Many are labeled by words in ordinary language, but others are not. Recall of images from the long-term banks with little or no linkage is just memory. Recall with linkages, and especially when tinged by the resonance of emotional circuits, is remembrance.

The capacity for remembrance by the manipulation of symbols is a transcendent achievement for an organic machine. It has authored all of culture. But it still falls far short of the demands placed by the body on the nervous system. Hundreds of organs must be regulated continuously and precisely; any serious perturbation is followed by illness or death. A heart forgetful for ten seconds can drop you like a stone. The proper functioning of the organs is under the control of hard-wired autopilots in the brain and spinal cord, whose neuron circuits are our inheritance from hundreds of millions of years of vertebrate evolution prior to the origin of human consciousness. The autopilot circuits are shorter and simpler than those of the higher cerebral centers and only marginally communicate with them. Only by intense meditative training can they occasionally be brought under conscious control.

Under automatic control, and specifically through balance of the antagonistic elements of the autonomic nervous system, pupils of the eye constrict or dilate, saliva pours out or is contained, the stomach churns or quietens, the heart pounds or calms, and so on through alternative states in all the organs. The sympathetic nerves of the autonomic nervous system pump the body up for action. They arise from the middle sections of the spinal cord, and typically regulate target organs by release of the neurotransmitter norepinephrine. The parasympathetic nerves relax the body as a whole while intensifying the processes of digestion. They rise from the brain stem and lowermost segment of the spinal cord, and the neurotransmitter they release to the target organs is acetylcholine—also the agent of sleep.

Reflexes are swift automatic responses mediated by short circuits of neurons through the spinal cord and lower brain. The most complex is the startle response, which prepares the body for an imminent blow or collision. Imagine that you are surprised by a loud noise close by—a car horn blasts, someone shouts, a dog charges in a fury of barking. You react without thinking. Your eyes close, your head sags, your mouth opens, your knees buckle slightly. All are reactions that prepare you for the violent contact that might follow an instant later. The startle response occurs in a split second, faster than the conscious mind can follow, faster than can be imitated by conscious effort even with long practice.

Automatic responses, true to their primal role, are relatively impervious to the conscious will. This principle of archaism extends even to the facial expressions that communicate emotion. A spontaneous and genuine smile, which originates in the limbic system and is emotion-driven, is unmistakable to the practiced observer. A contrived smile is constructed from the conscious processes of the cerebrum and is betrayed by telltale nuances: a slightly different configuration of facial muscle contraction and a tendency toward lopsidedness of the upward curving mouth. A natural smile can be closely imitated by an experienced actor. It can also be evoked by artificially inducing the appropriate emotion—the basic technique of method acting. In ordinary usage it is modified deliberately in accordance with local culture, to convey irony (the pursed smile), restrained politeness (the thin smile), threat (the wolfish smile), and other refined presentations of self.

Much of the input to the brain does not come from the outside world but from internal body sensors that monitor the state of respiration, heartbeat, digestion, and other physiological activities. The flood of "gut feeling" that results is blended with rational thought, feeding it, and being fed by it through reflexes of internal organs and neurohormonal loops.

As the scenarios of consciousness fly by, driven by stimuli and drawing upon memories of prior scenarios, they are weighted and modified by emotion. What is emotion? It is the modification of neural activity that animates and focuses mental activity. It is created by physiological activity that selects certain streams of information over others, shifting the body and mind to higher or lower degrees of activity, agitating the circuits that create scenarios, and selecting ones that end in certain ways. The winning scenarios are those that match goals preprogrammed by instinct and the satisfactions of prior experience. Current experience and memory continually perturb the states of mind and body. By thought and action the states are then moved backward to the original condition or forward to conditions conceived in new scenarios. The dynamism of the process provokes labeling by words that denote the basic categories of emotion—anger, disgust, fear, pleasure, surprise. It breaks the categories into many degrees and joins them to create myriad subtle compounds. Thus we experience feelings that are variously weak, strong, mixed, and new.

Without the stimulus and guidance of emotion, rational thought slows and disintegrates. The rational mind does not float above the irrational; it cannot free itself to engage in pure reason. There are pure theorems in mathematics but no pure thoughts that discover them. In the brain-in-the-vat fantasy of neurobiological theory and science fiction, the organ in its nutrient bath has been detached from the impediments of the body and liberated to explore the inner universe of the mind. But that is not what would ensue in reality. All the evidence from the brain sciences points in the opposite direction, to a waiting coffin-bound hell of the wakened dead, where the remembered and imagined world decays until chaos mercifully grants oblivion.

Consciousness satisfies emotion by the physical actions it selects in the midst of turbulent sensation. It is the specialized part of the mind that creates and sorts scenarios, the means by which the future is guessed and courses of action chosen. Consciousness is not a remote command center but part of the system, intimately wired to all the neural and hormonal circuits regulating physiology. Consciousness acts and reacts to achieve a dynamic steady state. It perturbs the body in precise ways with each changing circumstance, as required for well-being and response to opportunity, and helps return it to the original condition when challenge and opportunity have been met.

The reciprocity of mind and body can be visualized in the following scenario, which I have adapted from an account by the neurologist Antonio R. Damasio. Imagine that you are strolling along a deserted city street at night. Your reverie is interrupted by quick footsteps drawing close behind. Your brain focuses instantly and churns out alternative scenarios—ignore, freeze, turn and confront, or escape. The last scenario prevails and you act. You run toward a lighted storefront further down the street. In the space of a few seconds, the conscious response triggers automatic changes in your physiology. The catecholamine hormones epinephrine ("adrenaline") and norepinephrine pour into the bloodstream from the adrenal medulla and travel to all parts of the body, increasing the basal metabolic rate, breaking down glycogen in the liver and skeletal muscles to glucose for a quick energy feed. The heart races. The bronchioles of the lungs dilate to admit more air. Digestion slows. The bladder and colon prepare to void their contents, disencumbering the body to prepare for violent action and possible injury.

A few seconds more pass. Time slows in the crisis: The event span seems like minutes. Signals arising from all the changes are relayed back to the brain by more nerve fibers and the rise of hormone titers in the bloodstream. As further seconds tick away, the body and brain shift together in precisely programmed ways. Emotional circuits of the limbic system kick in—the new scenarios flooding the mind are charged with fright, then anger that sharply focuses the attention of the cerebral cortex, closing out all other thought not relevant to immediate survival.

The storefront is reached, the race won. People are inside, the pursuer is gone. Was the follower really in pursuit? No matter. The re-

1° + 2° emotions

public of bodily systems, informed by reassuring signals from the conscious brain, begins its slow stand-down to the original calm state.

Damasio, in depicting the mind holistically in such episodes, has suggested the existence of two broad categories of emotion. The first, primary emotion, comprises the responses ordinarily called inborn or instinctive. Primary emotion requires little conscious activity beyond the recognition of certain elementary stimuli, the kind that students of instinctive behavior in animals call releasers—they are said to "release" the preprogrammed behavior. For human beings such stimuli include sexual enticement, loud noises, the sudden appearance of large shapes, the writhing movements of snakes or serpentine objects, and the particular configurations of pain associated with heart attacks or broken bones. The primary emotions have been passed down with little change from the vertebrate forebears of the human line. They are activated by circuits of the limbic system, among which the amygdala appears to be the master integrating and relay center.

Secondary emotions arise from personalized events of life. To meet an old friend, fall in love, win a promotion, or suffer an insult is to fire the limbic circuits of primary emotion, but only after the highest integrative processes of the cerebral cortex have been engaged. We must know who is friend or enemy, and why they are behaving a certain way. By this interpretation, the emperor's rage and poet's rapture are cultural elaborations retrofitted to the same machinery that drives the prehuman primates. Nature, Damasio observes, "with its tinkerish knack for economy, did not select independent mechanisms for expressing primary and secondary emotions. It simply allowed secondary emotions to be expressed by the same channel already prepared to convey primary emotions."

Ordinary words used to denote emotion and other processes of mental activity make only a crude fit to the models used by the brain scientists in their attempts at rigorous explanation. But the ordinary and conventional conceptions—what some philosophers call folk psychology—are necessary if we are to make better sense of thousands of years of literate history, and thereby join the cultures of the past with those of the future. To that end I offer the following neuroscience-accented definitions of several of the most important concepts of mental activity.

What we call *meaning* is the linkage among the neural networks created by the spreading excitation that enlarges imagery and engages emotion. The competitive selection among scenarios is what we call *decision making.* The outcome, in terms of the match of the winning scenario to instinctive or learned favorable states, sets the kind and intensity of subsequent emotion. The persistent form and intensity of emotions is called *mood.* The ability of the brain to generate novel scenarios and settle on the most effective among them is called *creativity.* The persistent production of scenarios lacking reality and survival value is called *insanity.*

The explicit material constructions I have put upon mental life will be disputed by some brain scientists, and reckoned inadequate by others. That is the unavoidable fate of synthesis. In choosing certain hypotheses over others, I have tried to serve as an honest broker searching for the gravitational center of opinion, where by and large the supporting data are most persuasive and mutually consistent. To include all models and hypotheses deserving respect in this tumultuous discipline, and then to clarify the distinctions among them, would require a full-dress textbook. Undoubtedly events will prove that in places I chose badly. For that eventuality I apologize now to the slighted scientists, a concession I comfortably make, knowing that the recognition they deserve and will inevitably receive cannot be blunted by premature omission on the part of any one observer.

THE SUBJECT thus qualified, I will next describe the deeper problems that must be resolved before the physical basis of mind can be said to be truly solved. The one universally judged to be the most difficult of all is the nature of subjective experience. The Australian philosopher David Chalmers recently put the matter in perspective by contrasting the "easy problems" of general consciousness with the "hard problem" of subjective experience. In the first group (easy, I suppose, in the sense that Mont Blanc is more readily climbed in beachwear than Everest) are the classical problems of mind research: how the brain responds to sensory stimuli, how it incorporates information into patterns, and how it converts the patterns into words. Each of these steps of cognition is the subject of vigorous contemporary research.

The hard problem is more elusive: how physical processes in the brain addressed in the easy problems give rise to subjective feeling.

What exactly does it mean when we say we *experience* a color such as red or blue? Or experience, in Chalmers' words, "the ineffable sound of a distant oboe, the agony of an intense pain, the sparkle of happiness or the meditative quality of a moment lost in thought. All are part of what I am calling consciousness. It is these phenomena that compose the real mystery of the mind."

An imaginary experiment proposed by the philosopher Frank Jackson in 1983 illustrates the supposed unattainability of subjective thought by the natural sciences. Consider a neurobiologist two centuries hence who understands all the physics of color and all the brain's circuitry giving rise to color vision. But the scientist (call her Mary) has never experienced color; she has been cloistered all her life in a black-and-white room. She does not know what it is like for another person to see red or blue; she cannot imagine how they feel about color. According to Jackson and Chalmers, it follows that there are qualities of conscious experience that cannot be deduced from knowledge of the physical functioning of the brain.

Although it is the nature of philosophers to imagine impasses and expatiate upon them at book length with schoolmasterish dedication, the hard problem is conceptually easy to solve. What material description might explain subjective experience? The answer must begin by conceding that Mary cannot know what it feels like to see color. The chromatic nuances of a westering sun are not hers to enjoy. And for the same reason she and all her fellow human beings *a fortiori* cannot know how a honeybee feels when it senses magnetism or what an electric fish thinks as it orients by an electric field. We can translate the energies of magnetism and electricity into sight and sound, the sensory modalities we biologically possess. We can read the active neural circuits of bees and fish by scanning their sense organs and brains. But we cannot feel as they do—ever. Even the most imaginative and expert observers cannot think as animals, however they may wish or deceive themselves otherwise.

But incapacity is not the point. The distinction that illuminates subjective experience lies elsewhere, in the respective roles of science and art. Science perceives who can feel blue and other sensations and who cannot feel them, and explains why that difference exists. Art in contrast transmits feelings among persons of the same capacity. In other words, science explains feeling, while art transmits it. The majority of human beings, unlike Mary, see a full color spectrum, and art: means to transmit feeling

they feel its productions in reverberating pathways through the forebrain. The basic patterns are demonstrably similar across all color-sighted human beings. Variations exist, owing to remembrances that arise from the personal memories and cultural biases of different people. But in theory these variations can also be read in the patterns of their brain activity. The physical explanations derived from the patterns would be understandable to Mary the confined scientist. She might say, "Yes, that is the wavelength span classified by others as blue, and there is the pattern of neural activity by which it is recognized and named." The explanations would be equally clear to bee and fish scientists if their species could somehow be raised to human levels of intelligence.

Art is the means by which people of similar cognition reach out to others in order to transmit feeling. But how can we know for sure that art communicates this way with accuracy, that people really, truly *feel* the same in the presence of art? We know it intuitively by the sheer weight of our cumulative responses through the many media of art. We know it by detailed verbal descriptions of emotion, by critical analyses, and in fact through data from all the vast, nuanced, and interlocking armamentaria of the humanities. That vital role in the sharing of culture is what the humanities are all about. Nevertheless, fundamental new information will come from science by studying the dynamic patterns of the sensory and brain systems during episodes when commonly shared feelings are evoked and experienced through art.

But surely, skeptics will say, that is impossible. Scientific fact and art can never be translated one into the other. Such a response is indeed the conventional wisdom. But I believe it is wrong. The crucial link exists: The common property of science and art is the transmission of information, and in one sense the respective modes of transmission in science and art can be made logically equivalent. Imagine the following experiment: A team of scholars—led perhaps by color-challenged Mary—has constructed an iconic language from the visual patterns of brain activity. The result resembles a stream of Chinese ideograms, each one representing an entity, process, or concept. The new writing—call it "mind script"—is translated into other languages. As the fluency of its readers increases, the mind script can be read directly by brain imaging.

In the silent recesses of the mind, volunteer subjects recount episodes, summon adventure in dreams, recite poems, solve equations, recall melodies, and while they are doing this the fiery play of their neuronal circuitry is made visible by the techniques of neurobiology. The observer reads the script unfolding not as ink on paper but as electric patterns in live tissue. At least some of the thinker's subjective experience—his feeling—is transferred. The observer reflects, he laughs or weeps. And from his own mind patterns he is able to transmit the subjective responses back. The two brains are linked by perception of brain activity.

Whether seated across from one another at a table, or alone in separate rooms or even in separate cities, the communicants can perform feats that resemble extrasensory perception (ESP). But only superficially. The first thinker glances at a playing card he holds cupped in his hand. With no clue other than the neural imagery to guide him, the second thinker reads the face of the card. The first thinker reads a novel; the second thinker follows the narrative.

Accurate transmission of the mind script depends as much as conventional language does on the commonality of the users' culture. When the overlap is slight, the script may be limited in use to a hundred characters; when extensive, the lexicon can expand to thousands. At its most efficient, the script transmits the tones and flourishes indigenous to particular cultures and individual minds.

Mind script would resemble Chinese calligraphy, not only a medium employed for the communication of factual and conceptual information, but also one of the great art forms of Eastern civilization. The ideograms contain subtle variations with aesthetic and other subjective meanings of their own shared by writer and reader. Of this property the Sinologist Simon Leys has written, "The silk or paper used for calligraphy has an absorbent quality: the lightest touch of the brush, the slightest drop of ink, registers at once—irretrievably and indelibly. The brush acts like a seismograph of the mind, answering every pressure, every turn of the wrist. Like painting, Chinese calligraphy addresses the eye and is an art of space; like music, it unfolds in time; like dance, it develops a dynamic sequence of movements, pulsating in rhythm."

AN OLD IMPASSE nonetheless remains: If the mind is bound by the laws of physics, and if it can conceivably be read like calligraphy, how

Free will

can there be free will? I do not mean free will in the trivial sense, the ability to choose one's thoughts and behavior free of the will of others and the rest of the world all around. I mean, instead, freedom from the constraints imposed by the physiochemical states of one's own body and mind. In the naturalistic view, free will in this deeper sense is the outcome of competition among the scenarios that compose the conscious mind. The dominant scenarios are those that rouse the emotion circuits and engage them to greatest effect during reverie. They energize and focus the mind as a whole and direct the body in particular courses of action. The self is the entity that seems to make such choices. But what is the self?

The self is not an ineffable being living apart within the brain. Rather, it is the key dramatic character of the scenarios. It must exist, and play on center stage, because the senses are located in the body and the body creates the mind to represent the governance of all conscious actions. The self and body are therefore inseparably fused: The self, despite the illusion of its independence created in the scenarios, cannot exist apart from the body, and the body cannot survive for long without the self. So close is this union that it is almost impossible to envision souls in heaven and hell without at least the fantastical equivalent of corporeal existence. Even Christ, we have been instructed, and Mary soon afterward, ascended to heaven in bodies—supernal in quality, but bodies nonetheless. If the naturalistic view of mind is correct, as all the empirical evidence suggests, and if there is also such a thing as the soul, theology has a new Mystery to solve. The soul is immaterial, this Mystery goes, it exists apart from the mind, yet it cannot be separated from the body.

The self, an actor in a perpetually changing drama, lacks full command of its own actions. It does not make decisions solely by conscious, purely rational choice. Much of the computation in decision making is unconscious—strings dancing the puppet ego. Circuits and determining molecular processes exist outside conscious thought. They consolidate certain memories and delete others, bias connections and analogies, and reinforce the neurohormonal loops that regulate subsequent emotional response. Before the curtain is drawn and the play unfolds, the stage has already been partly set and much of the script written.

The hidden preparation of mental activity gives the illusion of free will. We make decisions for reasons we often sense only vaguely, and seldom if ever understand fully. Ignorance of this kind is conceived by the conscious mind as uncertainty to be resolved; hence freedom of choice is ensured. An omniscient mind with total commitment to pure reason and fixed goals would lack free will. Even the gods, who grant that freedom to men and show displeasure when they choose foolishly, avoid assuming such nightmarish power.

Free will as a side product of illusion would seem to be free will enough to drive human progress and offer happiness. Shall we leave it at that? No, we cannot. The philosophers won't let us. They will say: Suppose that with the aid of science we knew all the hidden processes in detail. Would it then be correct to claim that the mind of a particular individual is predictable, and therefore truly, fundamentally determined and lacking in free will? We must concede that much in principle, but only in the following, very peculiar sense. If within the interval of a microsecond the active networks composing the thought were known down to every neuron, molecule, and ion, their exact state in the next microsecond might be predicted. But to pursue this line of reasoning into the ordinary realm of conscious thought is futile in pragmatic terms, for this reason: If the operations of a brain are to be seized and mastered, they must also be altered. In addition, the principles of mathematical chaos hold. The body and brain comprise noisy legions of cells, shifting microscopically in discordant patterns that unaided consciousness cannot even begin to imagine. The cells are bombarded every instant by outside stimuli unknowable by human intelligence in advance. Any one of the events can entrain a cascade of microscopic episodes leading to new neural patterns. The computer needed to track the consequences would have to be of stupendous proportions, with operations conceivably far more complex than those of the thinking brain itself. Furthermore, scenarios of the mind are all but infinite in detail, their content evolving in accordance with the unique history and physiology of the individual. How are we to feed that into a computer?

So there can be no simple determinism of human thought, at least not in obedience to causation in the way physical laws describe the motion of bodies and the atomic assembly of molecules. Because the individual mind cannot be fully known and predicted, the self can go on passionately believing in its own free will. And that is a fortunate circumstance. Confidence in free will is biologically adaptive. Without it the mind, imprisoned by fatalism, would slow and deteriorate.

Thus in organismic time and space, in every operational sense that applies to the knowable self, the mind *does* have free will.

FINALLY, given that conscious experience is a physical and not a supernatural phenomenon, might it be possible to create an artificial human mind? I believe the answer to this philosophically troubling question to be yes in principle, but no in practice, at least not as a prospect for many decades or even centuries to come.

Descartes, in first conceiving the question over three centuries ago, declared artificial human intelligence to be impossible. Two absolutely certain criteria, he said, would always distinguish the machine from a real mind. It could never "modify its phrases to reply to the sense of whatever was said in its presence, as even the most stupid men can do," and it could never "behave in all the occurrences of life as our reason makes us behave." The test was recast in operational terms by the English mathematician Alan Turing in 1950. In the Turing test, as it is now generally called, a human interpreter is invited to ask any question of a hidden computer. All he is told is that either another person or a computer will answer. If, after a respectable period of time, the questioner is unable to tell whether the interlocutor is human or machine, he loses the game; and the mind of the machine is accorded human status. Mortimer Adler, the American philosopher and educator, proposed essentially the same criterion in order to challenge not just the feasibility of humanoids but also the entire philosophy of materialism. We cannot accept a thoroughly material basis for human existence, he said, until such an artificial being is created. Turing thought the humanoid could be built within a few years. Adler, a devout Christian, arrived at the same conclusion as Descartes: No such machine will ever be possible.

Scientists, when told something is impossible, as a habit set out to do it. It is not, however, their purpose to search for the ultimate meaning of existence in their experiments. Their response to cosmic inquiry is most likely to be: "What you suggest is not a productive question." Their occupation is instead exploration of the universe in concrete steps, one at a time. Their greatest reward is occasionally to reach the summit of some improbable peak and from there, like Keats' Cortez at Darien, look in "wild surmise" upon the vastness beyond. In their ethos it is better to have begun a great journey than to have finished it, better to make a seminal discovery than to put the final touches on a theory.

The scientific field of artificial intelligence, AI for short, was inaugurated in the 1950s hard upon the invention of the first electronic computers. It is defined by its practitioners as the study of computation needed for intelligent behavior and the attempt to duplicate that behavior using computers. A half century of work has yielded some impressive results. Programs are available that recognize objects and faces from a few select features and at different angles, drawing on rules of geometric symmetry in the manner of human cognition. Others can translate languages, albeit crudely, or generalize and classify novel objects on the basis of cumulative experience—much in the manner of the human mind.

Some programs can scan and choose options for particular courses of action according to preselected goals. In 1996 Deep Blue, an advanced chess-playing computer, earned grand master status by narrowly losing a six-game match to Gary Kasparov, the reigning human world champion. Deep Blue works by brute force, using thirty-two microprocessors to examine two hundred million chess positions each second. It finally lost because it lacked Kasparov's ability to assess an opponent's weakness and plan long-term strategy based in part on deception. In 1997 a reprogrammed and improved Deep Blue narrowly defeated Kasparov: the first game to Kasparov, the second to Deep Blue, then three ties and the final game to Deep Blue.

The search is on for quantum advances in the simulation of all domains of human thought. In evolutionary computation, AI programmers have incorporated an organismlike procedure in the evolution of design. They provide the computers with a range of options in solving problems, then let them select and modify the available procedures to be followed. By this means the machines have come to resemble bacteria and other simple one-celled organisms. A truly Darwinian twist can be added by placing elements in the machines that mutate at random to change the available procedures. The programs then compete to solve problems, such as gaining access to food and space. Which mutated programs will be born and which among the neonates will succeed are not always predictable, so the "species" of machines as a whole can evolve in ways not anticipated by the human designer. It is within the reach of computer scientists to create mutable robots that travel about the laboratory, learn and classify real resources, and thwart other robots in attaining their goals. At this level their programs would be close to the instinctive repertories not of bacteria but of simple multicellular animals such as flatworms and snails. In fifty years the computer scientists—if successful—will have traversed the equivalent of hundreds of millions of years of organic evolution.

But for all that advance, no AI enthusiast claims to have a road map from flatworm instinct to the human mind. How might such an immense gap be closed? There are two schools of thought. One, represented by Rodney Brooks of the Massachusetts Institute of Technology, takes a bottom-up approach. In this version, the designers would follow the Darwinian robot model to higher and higher levels, gaining new insights and elaborating technology along the way. It is possible that in time, humanoid capability might emerge. The other approach is top-down. Favored by Marvin Minsky, a founding father of AI and colleague of Brooks at MIT, it concentrates directly on the highest-order phenomena of learning and intelligence as they might be conceived and built into a machine without intervening evolutionary steps.

In the teeth of all pessimistic assessments of human limitation likely to be raised, human genius is unpredictable and capable of stunning advances. In the near future a capacity for at least a crude simulation of the human mind might be attained, comprising a level of brain sciences sophisticated enough to understand the basic operations of the mind fully, with computer technology advanced enough to imitate it. We might wake up one morning to find such a triumph announced in the *New York Times*, perhaps along with a generic cure for cancer or the discovery of living organisms on Mars. But I seriously doubt that any such event will ever occur, and I believe a great majority of AI experts are inclined to agree. There are two reasons, which can be called respectively the functional obstacle and the evolutionary obstacle.

The functional obstacle is the overwhelming complexity of inputs of information to and through the human mind. Rational thought emerges from continuous exchanges between body and brain through nerve discharges and blood-borne flow of hormones, influenced in turn by emotional controls that regulate mental set, attention, and the selection of goals. In order to duplicate the mind in a machine, it will not be nearly enough to perfect the brain sciences and AI technology, because the simulation pioneers must also invent and install an entirely new form of computation—artificial emotion, or AE.

The second, or evolutionary, obstacle to the creation of a humanoid mind is the unique genetic history of the human species. Generic human nature—the psychic unity of mankind—is the product of millions of years of evolution in environments now mostly forgotten. Without detailed attention to the hereditary blueprint of human nature, the simulated mind might be awesome in power, but it would be more nearly that of some alien visitor, not of a human.

And even if the blueprint were known, and even if it could be followed, it would serve only as a beginning. To be human, the artificial mind must imitate that of an individual person, with its memory banks filled by a lifetime's experience—visual, auditory, chemoreceptive, tactile, and kinesthetic, all freighted with nuances of emotion. And social: There must be intellectual and emotional exposure to countless human contacts. And with these memories, there must be meaning, the expansive connections made to each and every word and bit of sensory information given the programs. Without all these tasks completed, the artificial mind is fated to fail Turing's test. Any human jury could tear away the pretense of the machine in minutes. Either that, or certifiably commit it to a psychiatric institution.