THE RISE AND FALL OF REALITY by Dan Bruiger ©2008

Table of Contents

"The brain is not an organ of thinking but an organ of survival, like claws and fangs.
It is made in such a way as to make us accept as truth that which is only advantage"
-A. Szent-Gyorki

INTRODUCTION

There are two great mysteries: why there is anything at all and why there is conscious experience of it. These are related in the fact that only in consciousness do we know anything of the world. The being of the world and the being of the conscious self have in common a certain insubstantiality. Of what, ultimately is the physical universe made? Physics inclines more and more to answer that the difference between something and nothing is ultimately very subtle. And of what are our perceptions, thoughts and dreams made? Any adequate answer will surely involve us in even greater subtleties.

The ancient sage wondered whether he had awakened from dreaming of himself as a butterfly, or whether he was indeed a butterfly dreaming himself to be a man. The metaphor of dreaming recognizes the subjective status of experience, a realm independent of the physical world, which is the object of scientific inquiry. There are fashions in paradigms and philosophies, even in science, an ebb and flow of faith in realism and in rationality itself, as well as in specific theories. By placing such changes in the context of fashion, one is already attempting to establish a framework larger than competing camps—a meta-rational perspective within which cycles of interest in one or another world view have their births and suffer their declines.

Are we free spirits dreaming this conditional existence, or material organisms dreaming of unconditional freedom? Are we to understand the experience and the journey of individual and collective too in the progressive terms of modern science or in terms of the world's timeless spiritual teachings? Which is fundamental, mind or matter, subject or object? Facing such questions intellectually, the ancient sages might have hedged their bets, perhaps joining with modern physicists to merge butterfly and man in some monstrous superposed state. Usual strategies in philosophy have considered ether mind or matter to be primary and real, its complement derivative or illusory. My approach will be to seek a higher ground on which a grasp of the true nature and difficulty of the problem itself constitutes a solution of sorts—as close, perhaps, as the self-conscious intellect can come to understanding the dilemma of its own existence. If we are spiritual beings whose essential stuff is idea, the great mystery is our relationship to bodily life in this physical world. And if we are material beings, the mystery is the inner theater of our conscious experience and our capacity for self-transcendence.

The story line of spiritual idealism holds that mind consciously chooses its identification with matter, is ensnared by it, or even *creates* matter. The story line of scientific materialism holds that mind arises from matter through evolving complexity, driven by reproductive selection. The concept of intentionality, however, is capable of linking such divergent outlooks: something within the system is chosen to symbolically represent something else (whether outside it or within it). The choice is free insofar as the system is considered formally, as an abstraction. It is bound in that it is driven concretely by natural forces. In the spiritual story, we fell into bondage to matter and we struggle to regain a free state. In the materialist version, we began in bondage to matter, have conceived the possibility of freedom, and strive with uneven success to make it actual. Such parallel explanations promise to meet on a receding horizon.

The search in both worlds is for underlying unity and first principles. Science inquires about the physical world, the object of experience. Religion inquires about the subject and its relationship to experience at large. Both science and spiritual practice seek to purify experience of subjective idiosyncrasies; both aim at ultimate truth. Science seeks the underlying unity of objects, while religion seeks the underlying unity of subjects. Science inquires into the ultimate truth of the object, spirituality into that of the subject.

These dualities—of subject and object, inner and outer, mental and physical, "mere experience" versus "reality"—are known collectively in philosophy as the Mind-Body Problem. At root the "problem" is the awareness that the contents of consciousness may refer either to something material in the world or to some artifact of the mind. For instance, I may point to a tree in my yard. But the tree may appear in a dream as well, and I may "point" to it as an element of my personal experience, or to its greenness as a "quality," for which I feel obliged to take some responsibility since I can, under some circumstances, experience "green" detached from any object. Because this ambiguity presupposes consciousness of self, the mind-matter split is an inevitable concomitant of subjectivity and selfhood.

Self-consciousness may be defined as the ability to place a frame around some portion of experience, so as to consider its contents as a domain distinct from the world. Standing back to appreciate experience *as such*, rather than as an attribute of reality, defines the observer as a presence distinct from the scene observed. The very notion of experience presupposes self-consciousness and becomes a category in its own right.

Traditionally, the Realist (or Materialist) believes in the physical world as fundamental. The Realist points to the scene itself, thereby ignoring the frame and treating the view as a transparent window upon an objective world beyond. The Idealist points to consciousness as primary, to the frame itself, treating the view as a painting, movie, or "show." One focuses on the role of the world in determining experience ("the view"), the other on the role of the self. My contention will be that there is a deeper level in which these viewpoints are reconciled as shifting forces in a dialectic.

By temperament one finds oneself in natural sympathy with either Realism or Idealism. The split between mind and matter thus divides thought about it as well. For the Idealist, the territory is but an illusory projection of the map. For the Realist, the map is but an artifact, sketched from the territory. But either premise can lead to its own demise. For example, if the territory is an illusion, where do maps come from and why should they be considered real? And if the map is an artifact, it is nevertheless only through it that we know the territory. In this regard the Mind-Body Problem resembles the classic paradoxes of logic. As with such paradoxes, self-reference is at the heart of the Mind-Body Problem. What makes paradox is the inability to get above its logical level—to transcend or synthesize its contradictory implications. This is the situation confronting the mind conscious of its own subjectivity.

In either an Idealist world or a Realist world there would be no Mind-Body Problem. But subjective consciousness has cast us irreversibly into mixed terrain. There is no way to determine whether experience (or its extension, thought) is an accurate or literal representation of the world-in-itself, unsullied, so to speak, by mind. Experience and the world cannot be hung up side by side for comparison. For, everything we know of the world comes to us *in* experience, while every content of experience is in some way an image or reflection of the world. Dwelling in the map, nothing secure is known of the territory of which it is allegedly the map. And dwelling in the territory, no amount of climbing about in the machinery of the brain (to paraphrase Leibniz) gives any hint of what the color blue should be like as an experience, or why there should exist any such thing as experience or consciousness at all.

The Mind-Body Problem is the utter disjunction between subject and object, between first and third-person points of view. This points to logical entanglement as the culprit rather than some empirical fact. If so, conscious experience cannot be explained in terms of some causal power of the brain (nor as a quantum effect, etc.), anymore than in terms of a non-physical vital force. How the behavioral output of the brain emerges from the complex activities of many nerve cells, responding to inputs from the world, may be a challenging problem, but it is probably well within the purview of science. What remains elusive is how the complex behavior of neurons gives rise to the *experience* of color, for example, or to the very sense of selfhood and consciousness. Or, how the experience of *willing* a muscular movement leads to the movement. The subjective self and its experience are unquestionably correlated to physical brain-body states, but they are not physical objects in the world to be described from an observer's third-person point of view.

Subjective experience cannot be explained scientifically for the simple reason that science deals only with objective description. How can there be a strictly scientific explanation of consciousness, which

appears bound up with notions of personhood, selfhood, mind, spirit—the ineffable ethos of the subject? The existence of the conscious self, and the perspective from which "looking" takes place, must remain a mystery to an outlook that only looks outward toward objects. The places in science where this outlook is inevitably confronted by the role and state of the observer are famously problematic: relativity, quantum physics, and cosmology.

Experience was at one time explained as a spectacle appearing to an inward "soul." But to treat it as a kind of inner cinematic record of the external world implies an audience for this show. Then *that* observer's experience must in turn be explained, in an infinite regression of the metaphor. The mystery of experience cannot be explained by multiplying the subject. The notion of a subject or perceiving soul cannot *explain* perception; it merely stands in for the process to be explained. Does my soul have its own eyes to see, or its own brain, apart from those of my body?

Is it possible to transcend the Mind-Body Problem, to find a common ground between mind and matter, between Idealism and Realism, between metaphysics and physics? Theories of "emergence" and "neutral" concepts of information processing in complex adaptive systems seem promising within the scientific framework, where mind is studied as a self-regulating representational system. But there is still something missing in this picture, which remains a third-person description. Following such extrinsic metaphors, one never quite closes the gap with subjective experience, the first person. It is never quite clear how bustling neurons produce the intrinsic blueness of the color blue. I will argue that what is needed for this closure is a grasp of (a) the gratuitous nature of intentionality, and of (b) specific conditions imposed by the context of embodied participation in an evolutionary history. A cognitive system is a symbolic system, like a mathematical language that is an *interpreted* formal system, or like a real language that refers to the world. From such a perspective, pain can be understood to have its specific experienced quality because of its particular meaning to the organism as information about the state of its tissues. Ultimately, we should like to understand why the sky looks blue to us—rather than green, say—that is, why we have the subjective experience we do in response to light of a certain wavelength. This sort of explanation goes beyond the functioning of causal systems, to embrace the evolutionary advantages of particular intentional connections within evolved systems of meaning. The very nature of intentionality takes us beyond the science of passive matter.

Whatever the details, such explanation can only be based on the reasonable assumption that cognition is neither entirely determined by a world of external causes, nor entirely by the organism's internal symbolic conventions. It is rather a specifiable interaction in which both organism and environment contribute to the creation of experience and meaning. Cognition is rooted in the particular programs of the organism as well as in the environment, and is always a biased and incomplete account of the world, whose "true" face cannot be known or even assigned a meaning. (This much we owe to Kant). One cannot strictly speak of the accuracy of the map, but only of its adequacy as a tool of survival, or of reproductive success, since map and territory cannot be compared in the way common sense would like. We are hopelessly immersed in the map, and only able to perceive the territory through its mediation. On the other hand, the map itself can only be perceived as part of the territory. "Raw" (unmapped) territory is a concept that does not compute. It involves us in endless tail-chasing, for it can only mean the unseeable face of the world-in-itself, while we *see* by means of formal configurations of the map. Any concept of the territory is therefore but a feature of the map, which itself must be located somewhere in the territory…! Undaunted by such recursive circularity, mind leaps ahead to achieve an experience of a real external world compounded from its inner map; it somehow creates meaning from the babble of the senses.

Now, a formalism need not be interpreted in any particular way, nor even at all. Like a game, it is also gratuitous, a self-enclosed world in its own right. In that light, its premises, like the rules of a game, are arbitrary. Existence is profoundly irrational in just this way. On the positive side, our little life is whimsical and undetermined, the exuberant play of free beings. On the darker side, it can seem a tale of sorrowful entrapments in absurd circumstances or irrationality, the inertial ruts and insidious logics that dominate all scales of human experience and behavior. The face of the arbitrary has its comic and tragic masks.

The Real sets a compass of apparent order, meaning, direction and certainty on the face of the arbitrary, the Unknown. But the *sense* of reality is routinely used to endorse unconscious choice, to lend validity to opinion or motive. Natural reality has, like a good parent, given us structure and stable bearings in the time of our evolutionary childhood. But as we come of age, individually and collectively, it is up to us to define ourselves and make the world we want to inhabit.

Part One: MATTER AND MIND

1. *The Real*

The "Mind-Body Problem" has been defined in many ways throughout the course of Western thought. Perhaps this is because the only tool that philosophy really has is the ability to reorganize thinking. One definition could be: the divergence between a materialist and an idealist view of experience. Which is primary and real: the physical world or our personal experience of it? Clearly, this may be a chicken-or-egg sort of problem.

Alternatively, one could talk of "objective" and "subjective." A cautionary note about these terms is in order: they are biased in our materialist culture, where we tend to identify objectivity with reality. We have a genetic investment in "reality" as representing survival (Freud's Reality Principle). The subjective, on the other hand, tends to be pejoratively identified with the imaginary and irrational or the merely personal, private, and interior.

In one sense, the Mind-Body Problem is the confusion that arises from thinking of the subject (mind) as a subtle object (body). For instance, a memory is an "object" of consciousness, inviting us to wonder what *sort* of object it is, compared to the event it is a memory *of*. How does a mental image of a chair compare to the chair itself, how do we classify them, and how might they interact?

Materialism holds that the physical world presented to us in experience as *out there*, solid, and real, is what truly exists. The body, as a component of that physical reality, is the part of us that is "objectively" real. The world outside the body interacts with the nervous system, initiating complex material responses within the body's brain. These lead to objectively observable behavior, but also to subjective experience. In the materialist view, it seems that events in the world, the body, and the brain must *cause* our experience, which, however, is nonmaterial, even though somehow a byproduct of material events. It is this sense of an interior, non-physical realm we call subjective that gives rise to the MBP. For, how can the physical and the non-physical interact causally? How can one produce the other? In a materialist perspective, what place is there for a non-physical realm at all? We are therefore tempted to imagine the subjective realm as occupying a quasi-material existence along side the physical, somewhat on the same footing, and to try to ponder the interactions between them. (Isn't this how we conceive the "soul", or the "subtle bodies" of metaphysics—as quasi-physical entities subtler than physical matter?)

In contrast, idealist world-views hold that the material world is not "real," but a creation of the mind, or an illusion imposed upon it. Causality is reversed—and with it responsibility. For it is not physical reality that causes experience "in the mind," but the mind that causes the appearance of a physical reality "out there." Mind becomes primary and causative, matter becomes derivative and illusory.

The general thrust of cultural and technological development that we know as civilization can be viewed as a program of withdrawal from Nature and the body into certainties of the mind. No doubt this was originally motivated by the vulnerabilities of the physical self. Only recently are we coming to terms

with the vulnerabilities of the mind, which has become a new raw material for manipulation and exploitation. Were Descartes alive in the Information Age, he would have fresh arguments for skepticism.

Both the materialism of science and the idealism of spiritual traditions have sought the eternal and changeless beyond the realm of experience, but have pursued diverging strategies. Science looks for unifying principles behind the diversity of sensory experience by braceketing the subject of that experience as irrelevant. Spiritual idealism holds that the one underlying principle *is* the subject—consciousness, "I," the Self. From such a perspective, the MBP is less an ontological problem (a "category mistake") than a question of mistaken identity. We erroneously think we are the body, identifying with its perspectives and priorities. From a materialist perspective, however, there is an external reality that holds sovereign sway over the body's priorities, since the body is part of that reality and contingent upon events within it. Hence the widespread belief within our culture that mind is at the mercy of the brain, an organ of the body subject to physics and chemistry. Idealism reverses this sequence: mind invents the world, including the body and its brain, and their possible defects.

With this preamble, let us consider the "realness" we associate with physical objects to be a category of the mind, such as Kant held space and time to be. Consider it a *quality* with which mind imbues sensory experience, because mind is programmed to serve the body's need to survive, reproduce, and maintain itself in a zone of well-being. To survive, the orgaanism must take the external world seriously—which it does by considering it *real*.

The organism is in constant exchange with the rest of the physical world. Each creature receives impressions from the environment, which it is the job of its brain to interpret in such a fashion as to allow the creature's genetic success. Natural selection has guaranteed that by definition only those organisms exist that play this "game" successfully.

To indulge the game metaphor further, a game is a structured activity within a playing space or field, with playing pieces and rules. Chess and Monopoly are examples. So are algebra and geometry, computer programs, computer hardware and other machines, human relationships, including business and economics, and just about any system you can think of that involves ordered, structured activity. In the *game of natural selection*, the playing piece is the organism and the playing field is the environment. The rules are numerous (perhaps infinite), including the laws of physics, chemistry, and genetics. And there are many levels of rules, including what might be called the logic of the organism: how is the organism "programmed" to cope with the environment? The algorithms governing its behavior are rules and strategies in the game of survival. In a computer game, some software is embodied in the circuitry of the hardware. So with the brain/body: its "logic" is partially hard wired in its neurology and chemistry. Whereas a human programmer created the design of the machine, however, eons of cumulative genetic trial and error honed the skill of the organism at the game of life.

The brain is an organ of survival, and some creatures rely more heavily on it than others. The big-brained creatures tend to be visually and acoustically well-developed. What these two senses have in common is space-orientation, with the movements of pursuit and evasion. The concept of space also implies the independent existence of distinct objects—and therefore ideas of objectivity and the material reality of the world.

Of course, we take for granted that the world simply *is* as our big brains depict. We believe that the space we perceive is real, that the objects in it are real, and so forth, without questioning what "real" means—or what "meaning" means, as far as that goes. But perhaps the 3-D realness we automatically ascribe to perception is itself a survival strategy, like physical pain. Clearly there is a reason why pain "hurts." The hurting is the way we experience the mind's judgment (and alarm!) that the body's tissues are sustaining damage. Our physical existence cannot be free from pain, because it is the alert of pain that has allowed us to survive. Similarly, the realness with which the mind imbues physical reality is the way we subjectively experience its appraisal that the playing field of the world is a very serious place, full of consequences for the body. This too has allowed us to survive. It expresses the belief that the world holds over us the power of life and death, pleasure and pain. It reflects our earnest engagement in the game of survival. Like pain, the sense of reality is functional. We are here to casually reflect about all

this only because our ancestors took the Reality Principle, as Freud called it, so seriously that they lived long and well enough to reproduce.

2. *The Situation*

Our brains are big enough that we can theorize about how brains work and what kind of epistemic situation we are in as perceiving organisms. We can paint a picture about how minds paint pictures, and grasp that a picture a mind paints is indeed a creative invention with survival value, rather than a transparent window on the world. The Situation is that even the picture about picture painting, so to speak, is an artifact hanging on the wall of the mind, and not (as we naturally are inclined to think) an opening in that wall revealing a "real" world that is the proper subject of painting!

If all this tail chasing makes the head spin, the Situation is the very problem that our brains are big enough to think in circles. The problem, if we are to call it such, is that we are conscious of our consciousness, as opposed to simply being aware of the world. In this inevitable whirlpool of self-awareness, we back paddle through metaphor to try to comprehend The Situation. Aware of being aware, we *cannot* think of experience simply in terms of the world. Now that we know about painting, all scenes are suspect. The beautiful vista, examined closely, might turn out to be a clever backdrop, a *trompe l'oeil*. Plato described The Situation as like that of prisoners in a cave from birth, who know nothing of the outside world, but only experience shadows cast upon the cave walls by various objects from outside, which are never themselves seen. Descartes described it in terms of a clever demon who could systematically falsify all the information coming into a brain through the senses, creating the airtight illusion of a body and a material world.

Of course, this is all against common sense, which insists we can certainly tell the difference between a thing and a painting of it! The problem is that our metaphors (and our common sense) may already have deception built into them. The difference between the original and a copy, between fact and illusion is something we think we already know about. It is a distinction *within* our experience of the world, already presumed to be real.

Imagine yourself, however, inside a vast theater set. As long as you aren't too careful, the illusion of a scene from the real world can be maintained. But eventually you may encounter some detail that is inconsistent: the back side of a brick wall that is only boards leads you to look more closely and see that the bricks are just paint on canvas. This is because, with finite resources, only so much care was invested in creating the illusion. (By the same token, only so much care can be invested to detect the difference between real bricks and fake ones.) But now imagine a *Star Trek* Holodeck illusion—a *complete* simulation down to the last detail. Imagine further that you have never been outside the Holodeck, and have never been given any clue that there even *is* an outside. All your life you have grown up and interacted with fictional parents and friends, teachers, associates and a cast of millions who have gone about their business and aged as you have in an environment that is totally consistent within, yet fake from an outside perspective. It is the only reality you know. But can you know whether it is real?

Even in this metaphor, of course, we are actually biased by our prior knowledge that *Star Trek* is a fictional story, and that even within the story the Holodeck is a fictional reality. So we are *assuming* a primary reality that is true (we go and turn on the TV, have dinner in front of it while watching the show, etc.). But suppose the truth is that we really *did* grow up on a Holodeck, and the program running there (even at this moment) is your current experience. This experience includes the memory of making dinner in your holographic home, tuning in to *Star Trek* on your illusory TV, and all the ideas you have about Holodecks and primary and illusory realities. Now what? Pinch yourself? Perhaps your body too is a Holodeck illusion, and in the primary reality there is no such thing as bodies. Welcome, then, to The Situation! Of course, there is so much to *do* and experience within the Hologram that we can spend a lifetime there without ever being concerned that it might all be a highly consistent illusion. Such is life! Even when the dream is a nightmare, it *is* fascinating!

The Situation for the human brain/body is that it finds itself to be a separate unit, a distinct

playing piece in a playing space that seems to be a physical world of discrete objects, in a game that seems to be about survival. Survival means the success of genes, echoed in the well-being of the individual and its extensions, as defined in the context of the game. That is, if the purposes of the individual (which may be merely the purposes of genes) are accomplished or furthered, the experience is *defined* as pleasure. And if they are hindered, the experience is pain.

Thus, to identify with the organism is by and large to be what is called, in Game Theory, a "rational player": one who takes the game seriously and plays in earnest to win. The player identifies with the playing piece, adopts the goal of winning, and accepts the rules and conditions of the game, as well as principles of play that could lead to success. In an ordinary game, say of Chess or Monopoly, the players may become quite involved in the game, but they never lose sight of the fact that it is just a game; and when the game is over they carry on with their real life. We would question the sanity of a person who thought he actually *was* a Chess knight, for example, who could only jump about in L-shaped leaps and who was permanently obsessed by the moves and logic of Chess. Nevertheless, in life we *do* take the game so seriously that we think we *are* the roles we play, the social character with whose purposes and limitations we identify; and we are obsessed with the "moves" of our daily lives. Certainly we do give credence that our daily experience is real and not a dream. But are we not like actors who have become so caught up in the script, in the action of the play, that they have forgotten their conscious choice to "willingly suspend disbelief"?

If so, how did this all come about? Well, to begin with the obvious, realism is the ground state for a player in the natural drama, which is not a game to dabble at: the world outside our skins demands our attention and respect. As in any competitive sport, only those players who play in earnest are going to stay in the game. So, if you find yourself here, it's because you come from a long line of successful players. And basically that means your lineage has agreed to bracket any larger context for your existence and get thoroughly caught up in the game. Our very existence speaks volumes about the kind of creature we must be.

The game—the drama—is only possible, meaningful, interesting, and *real* because we adopt the rules and agree to start somewhere, embracing the fantasy "world" generated by its premises. Logic has power only because we believe in the laws of deduction, which enable us to leapfrog from one idea to another. What can be proven is only as reasonable as the initial assumptions. Similarly, in geometry one begins with certain axioms that cannot themselves be proven, but must simply be taken for granted in order to construct theorems. We embrace the game, or the system, in other words, through a voluntary gesture. The "initial conditions" are supplied by the game of life: the evolutionary history of the biosphere. All we need do is click our heels (or click our computer mouse) as though to say "yes, indeed, I do accept these rules and initial conditions", and before we know it we are swept into the world of that system. We are in The Situation because we have *agreed* to it—or, rather, our genes have agreed on our behalf. We have signed a contract of loyalty to the Corporation of Physical Existence, and now we are company men and women. We *are* the body living *in* the world because we have accepted the purposes and point of view of the body as a playing piece in the game of life, on the playing field of the material and social world. It seems that once in Oz, however, it is not so easy to get back home. Once we enter a system, an entrapment occurs, because such a system or game is by nature closed upon itself and complete within itself. It's a kind of logical black hole.

Imagine a life-sized Monopoly world where you go around buying and selling commodities and competing to build an empire. (Any resemblance to the so-called real world is hardly coincidence. Monopoly was invented during the Great Depression, and became popular because people wanted a fantasy relief from their real economic situation.) The only permissible activities in this world, and the only things that exist in it, are those defined for the game of Monopoly. Not only would it be rude, but you *cannot* pick your nose because that is not defined in the rules of Monopoly. You also cannot rest, make love, go out to dinner or raise children, because these things have no defined place in the game. They simply do not exist in that world. Nor does there exist any larger world beyond the borders of the game, because that is not defined either. In truth, of course, you do live in a larger world from which you momentarily enter or exit the world of Monopoly at will. But imagine you got stuck there and couldn't

remember even that you had a real life. Is there anything in the world of the game to remind you of this fact, any hints suggesting a larger world beyond? Or does every action in the game simply recycle you back into the game, like a dream from which you cannot awaken?

Monopoly may be a crude example. Just as with the metaphor of growing up on a Holodeck, we can easily imagine more sophisticated virtual worlds. The point, however, is that every game defines such a virtual world, and every such world is a product of some game. The limit of the game we are playing is the limit of the world we can perceive.

Now, if the game or system is sophisticated enough, there *is* a way out. The more complex the simulated world, the more perfect and seamless the illusion. At the same time, however, the more complex the game or system upon which the simulation is based, the more powerful it is in its own right. A system can ultimately be powerful enough to find its own zipper, to to speak, and break out of itself. And the degree of power that is required seems to be the capacity to self-refer.

Returning to Monopoly, one might draw a Community Chest card, for instance, that reads: "Pass Go and collect $200". It is an instruction about what lawful move to make next. It's by following such instructions, one after another, that we remain in the game, however, because each instruction only points to the next, and never to anything outside of the game. But if Monopoly could self-refer, you might draw a strange Community Chest card that reads: "This Community Chest card reads: 'This Community chest card reads: This...' [etc...]" You have suddenly entered an infinite hall of mirrors. You can never get to the end of the statement, to figure out what it "finally" says. Moreover, each reiteration of this infinite regression proclaims indirectly that "This is a Community Chest card in the *game* of Monopoly"! It is not just another instruction about what to do next. There is information about the fact that the card is an element of a game, and so our attention is drawn to the fact that we are situated in a game or logical system. Once we realize this, effectively we have already escaped from the system, because that awareness itself implies the existence of some world (or bigger system) outside the one we are in. To be truly trapped is to be ignorant of imprisonment. Once we have found the zipper, we are already out.

So, another way to look at The Situation is this: the cosmos we consent to embrace proves to be an evolving system. As part of it, some biological organisms eventually reach a complexity permitting them (us) to disengage from the identification, to transcend the limits of the game, at least mentally. As organisms, we are able to wake up from the trance induced by organic existence, at least enough to ask these questions about its nature.

What mind considers real and meaningful derives naturally and historically from values related to the body's well-being. Upon this foundation mind has built elaborate cultural structures within which it may claim a certain freedom, and even *disclaim* its animal heritage. It has defined its *own* games, apart from those inherited from Nature, and through which it attempts to circumvent The Situation.

Here we are on paradoxical gorund, however. The Situation continues to present an apparent tradeoff: either we are imprisoned within Nature and bodily experience, benefitting from its ready-made meanings, or else we claim the freedom of disengagement but are then lost in the void of the meaningless and arbitrary, if not imprisoned within our own artifice. The concept of the game, however, suggests a different relationship—a third alternative or middle ground: voluntary play within apparently serious reality, whether natural or cultural. The player agrees to play as though in earnest, consciously adopting the premises of the game, not in order to win but for sheer delight of playing. (Playing to win, in contrast, means bringing the game to a close, which is but an obstacle to play for its own sake.) One plays according to rules and within defined boundaries, whether imposed by nature or society. But one may also play *with* rules and boundaries, to the degree that they involve our willing participation. In the one case, the game determines our possible moves; in the other, we ourselves define the game as we go along, and also retain our identity as players.

Game is in essence a mathematical concept, and such play is the very nature of mathematics, which lends itself to serious application with powerful results. Through mathematics we have gone to the moon. But it forever transcends any particular application, formulation, system, concept, or use. It remains inherently gratuitous, playful, and flexible, while capable of the meanings we lend it. It shows us how to reconcile the desire to live in freedom with the desire for meaning and truth.

3. *The Subjective Frame*

Any *concept* of the Real can only be the notion of a self-conscious mind. The category of the object arises in subjective consciousness together with its polarity, the category of the subject. The story of the Real, therefore, is bound up with that of the self, and begins with the fact that you and I are conscious beings aware of our own consciousness.

From a third-person point of view, cognition is information processing, for monitoring the world and shaping appropriate action upon it. But cognition in the first person is simply experience. The world must therefore figure prominently in experience. Yet we all know that some experience consists of things that cannot be identified as features of the external world. These incongruous elements suggest a realm that is parallel to the objective one, if not quite on the same footing. We want to know of what kind of *stuff* thoughts, hallucinations, feelings, after-images, and referred sensations (e.g. "phantom limbs") are made, if they are not physical things in the external world. Even when the object of awareness is the physical world, we may still have the sense of this awareness as experience happening in the mind, in much the way that news footage happens here and now, on the flat screen of your TV, even though it depicts three-dimensional events recorded far away. The outward orientation of mind tempts one to objectify this mental realm as a kind of rarefied substance awkwardly sharing the cosmos with its material counterpart. But subjectivity is a *mode*, and not only a content, of experience.

Self-consciousness invokes the awareness *of* experience, considered *as* subjective. It is precipitated by the presence of a subjective "frame" bounding something whose objective reality would not otherwise be questioned. This frame is some cue in the experiential field. Without such a cue, whatever is experienced is experienced "naively" as a part or quality of the world; with it, one conceives some portion or aspect of experience in relation to oneself as a perceiving subject. You are not normally aware of the outline of your visual field, for instance. But sometimes this jumps into awareness, reminding you that your own presence is literally as plain as the nose on your face! Alternatively, you might become aware of some bodily sensation as a subjective artifact rather than an external reality. This could be, for instance, a sudden shooting pain that makes no sense as a report of injury; or the scintillating scotomata that, under specific conditions of lighting, appear as tiny squiggles of light or dark in the visual field. Similarly, one normally believes that the color seen when looking at an orange should be ascribed to the surface of the orange. The fact that one can also experience a similar sensation with eyes closed (for instance, after looking at an intense source of blue-violet light) leads to the conclusion that the sensation of color is an artifact of the nervous system, and this so *even when looking at the actual fruit*. That is, even normal perception is a construction, in the subjective domain, even when it refers to the external world.

Irregularities or breakdowns in normal perception are cues that tag impressions in such a way as to identify them as originating with the subject. Thus an inner realm of subject-generated experience is implied by optical and other illusions, hallucinations with psychotropic drugs, various conditions of neurologically injured patients, normal imagination and memory, dreams, somatic sensations, after-images, etc. These diverse experiences all appear *not* to be perceptions of anything physical, even if caused by physical stimuli or referring indirectly to physical events. Perceptual anomalies betray the mediating presence of mind, just as the waviness of old glass in a window betrays the existence of the intervening pane.

Self-consciousness is a form of self-reference, the paradoxes of which may be viewed as boundary issues between categories or logical levels. Hence, the "self-consciousness" of the artist who continues the painting onto the frame; or paints the frame to look like part of the room; or depicts a hand with brush painting the painting; or who hangs a painting that is an image of the room it occupies, an image which includes itself. The frame of a painting is supposed to set off the picture as a separate domain from its surrounding. The image is distinguished from the reality, the map from the territory. But the status of painting and surrounding as domains, or distinct logical levels, may be ambiguous. The

painting is an object *in* the room, but it is also a window *onto* another scene that is of a different order. Everything outside of the frame may be taken as real, logically primary. A border to this primary domain, however, implies that it is not all that exists. The presence of an *image* within this primary domain ultimately calls into question the domain itself. Before the appearance of this border, there can be no such suspicion, no subjectivity—indeed, no self. But once a cognitive system can self-refer, it may find an element within itself that is incongruous with the ostensibly objective world. One's experience then becomes a domain in its own right, first within the borders of the subjective frame—as imagination or dream, for example. But then, (as suspicion spreads) this domain may expand to encompass the domain of the 'world' as a mere subset of 'experience.' As opposed to the world itself, one's experience of the world begins to seem a creative endeavor that intimately involves the self.

This development may be likened to the discovery of irrational numbers, an anomaly that did not fit into the domain of rational numbers (fractions). In order to accept and use them as legitimate mathematical entities, mathematicians opted to expand the domain to include them in a broader definition of number. (The expanded continuum then included the rational numbers as a subset of all possible decimals—the "real" numbers.) Similarly, the recognition of elements of experience that cannot be considered part of the objective external world led to an expansion of the human cognitive domain. The domain of 'the world' was forced by self-consciousness to include such elements in an expanded domain of 'experience.' And the category of experience in turn seems to imply a subject, an experiencer conscious of its role in relationship to objects of experience, both physical and mental. Thus arises the split of subject and object.

In the beginning was *the world*. Just as our eyes cannot see themselves without the aid of a mirror, the mind unaware of itself (as hypothetical as that must remain for us) could not have figured within its own experience, or known its role in producing and regulating that experience. All that we moderns call subjective would have appeared to it simply as *real* or *there*. There could be no subjective domain, only the view through the window. The development of human consciousness, as inferred from ancient myths, seems to have passed through stages similar to those through which the consciousness of an individual develops. We are free to speculate that the early human psyche, like that of a newborn child, merged in an undifferentiated unity with its mother and the world, would have distinguished poorly between inside and outside. From our current subjectified perspective, we may call this 'projection.' But from the perspective of the pre-subjective mind, everything we would presently describe as taking place within the psychological self would be experienced as external objective events. Without a self, only the world exists. (Although, of course, without a self, the *category* of 'world' would no more exist than the category of 'self.')

Just as the maturing child is ambivalent about its dependency, the rise of subjective consciousness must also have involved a conflicted struggle to break away from the securities of a less reflective state. The child begins to assert its will beyond the sphere of voluntary bodily movement, through experimenting with objects and testing itself against other wills. The young ego begins to disengage its own identity and exercise itself through interaction with the world, rebelling against its helpless dependency and against its mother. Similarly, the developing human consciousness must have included rebellion against Mother Nature, against the instinctual body-mind, against controlling parental gods, and against the confines of reality.

The development of subjective consciousness depends on the ability to observe and manipulate internal mental objects as well as objects in the environment. Conscious control *of* the mind means wresting oneself from possession *by* its contents. By bracketing such contents as subjective, one gains control over them; the relation to instinct is loosened, distance and mediacy are acquired, and something purely cognitive is distilled from compelling emotional contents. In our present subjective and overly-mental society, it is difficult to appreciate what must have been a very different situation at the dawn of humanity. While we long to recapture some instinctual vitality, in pre-subjective times the great task must have been to tame the mind's terrors, as well as nature's. This would have meant to gain freedom, not only from actual brutalizing realities, but also from overwhelming perceptions and feelings *experienced as apparent reality*.

Possession *by* a mental content means seeing the world *through* it, so that it is experienced as a feature of objective reality, absolute and imbued with self-evident truth and meaning—about which one can do little. Subjective consciousness is the capacity to take such contents back into the psyche where they can be directed by a conscious will, appropriated as palpable instruments of the self. Through the awareness of awareness, the organism is able to change its internal structure, to voluntarily re-tool itself, so to speak.

Subjective consciousness achieves a more flexible mental instrument—yet ultimately in the service of the body. By experiencing its own experiencing, mind can transcend its rigidities, stepping beyond its habitual categories, perspectives, and assumptions to see them as such. It is then in a position to modify them from a new foothold of relative detachment and freedom. Such a foothold is *only* relative—a movement from lesser to greater objectivity, not a static quality or thing. The paradox of objectivity is that it can only be attained through subjective consciousness. As new truths are conjured from this vantage point, they in turn must be relativized.

4. *Some Mind-Body Problems*

The subjective frame cleaves the experiential field in two, as it were. Some aspects of experience belong to the world, some to the self. Object and subject are disjunct categories: not apples and oranges but apples and *ideas* about apples. It also gives rise to two perspectives. One looks *upon* the object—implicitly from all possible points of view, both spacial and cognitive (while actually from one point of view at a time). The other looks *from* the subject—literally from a dimensionless point of view (while potentially it is all such perspectives). The object—even the object of thought—may have extension in space and time and various sensory qualities. The subject has no qualities and no extension. What it does have is *intention*—which, we shall see, is the very source of experienced qualities.

The subject-object dichotomy itself can be examined from either viewpoint. We can look *through* our experience *at* the organism as physical, and note that its responses are a product both of an objective environment, of which it is a part, and of its own internal programs. We can then infer an experience that is an inner representation of the environment and its own state, and try to grasp the process whereby this representation is constructed. *Or*, we can look *at* our own field of experience, to examine its internal logic and structure, in order to comprehend the "deduction," from sensory input, through which the external reality has been unconsciously constructed as image. In the first case, we begin with the physical world as primary and given (with our bodies and brains as part of it), and try to understand the arising and place of consciousness in that scheme of things. In the second, we begin with consciousness as primary and given, and try to understand the system of cognitive beliefs that, in their sum, give us the experience of a real world. Neither can be proven the "right" approach, for they are simply the same situation viewed from complementary perspectives. There is no ground firmer than either perspective on which to demonstrate such a proof. The best we can do is to hold that the contents of consciousness (whether viewed as phenomenal experience or as perception of the world) are co-created by both self and world.

This co-creation of experience is something like an algebraic product of two factors. This "equation in two variables" can be written:

$$E = f(o, s)$$

where o is the object (world) and s is the subject (self). Of course, this equation can never be "solved" except by arbitrarily holding one variable constant while observing how experience depends on the other. In actuality, the two factors always function together. Nevertheless, physics attempts to simplify the epistemic Situation by eliminating the subject, making experience a straightforward function of causes in the physical world. Psychology, on the other hand (to the degree that it is not merely an extension of physics), seeks to control the input of the environment in order to observe the dependency of experience

on the subject's mind or internal programming.

This formula can implicitly express either a materialist bias (as it does above) or an idealist bias (as below). The world (and the body as part of the world) may be taken to be primary and objective, while experience is a dependant function of it. Or the relation can be inverted, with experience as given, and with both the world and any concept of self as constructs derived from phenomenal experience:

$$(o, s) = f(E)$$

Choosing between these expressions is a moot point, since there is no absolute ground on which to decide between idealism and materialism. Taken together, however, they lead us to conclude that the epistemic Situation of embodied mind implies some limitations to our knowledge. These are that reality cannot be known: (1) as it is "in itself," independent of knowing minds; (2) in its absolute entirety; (3) impartially, apart from biases and interests of particular minds. And yet mind strives to overcome all three limitations. It may hold its experience to represent the world, hope to know everything, and aspire to an impartial objectivity. The three limitations can be restated more formally as follows:

1) The set (o,s) cannot be known except as elements of E. This says that there is no transcendent source of knowledge, unmediated by a subject.

2) The set (o,s) may contain elements that are not in the set E, and vice-versa. This asserts that knowledge or experience can never be a complete account of reality; nor can the reality of either the world or of the subject ever be completely represented in experience or otherwise be mapped by mind (at least not, as we shall see in later sections, if the world is natural and not an artifact). It also implies that experience may contain elements that cannot be accounted for either by the world or the conscious subject as it knows itself.

3) Generally, $f_1(o,s)$ and $f_2(o,s)$ are not identical, and there is no absolute basis for preferring one over another. This says that all minds are different, partial and relative.

These considerations are reflected in facts of everyday life. As a consequence of living in bodies, we are preferentially bound up with our own existence, and literally confined to our point of view in space and time; we are separate and distinct from others, unable to experience other minds from within, nor our own from without. As social beings, we have developed empathy to compensate for the one condition, detachment for the other.

The concept of mental illness is but an exaggeration of the gulf between these conditions, as may be the mental illness itself. When viewed externally, mental illness is a departure from normal behavior, perhaps a nervous or chemical disorder. Viewed from within its own terms, it may be a novel way of organizing experience, following its own reasons, and perhaps a glimpse of truths for which normal consciousness is not prepared

The difference in perspectives also divides knowledge or intellectual inquiry into two broad categories: ontology, which asks what is, and epistemology, which asks how the mind obtains and justifies its knowledge of what is. These have been retained as branches of philosophy, but the distinction between world and self has also given rise to specialized sciences such as physics and psychology.

The Mind-Body Problem is elusive and confusing because *we ourselves* are at the core of it. A function of our way of looking, it is wherever we turn our gaze. When the subjective viewpoint is reified through objectifying eyes, mind appears as a subtle kind of matter. Mind and matter are then seen as two complementary, but incompatible, kinds of "stuff." How do they relate? How, for example, can the will to move one's hand bring about the hand's movement; or how can electrochemical discharges in nerves bring about the sensation of touch? We might be tempted to agree with Gilbert Ryle's famous dismissal of the whole quandary as a "category mistake." But the problem does not go away by heroically declaring it void and meaningless. Its very persistence is a sign of the fundamental and unique dilemma of the self-conscious mind, located upstream, so to speak, from all ordinary strategies of thought. Our task should be to understand the dilemma itself, and the principal obstacle to such understanding in this

case is the pernicious fact that mind cannot stand outside the very problem it attempts to model. For analysis, like perception, takes an external perspective, whereas the subject is not a *thing* but the internal perspective.

The MBP is actually a complex of inter-related dilemmas. These fall roughly into three types. The first concerns the classical incompatibility of mind and matter as categories of being. Secondly, there is the problem of reliable knowledge for subjective beings. Finally, there are difficulties posed by dealing with other subjective minds. These overlap and interpenetrate, and all stem ultimately from the fact of self-consciousness, without which there would be no such dilemmas.

5. *The Ghost in the Machine*

Immaterial consciousness seems to inhabit a physical body and look out upon a material universe. Here is the problem of accounting for consciousness in terms of an external perspective—Leibniz' famous "ghost in the machine":

> It must be confessed... that perception and that which depends on it are inexplicable by mechanical causes... And supposing that there were a machine so constructed as to think, feel and have perception, we could think of it as enlarged and yet preserving the same proportions, so that we might enter into it as into a mill. And this granted, we should only find on visiting it, pieces which push one against another, but never anything by which to explain a perception.[1]

Notice that 'perception' is an ambiguous term: it can refer either to a form of behavior (third-person description) or to a subjective report of phenomenal experience (first person). Accordingly, Leibniz could be speaking of two distinct problems: how scientifically to characterize brain processes, or how to account for the very existence of a first-person perspective (especially in third-person terms)—what has been termed 'the hard problem'. I assume he means the latter, since mechanistic explanations of behavior were already current.

The major philosophical attempts to reconcile the first and third person perspectives have been of two broad types, according to whether they are admitted on an equal footing or whether one is reduced to the other or is otherwise dismissed as unimportant. Dualism holds that there are two kinds of being in the universe—mental and physical—neither of which is reducible to the other. Monism, by contrast, takes one as primary and attempts to show how the other arises from it. There have been three major brands of monism: materialism (or realism), idealism, and neutral monism (emergence). The first argues that only the physical world exists. In modern times, this is the approach of science in general, and stimulus-response psychology in particular, according to which the mind is a reflex machine operating by association. Consciousness is held to be an "epiphenomenon" without causal power, when it is acknowledged at all.

Idealism takes the opposite stance: only the mental or "phenomenal" world exists. Matter is but an idea, a theory, construct, or illusion of the mind to help make sense of the flow of experience. Western idealism, through Kant, has been taken up in modern times by cognitive psychology. Eastern idealism asserts a common ground of both the inner and outer worlds. Both matter and mind are "illusory"—not that they do not exist but that one's unreflective relationship to them is a source of error and suffering to be transcended.

Theories of emergence often build on materialism. In one version, mind, like life itself, is described as a property of the organization of matter at a certain level of complexity. This modern view relates mind to concepts like order, self-organization, and information.

The universe as a complex system appears to manifest intelligent design. To understand Nature as something like the work of human hands must have been one of the earliest and most natural reflections. But to infer that such intelligent design is literally the work of a divine Artisan is to duplicate

the MBP on a cosmic scale. This makes God a ghost in his Creation, separate from it in just the way humans have distanced themselves from Nature. The human separation from Nature may be regarded as a symptom of mind-body dualism. Culture, while material, is a humanly-defined domain deriving from imagination and creativity, as distinct from matter and Nature. Our urban environments are literally mental constructs before they are physical constructions.

The notion of free will has traditionally depended on a concept of a non-material self above the deterministic laws of the universe, just as God is above them or outside them. The mechanistic view has gradually encroached on this notion, threatening to explain human behavior as a mere cog in the machinery of matter, governed as strictly by causal laws as every other part of the cosmos. This is reflected even in jurisprudence, when it is deemed unreasonable to hold a person responsible for actions determined by forces outside their conscious control—even if not outside their psyche. This aspect of the MBP becomes the search for meaningful concepts of selfhood, free will and responsibility.

The mind-body split has influenced, and been influenced by, physical conceptions of matter and space. Without the metaphor of the container, there would be no inner and outer worlds. Without the notion of mechanism, there would be no abstract machine for the ghost to haunt. But what exactly gives us the idea of an inner space? Where is it located? To be sure, the body is sensed as containing a certain volume, and this may be felt to be inside and to belong to oneself. But the body can also be experienced as external to the self when the self is considered a pure point of consciousness. We identify with sensations such as pains, which seem engulfing in their urgency. At times, however, somatic sensations seem rather a part of the space contained by the body, a part of the physical world, surrounding but separate from one's observing consciousness. The intimate or subjective senses are ambivalently experienced as qualities of the world or as bodily sensations. A certain flavor, for example, can seem to belong to whatever is tasted, but can also appear as a sensation in the mouth. This is manifestly *not* true of the distance senses. Except for anomalies recognized as perceptual artifacts, we experience seeing as taking place in the world outside the eyes; it is never experienced as a sensation on the body surface that the retina actually is. (Notwithstanding the painful sensitivity of the eyes to intense light, which is due to receptors connected with the diaphragm and not with the processing of a visual image.) Under particular conditions (for instance, with tinnitis, or with a loudness on the threshold of pain), sounds may be experienced as *sensations* in the ear, but normally they are perceived as *happenings* in the world, located in external space. The very idea of inside and outside originates in the difference between the normal functioning of near and the far senses; in particular, between vision and somatic sensation. Among the senses, we are highly identified with vision and generally convinced of the objective physical world it conveys. Through their ambivalence, on the other hand, the intimate and somatic senses at once introduce the idea of subjectivity and cast doubt, moreover, on the objectivity of the far senses. One could, however, turn the question the other way around and consider the puzzle (and the wonder!) to be the apparent externality and depth of space. How is it, in other words, that we see and hear things in the world rather than in our heads? How have we come to give such primacy to the visual sense, and founded upon it a concept of objectivity that excludes the body and the subject's input?

In the right conditions, the sense of touch can also be experienced as projected in space.[2] A tiny vibrator is applied to each of two fingertips on a person's hand. The rate of vibration of each stimulus can be varied independently, and the two can be coordinated so that the delay between them is slowly varied. When they are very much out of phase, the person feels separate sensations in the two fingertips. If the delay is reduced to a certain interval, the sensations fuse and are at first localized in the finger receiving the first impulse. But as the delay is further reduced, the feeling of the stimulus moves into and across the space between the fingers! Vision and hearing—the projective spacial senses par excellence—similarly involve the analysis of frequencies and the comparative inputs of two sources. Clearly, the sense of externality and space is constructed.

Since any mental event enters awareness in a sensory-based form (that is, as an image of the world), it is not surprising that imagination takes place in an interior space. "Thinking" is a catchword for a broad range of mental activities, of which deduction and induction have been abstracted as formal techniques. But even these are in certain respects spacialized—as, for example, in the mathematical idea

of vector space or the Venn diagrams of set theory. As Kant maintained, the experience of extension, or space, may be so primary that it is the condition for all other experience. Even pure consciousness, without contents, is sometimes described as an experience of empty space.

6. *Skepticism, or the Problem of Knowledge*

Through observation and reflection, the early Greeks developed an attitude of skepticism toward the reliability of knowledge, particularly that gained through the senses. Hence Plato's allegory of the Cave, demonstrating "the degrees in which our nature may be enlightened or unenlightened":

> Imagine the condition of men living in a sort of cavernous chamber underground, with an entrance open to the light and a long passage all down the cave. Here they have been from childhood, chained by the leg and also by the neck, so that they cannot move and can see only what is in front of them... At some distance higher up is the light of a fire burning behind them; and between the prisoners and the fire is a... parapet... like the screen at a puppet show, which hides the performers while they show their puppets over the top... Now behind this parapet imagine persons carrying along various artificial objects, including figures of men and animals... which project above the parapet... Prisoners so confined would have seen nothing of themselves or of one another, except the shadows thrown by the firelight on the wall of the cave facing them... And they would have seen as little of the objects carried past... Now if they could talk to one another, would they not suppose that their words referred only to those passing shadows that they saw?... In every way, then, such prisoners would recognize as reality nothing but the shadows of those artificial objects.[3]

In this thought experiment Plato casts doubt on the validity of the mind's sensory impressions, imprisoned as it seems to be in the body. It is not specific facts or elements of experience that are questioned, but the general validity of sensory experience and empirical knowledge. He raises the question of whether it is possible to know reality behind appearances—that is, the world-in-itself. His answer is yes: we can know the eternal "forms" directly through intuition. But Plato goes beyond mere skepticism (and Platonism!) to boldly speculate on the nature of perceptual processes. In the passage above, it seems to be his intention to model cognition and render an account of how the naive mind comes to believe in a false reality. But since Plato is talking about normal perception, the allegory is a psychological theory as well as an epistemological critique. It is a forerunner of Kant's idealism and the modern representation theory of cognition, and perhaps a distant reflection of Eastern idealism. Plato is saying in effect that there must be some process going on in the mind through which it represents to itself the objects of the outer world. He recognizes that we cannot know the world "directly" through the sensory channels—that is, without the creative intermediary of the mind. In contemporary terms, the brain has access only to the pattern of neural firing on the sensory surfaces—just as the cave dwellers are privy only to the play of shadows. The prisoner in the cave of the skull makes theoretical judgments about the world outside the cave of the skull, projecting the shapes it sees, in the shadow show of experience, as events in that world itself.

Plato had his own ideas about the nature of that external world. Notice that he specifically calls the objects casting shadows *artificial*. For him, not only are phenomenal appearances mental constructs, but the "real" things giving rise to them are also. Platonism, we shall see, has remained a strong undercurrent in the scientific study of Nature. For several reasons, people have long refused to credit Nature fully with its own independent reality, which is reserved to creations of the mind.

Skepticism has lost credibility in Western philosophy by harping on the doubtfulness of common sense, rather than following Plato's lead in modelling the process by which we know what we think we know. Thus, universal skepticism is often dismissed by modern philosophers as illogical or pointless:

> Just as 'there can be false coins only when there are coins made of the proper materials by the proper authorities', so... there can be times when our senses deceive us only if there are times when they do not.[4]

But there is a broader sense in which *all* coin is false, since it is merely token of something else. It is, so

to speak, not a question of genuine apples versus fake apples, but of the fruit versus its name or image. The value of coin is symbolic and conventional, unlike the usefulness of the goods and services it can buy. Experience too is a domain of meaning, in the language or coin of the senses. While it too can be counterfeit, it is naive to think that truthful reporting of the senses means that reality is transparently presented as through a window. Our window on the world is of a different and subtler nature. The fact that an object is registered in experience may reflect the *presence* of a real thing, but *how* it is experienced reflects also the nature of the window as well, which is more like the view onto a stage or the portal offered by a television screen. Objects, the principal characters in our reality play, are pragmatic fictions. The concept of object breaks down at the limits of the human scale.

Behind the shadow of appearances, is there, as Plato believed, a world knowable by some means more reliable, direct, or true than peering through the window of the senses? The metaphor of counterfeit money fails here because it requires a domain of "real" values, a windowless perspective on the world. But if perception is a projected shadow, as Plato would have it, it is not (as he also thought) the shadow of anything that can be accessed independently of conditional mental processes, unless we are to believe that the mind does not depend on the nervous system. Rather, we see in the shadows what we need to see in order to keep our hand in the game of existence. Skepticism, far from being fruitless, renders a valuable service in making us look before we leap into action based on too literal faith in our perceptions or beliefs.

Knowledge of the external world is always relative, conditional, contingent. The only absolutely certain knowledge is tautological: the deductive truths of logic. It is not knowledge of the external world, but of human definitions. Plato believed in knowledge of an absolute realm of ideal forms transcending Nature. This he posited to explain the frailties of the senses and subjective perception, yet to retain at the same time the possibility to achieve certainty and know reality as it truly is. We can have sympathy with his aim yet doubt his reasoning.

Plato's metaphor does attempt to grasp what it is to be in The Situation—in the position of embodied self-conscious mind. The Mind-Body Problem has boggled philosophers for centuries because it is impossible for the mind to be simultaneously subject and object, seer and seen, thinker and thought. We know very well how to describe the world of matter. But self-perception is contrary to the mind's ingrained mode of experiencing all as *world*, as external to the cogitating self. Whatever enters the field of experience is an object of consciousness, and never the subject. In tagging some of these contents of consciousness are "subjective," we recognize that they are somehow a property either of the physical organism or of its mental functioning, rather than of the world external to the body. Here, then, are three levels of distinction. First, there is the physical world inside and outside of the skin. Then there is the component of experience that seems attributable to the organization of one's mind, versus the aspect that seems to reflect objective reality. Thirdly, there is the absolute subject, which is interior to (or separate from) all objects of experience, physical or mental.

Cognitive psychology has taken up Plato's theme through numerous experiments demonstrating the mind's creative contribution to experience. In one situation, reminiscent of the Cave, human subjects watched a shadow cast on a translucent screen by a three-dimensional wire form. While stationary, the wire's projection as a shadow was perceived merely as a dark line on the flat screen. However, as soon as the wire was made to rotate, its true nature as a three-dimensional object came to life. In another experiment, kittens reared in a grey and featureless environment seemed almost unable to distinguish objects at all when released into normal surroundings. Plato also wondered what would become of human subjects rescued from his infernal cavern:

> Suppose one of them set free...What do you think he would say if someone told him that what he had formerly seen was meaningless illusion...? Suppose further that he were shown the various objects being carried by... Would he not be perplexed and believe the objects now shown him to be not so real as what he formerly saw?[5]

A middle-aged man who had been blind from an early age had his eyes restored through a corneal graft. Like the fugitive from the Cave, this man had to relearn to see. He could visually identify things familiar

to him by touch, but had difficulty naming objects not already known in this way. Someone brought him before a complicated piece of machinery. He was unable to say anything at all about it until he closed his eyes and began to explore the device with his hands. After a while he stood back, eyes open, and proclaimed that *now* he could see it, and proceeded to explain its various parts and their functions.[6]

The irony is that, in Plato's system, neither the appearances in the Cave nor the "various objects being carried by" are real in the modern sense of scientific materialism, which very much garners its knowledge of Nature in the manner of the above blind man—by actively and tangibly interacting with external objects.

7. *Other Minds, Other Bodies*

Some philosophers argue that it is logically possible (conceivable) that the various colors, experienced by a person with normal color vision, could be systematically interchanged in the experience of someone else in a behaviorally undetectable way. For example, what I see as red you could *call* red but *see* as green. Certainly, I would not point to the leaves of a tree in springtime and call them red when everyone else calls them green, unless I simply cannot distinguish red from green—which would be detectable behavior. Even then I would be familiar with other people's usage of color names and aware of my own disability. Of course, the argument is sometimes made that one could have all the same *names* for objects in the world, and be capable of all the same distinctions, and still *experience* them differently. This also suggests the idea of "me" somehow experiencing "your" experience. Apart from the notion of my brain being connected to your eyes, I do not see much possible meaning for this idea. That it appears meaningful to some probably arises from their reifying experience as a substance or realm existing objectively and independently of subjects—to which various observers could have common access, as they do to the physical world. Experience is not a thing, however, but more like a language.

Pain is often considered paradigmatically subjective and private. I cannot directly experience your pain—obviously because my brain is *not* hooked up to your body. It is not that information about your body's tissue damage is unavailable to me in principle. The deeper issue is that I am not motivated to respond to your condition in the way I would be if it were my own. This is no question of privacy, or of an inalienable realm of one's "own" experience as privileged knowledge. Your tissues and mine are alike public domain, either visible or subject to probing with technology. (One day, in the unenviable future, our "brain states" may be public domain as well.) The *relationship* between mind and body differs, however, in first and third person experience: *this* brain is concerned for *this* body in a way that it is not for other bodies. It is a question of values (intentions) rather than information. Suppose my brain *could* be hooked up to your body. In that case it seems certain "I" *would* feel "your" pain as my own. If both our brains could be hooked up to both our bodies, allowing a simultaneous experience of the condition of each, whose pain would it be then? Indeed, whose body?

From a commonsense perspective, we all experience the same world, from differing vantages in space and time and in our own ways. If all minds (and brains) were qualitatively identical, and only numerically different, then we would all perceive the world identically, except for spacial-temporal perspective. There could be disagreements about the; owing to a different space-time locus, for example, I could be experiencing pain while you are not. My foot could have been run over by a passing vehicle while you, standing a few feet away, went unscathed. This does not make the event of the injury a private experience. My pain (in the foot) appears to be my private experience because mine is the only brain connected directly with nerves to that foot. My pain is an awareness of something physical and quite public—namely, damage to this body. Damage to *another* body will not be experienced by me as physical pain, but it *will* be experienced by me in some manner if I come to know of it of it. Similarly, I could be seeing the blue of the sky while you (on the other side of the world) are seeing the dark of night. This does not make vision a private experience for either of us. If my eyes, halfway around the earth from you, could be miraculously connected to your brain, then you too would see the blue sky where my body is. Similarly, if nerves leading to your pain centers were connected to receptors in my wounded

foot, you would experience "my" pain. In that case, what could 'privacy' mean? Is it really a private screening of *blueness* we are separately experiencing or is it the common objective *sky*? Are we having separate but equal pains, or is it the objective damage to this foot we both feel (and see as well)? What makes experience personal is that it is *for* and *by* this body. This is because the connections—the pathways of information processing, both intentional and causal—lie within this organism. By analogy, it is grammatically permissible to say that my knowledge of physics is for my personal use and information, as yours is for you. While intersubjectively available knowledge can be used for private purposes, one could not conclude, however, that this knowledge is intrinsically private, in the sense that it is not publicly accessible. Let alone can it mean that there is no public science of physics, or no intersubjective world it is about.

In regard to both the inverted spectrum argument and the supposed privacy of pain, we could also consider scenarios where my brain is hooked up to inclusively more of your nervous system than just the pain receptors or the retina. What if we swapped brains or bodies? Or, what if my brain was hooked up to your brain in tandem? What could that look like from the outside (let alone the inside)? What could it possibly mean for two brains to be integrated? Where would the joining nerves run exactly, given that brains are massively decentralized? As mind-bending as these thought experiments might be, the real difficulty posed by the apparent privacy of experience is not that it is a domain of privileged information, owing to hard connections between brain and body. Information in general is rather a public domain, wherever it is shared. Perhaps what actually makes the sense of privacy seem inviolable is the possibility of dissimulation that inheres in subjectivity, as well as in the ambiguities of language, so that the other may be kept guessing as to one's inner state. Deception, hiding, bluffing, managing the other's image of oneself, are subjective stratagems; but even these are no ultimate weapons of privacy. Knowledge of other minds, like other kinds of knowledge, is a question of adequate models and of sensitivity to the information available. In poker, the cards are hidden, but not the faces and actions of the other players.

The inside and outside perspectives are different modes of registering the same facts. The neurophysiologist's determination that my nerves and brain respond in certain patterns to light of a particular wavelength contsains the *same information* as my own determination that I am seeing *blue*. But a further difficulty arises because of the mind's imprisonment within its own systems of meaning. It is hard, for instance, to assess the words and actions of other people with whom we disagree, because we can never be certain to what extent the situation is really as we see and feel it to be, or to what extent we are coloring it with our own biases or filtering it through our own interests and defenses. Each person may tell an amazingly different story based on what are presumably the same objective events. As subjective beings, we know that the mind is "prejudiced"—to use Descartes' old fashioned term—a system operating on particular assumptions. Other subjective minds have other ideas, for which they claim truth as vigorously as we. And yet we know (as presumably they know) that another viewpoint is always possible and that, by invalidating the other's perspective, either party could be missing something vital to their own perception of the truth. In this very personal (but not private) sense, the problem of subjectivity is to distinguish what is of oneself (or of other selves) from what is of the world. This project is begun in early childhood, and never really completed. It is the task of sorting out what we impose upon (or exclude from) experience from what is imposed on it (or excluded from it) by others and by reality at large.

8. *The Reality Principle*

As a concept and a category that is a human construction, the Real refers to serious, survival-oriented cognitive behavior, as well, of course, as to the reality of the world. It applies to the behavior of biological organisms in general. The fact that we experience the world as real, concrete, independent and external is a sign of our sensitive dependence as organisms on a crucial environment. This is not to deny, but to affirm, the reality of that environment, which the category of the Real necessarily reflects.

The dependency on an environment places the organism in the logical relation of a finite

axiomatic system to a more inclusive system (environment). The former derives its semantic meanings from its context within the latter. As a dissipative system, the organism depends on an external energy source; as a biological system it is part of and dependent upon an ecology external to it, which shapes its nature; and as an "interpreted" formal system, modifications to its rules and structures must originate from a meta level.[7] Meaning for the organism, and its sense of reality, therefore derive from outside itself.

The external world is real to the degree that it as power to determine experience. Were we disembodied minds, not bound to the laws of organic life and the physical universe, we would be free to create experience in any arbitrary way. We would not only be free to follow whim, but *obliged* to, since it is the game of physical existence which provides the actual rationales for meaningful activity and the categories of experience. Even dreams are rooted in reality, their images drawn from waking experience, their messages speaking to us of our life in the real world. Without a real world to respond to, the mind would have to create an imaginary one in which to move—as happens in situations of sensory deprivation.

On the other hand, the human organism is not as straightforwardly tied to reality as simpler creatures appear to be. We have the remarkable ability to see our own programming, and thus transcend it, often leading at once to a looser and yet more objective relationship to the world. In the same freedom, we create imaginary or ideal worlds as we please, as alternatives to the natural one, laying claim to the right to make gratuitous choices. It is in this sense that "pleasure" opposes itself to "reality," rather than in the terms proposed by Freud:

> Just as the pleasure-ego can do nothing but wish, work for a yield of pleasure, and avoid unpleasure, so the reality-ego need do nothing but strive for what is useful and guard itself against damage.[8]

Freud here confuses the gratuity of imagination—wishing—with the search for pleasure, which he then opposes to the survival value of the useful. But pleasurableness has an ambivalent nature. On the one hand it is a quality attaching to objects or events; on the other, it is a cognitive judgment concerning them made by the organism with reference to their utility toward survival. The apparent opposition between a pleasure principle and a reality principle, which plays such a large role in Freud's thinking, dissolves and shifts ground if we grant that both pleasing/unpleasing and useful/unuseful are *survival-oriented* discriminations, albeit on differing levels of mental organization. To say that we seek pleasure is to say no more than that we seek, for good reasons, what we have been genetically conditioned to prefer. In other words, the organism does not primarily seek pleasure per se, but (generally) well-being. Pleasure is the built-in awareness of what is good for it (or perhaps, more precisely, good for the genes at the core of it). The subjectified notion of pleasure, as an experience valuable in its own right, is a product of modern subjectivity; it has come to seem like a kind of thing or consumable substance, rather than a state of well-being reflecting an environment. Sometimes the subjective experience of pleasure even contradicts well-being. Laboratory rats (which presumably do not have subjective consciousness) may seem to seek pleasure itself in stimulation, rather than some good which appears (to a human observer) objectively beneficial. But this appearance exists in the eye of the beholder, while the animal has simply judged the stimulus 'good', if mistakenly. Human beings, however, may indeed seek pleasure for its own sake, whether or not it is good for them, since they have abstracted it from "goodness." The pursuit of pleasure, in Freud's scheme, turns out to be an expression of the Reality Principle, however distorted by human subjectivity.

Furthermore, a reality principle implies more than utilitarian preoccupations and goal-orientedness, but encompasses the whole outward-leaning bent of the mind. This is the mind's tendency to project its own cognitive processes as an independent external world. Ideally, when this projective tendency is epistemically justifiable, it is the search for objectivity and truth. The many forms of realism and reification are instances of this broad principle. Rather than Freud's antinomy, let us contrast a reality principle with one which emphasizes the "self" component of experience—a principle of gratuity or subjective freedom.

Apart from pure interest in truth—which is disinterest—the concern for the Real is a concern for the future, the spacially and temporally distant, the non-actual. Experience of the real world is largely a construction of objects in visual space. Cognitively, the mind imposes upon the chaos of sensory input various patterns it has learned to anticipate. One worries less over the actual than over the imagined. Paradoxically, the whole purpose of the Real is to provide a controlling bulwark against non-actual contingencies. One is therefore caught between two modes: embracing the actual and monitoring the potential. The first implies engagement in the here and now, gratuitous play in the shadow of annihilation. The other implies goal-orientation, disappointment lurking behind the promise of mastery and well-being. Success is never guaranteed in the quest for mastery, which has a way of generating further problems. On the other hand, well-being is hardly assured simply by abdicating concern, by sticking one's head in the sand. The ego lives between a rock and a hard place, unable to achieve security yet bound by its mandate to try.

9. *Self-reference and Subjectivity*

The board-game "reality," in which the only things that exist are those defined for the game and allowed by its rules, is like a complete formal system. All possible outcomes derive purely and simply from the rules and initial starting positions. What could be experienced within a "world" defined by this system is coextensive with the game-world it defines. The question of a real world beyond its limit is undecidable.

Such a self-contained game-world is closed upon itself; the "point of view" of a playing piece cannot be transcended from within the game. The meta-viewpoint of a player, in contrast to that of the playing piece, however, is able to represent the game in the larger context of the real world the player occupies. This is not simply a matter of physical perspective, but of a more inclusive set of definitions, of an amplified system, embracing a larger and more complex reality within which the game-world is situated. Not so for the playing piece, whose reality is limited to the world of the game. We could likewise imagine entering the "world" defined by a piece of machinery.

Pre-subjective consciousness might be viewed as analogous to a complete formal system. The pre-subjective world is fixed and final and, like the board game, contains no more than what is prescribed by given rules, of which there can be no awareness as such. Only the world exists and the unquestioned truths that seem to be part of it. Nothing beyond this world *can* exist, since the system lacks the ability to self-refer or extend its own definitions. There is no subjective self to harbor doubt or side-step tradition. In contrast, consider a system capable of self-reference—a game in which the game itself is a defined element. Subjective consciousness is analogous to such a system; it constructs a meta-perspective or meta-language in a system that self-refers, pushing the mind over a threshold of complexity into Gödellian "essential incompleteness." In theory, self-conscious mind can transcend any model, value system, or world.

Gödel's proof regarding formal systems is a formal counterpart of subjective consciousness. It works by representing within itself statements about itself, just as subjective consciousness contains an image of itself. The "theorem" of subjective consciousness is limitative in the way that Gödel's is: there are thinkable thoughts, and haveable experiences, the truth or reality of which are undecidable in any particular mind or state of mind. There is always more than meets a given mind's eye. There is always the Unknown.

Subjective consciousness is a reserve against the hubris of thought, a foil for the reifying tendency of mind. But it also plays a positive role. The world can only be objectively appreciated (realized) by a mind that is conscious of its own role in creating experience and knowledge of it, just as truth can only be distinguished from provability in a system with essential incompleteness.

For the naively realist pre-subjective mind, the reality of the world must be self-evident, unquestionable, given once for all. The world coincides with the horizon of consciousness. For the subjectively conscious mind, by contrast, the identifying quality of reality is that it is always vaster than

what can be captured in thought or experience; some part of it always will lie over the horizon. Yet this fact remains indistinguishable from the effect of the self-reference itself. Subjective ideas and perceptions can always transcend themselves just because they self-refer; it is a self-expanding system. Some part of the subjective world always also lies beyond a given horizon. How then can one say for certain that this transcendental quality of reality is any more than some sort of projection of the mind's self-transcending capabilities? The transcendent realness of reality, after all, consists in the fact that it is not an artifact of mind, nor capturable by mind. But if "transcendence" were a quality that mind itself generates, then how could it be relied upon as the hallmark of reality?

Three distinct cases seem relevant to this conundrum: (a) The universe is finitely large and finitely detailed. Its complexity can be exhausted in human descriptions; eventually we will come to know everything. (b) The universe is infinitely large or infinitely detailed and complex, or both. No finite system of thought can encompass its totality, and reality will always remain elusive. (c) Quite independent of the nature of the world-in-itself, the nature of the mind as an open system implies that understanding of the cosmos can never be complete. Cognition, involving both subject and object, is troubled by the equivalent of Gödel's essential incompleteness, so that even if (a) is true we will nevertheless always surprise ourselves and knowledge will always be unfinished. And, if (c) is true, how then to decide between (a) and (b)? The apparent *depth* of reality could as well be a product of our perception—that is, of our own depth as self-conscious beings.

While the Mind-Body Problem is a dilemma of self-reference, for pre-subjective consciousness there is no dilemma, since there is only one category of existence: the world. Classical physics, in keeping with the premises of scientific method, eliminates the subject from its discourse. For it, as for the pre-subjective mind, only the (physical) world exists. Even when the processes under study are mental processes, they must be treated strictly as events in the physical world. In this way the domain described, like a complete formal system, is sealed off from a potentially larger domain of description. Science, as an expression of the Reality Principle, simply skirts the paradoxes of self-reference by limiting its domain. Classical physics is a formal counterpart of the pre-subjective (naively realist) attitude, while the scientific revolutions of the early part of this century mirror the rise of subjective consciousness.

10. *Paradox and the Subjective Frame*

Epimenides was the Cretan who infamously declared that "all Cretans are liars." The modern form of this paradox is the sentence:

(1) 'This statement is false.'

The sentence refers to itself and also negates itself. If it is true, then it must be false. If it is false, then what it asserts is true. (Notice that, if it read: 'This statement is true', there would be no difficulty, since the statement simply affirms itself without contradiction.)

Now a statement, P, is in effect the assertion that 'P is true'. This constitutes another statement, P_1, which in turn is equivalent to the assertion that 'P_1 is true'. Since this process can be carried on indefinitely, any statement is equivalent to an infinite recursion of statements about statements about statements... Thus:

'This statement is false' = "['This statement is false'] is true."

For assertions that do not self-refer, there is no contradiction:

"['Statement X is false'] is true" = "["['Statement X is false'] is true"] is true"... etc.

Whether X is true or false, every recursive version of it is also true or false, accordingly. But for the self-

referring statement, the situation is different when the statement is self-negating as well. Each recursive step reverses our verdict concerning its truth. By contrast, consider a statement such as:

(2) 'This statement cannot be proven in system X.'

Statement (2) is the "formal" or relativized equivalent of (1), since it is framed in terms of proof-within-a-system rather than in terms of truth. It contains no contradiction so long as the notion of contradiction has an "interpreted" meaning as "both true and false in some context larger than system X." Thus:

'This statement cannot be proven in system X' = "['This statement cannot be proven in system X'] is true,"

which is no contradiction. But when the concept of contradiction also assumes a relativized definition (provability instead of truth), then (2) *is* self-contradictory because it *and* its negation are both provable (that is, derivable from the axioms of system X). Substituting "provability-in-X" for "truth", (2) becomes:

"['This statement cannot be proven in system X'] can be proven in system X."

And substituting 'this statement' for what it refers to, (2) then becomes

"['This statement'] can be proven in system X."

Therefore if (2) can be proven in system X, then so can its negation, and so (2) *is* a contradiction. In other words, a statement like (2) is not a contradiction so long as there is a method of evaluating truth that lies outside the system. But it is a contradiction when it must be evaluated strictly from within X—in terms of provability rather than truth.[9] We could generalize to say that contradiction or paradox can arise in subjective consciousness, as it can in self-referring statements and logical systems, when one is limited to a closed framework of thought. To transcend the terms of the paradox, it is necessary to expand into a larger system in which the contradictory theorems can be evaluated as "truths." The possibility of self-reference generates paradox in the first place, and also provides a way out.

11. *The Ground of Experience*

The classical distinction in philosophy between intensional and extensional statements bears examination in the light of the previous section—and also, we shall see, the following one. 'Extensional' refers to statements P about the world, or to propositions asserting the truth of such statements: 'P is true'. 'Intensional' refers rather to statements of the form 'I propose that P is true'. The latter is an explicitly subjectified version of the former, in which the subject is implicit. To the pre-subjective mind *all* propositions are extensional. To the subjective consciousness, on the contrary, all propositions are implicitly intensional, though they may or may not be formally recognized as assertions of a particular mind. In other words, all facts may be questioned as beliefs. There are no "free-floating" truths, independent of minds.

The extensional statement P is the formal equivalent of the awareness, experience, or belief *that* P is the case. Furthermore, the notion that a proposition is either true or false (the Law of Excluded Middle) expresses a fundamental fact of experience in the world. That P is either true or false generalizes and abstracts the apparent truth that any object either exists or doesn't at a particular place and time.

The *experience* itself of certainty does not add to the validity of one's knowledge or the legitimacy of one's faith in a proposition—any more than one logically adds to the reliability of a statement by insisting it is true. (The exception is the legal oath, since breach of it constitutes a separately punishable offense.) On the contrary, such insistence can raise eyebrows in subjective culture, where the

everpresent possibility of falsehood requires skepticism in regard to all claims. I believe these considerations must apply also to Cartesian certainty, which is based on the supposedly irrefutable nature of private experience.

The bracketing of "experience" in subjective consciousness amounts to the proposition, 'this object is not a real object', where 'real' refers to some element in the experiential field that is normally identified as part of the world. In other words, it is realized that this element of experience belongs not to the world, after all, but to the self. While we subjective moderns find no contradiction in this, we can imagine that the pre-subjective mind might have found it deeply disturbing. In fact, we might speculate that it was exactly this discovery—of things in the phenomenal world that are not things in the external the world—which led to the expansion of the human cognitive domain to include the subjective and the concept of 'experience.'

But when 'this object' refers to something already recognized as belonging in the subjective domain, the proposition becomes paradoxical even for us. For it then becomes equivalent to 'this experience is not a real object of experience' or 'this thought is a real thought'. The apparent absurdity of such statements led Descartes to found his philosophy on the seemingly irreducible primacy of experience: whatever its interpretation, there can be no doubting the fact of having an experience. Must we not admit, however, that the numinous self-evidence of *experience* for the subjective mind could be as illusory as we imagine that of *the world* to be for pre-subjective thought?

'This object is not an object' and 'this experience is not an experience' share with 'this statement is false' the condition of lacking a domain, meta to them, in which they can be judged as non-self-negating propositions of the form 'this statement is not provable in system X'. We accept the existence of objects-which-are-not-real precisely because we have the meta-domain of subjective experience in which to judge the reality of phenomenal objects, just as we may accept self-negating statements as non-paradoxical when they can be reformulated in terms of provability. Of the configurations of sensory experience, not all constitute external objects. The fact that there is a remainder is the basis of subjective consciousness. But is there a meta-meta-domain within which to decide the status of an experience—not whether it is veridical, but whether it is really experience at all? While at first sight this question may seem nonsensical (as Descartes would insist), perhaps the role and significance of conscious experience can be fully appreciated only when the domain from which it arises is brought into clearer definition.

We have legitimized experience as a category yielding the very basis of a sense of self. One *has* experience, therefore one must exist. This is the meaning of Descartes' *cogito*. To be sure, the self remains unaffirmed by unconsciousness or non-experience. But is it really affirmed by conscious experience? Philosophically we have been led, by objects of experience that are not objects in the world, to question the very reality of the world. In a similar manner, might we be led to question the validity of the subjective domain as the repository of a self?

Descartes would argue that, even if a given experience does not prove the existence of a real object, it *does* prove the existence of a subject. But does that actually follow? The indisputable part of his assertion is this: while the *meaning* of an experience may prove elusive, the *occurrence* of the experience is unquestionable. However, Descartes takes the further leap of concluding that which he assumes in the first place: if experience occurs, then there must be a subject for that experience, an experiencer. Indeed, if we have *defined* experience as requiring a subject, then we are simply caught in circular reasoning.

Descartes' syllogism is popularly translated "I think therefore I am"—which presupposes an agent ("I") for "thinking." It ought to read instead: "There are experiences, therefore I am." In an even purer form it could read: "The phenomenal world is inconsistent (giving rise to the concept of experience), therefore *I* exists." Either way, it is difficult to see how the conclusion follows from the premise. Rather, to infer the existence of a subject from the occurrence of experience seems the same sort of mental act as to posit the existence of external objects on the basis of particular experience. In other words, the existence of the self parallels that of the world, in that both are logical inferences and constructions. If objects are philosophically questionable, then so are subjects.

Descartes' reasoning may be illogical, but as a psychological process it may well describe how *I*

concludes its own existence. And if *I* believes itself to be a subject for all experiences (in this field of consciousness), and an agent for all actions (of this body), then any such experience or action will be further proof to *I* of its own existence. This reasoning is on the same shaky footing as the claim that every morning I make the sun rise, and every sunrise proves my omnipotence. According to Descartes, *I* necessarily exists even when the content of experience does not veridically refer to the real world. But if one experiences a dream or hallucination, how is this confirmation of a real subject experiencing it?

The ego normally arises as a kind of personal theory about the coherence of experience, an explanation for the consistent association of certain contents of consciousness. For example, the proprioceptive and tactile sensations of the hand in picking up the cup of coffee, the sensations of taste in the mouth and of hot liquid swallowed in the throat, as well as the somatic effects of caffeine to follow, all constellate with visual sensations corresponding to hand, cup, etc., in an emerging sense of self or agency. *I*, as the passive witness and motor force behind this body, is how the brain makes sense of the association of these diverse inputs, including perceived outputs. The sense of I-ness may be a wonderfully pragmatic device, which mostly works well in the evolutionary scheme. But is it true? Does the self "really" exist?

There *are* people who attribute great powers to the self, believing they make the sun rise or control the weather, believing themselves the focus or the orchestrator of great cosmic happenings. But some mental patients, and some mystics, completely lose the sense of self, facing a bewildering confusion of sensation with no unifying scheme to make sense of it. Most people's experience, of course, lies safely between these extremes. But normality does not establish truth, nor make the self real. From the perspective of every spiritual tradition, entered deeply enough, the place and very existence of self is problematic and questionable.

12. *Epistemic Responsibility*

As noted in the previous section, 'extensional' refers to statements P that are made directly about the world, or to propositions asserting the truth of such statements. 'Intensional,' on the other hand, refers to statements of the form 'I propose that P is true.' In the latter case, the subject takes explicit responsibility for his or her assertions. The former expresses a third-person relationship to experience and the latter a first-person relationship.

The Mind-Body Problem is the gulf between these relationships. Because we are self-conscious beings, we can see that *all* perspectives are fundamentally in the first person, from the point of view of a subject. All experience is *someone's* experience and all propositions must be someone's assertions. Every view is a view from somewhere. But, of course, perception does not always occur in a self-conscious state. By habit we simply perceive *the world*, and are more or less unaware in the moment of the act of doing so. Thus, it appears that the world simply *is*, and our claims about it may reflect this fact. The world may seem to us as *we* perceive it, and that it would be so even if *no one* were present to perceive it. Thus, also, there appear to be two versions of reality or two domains of being: one's subjective experience (of the world) and the world itself. The first we have identified with mind, the second with material bodies, or external reality. But the distinction is generated by the fact that we sometimes do, and sometimes don't, claim responsibility for our first-person point of view.

This responsibility means, to begin with, realizing that the sequence of appearances that one takes to be the world unfolds in a center of consciousness. This is "I". There is the realization, moreover, that this center has some control not only over the appearances themselves but also over one's actions in response to them. Secondly, it means realizing that other bodies in the world may also be centers of consciousness, with similar responsibilities: other "I"s. Commonsense and legal notions of responsibility and personhood are intimately linked. To have the status of person means to have the responsibility and rights of citizenship. Lose one, and you lose the other.

A human being is held to be a person, idealized as the originator of his or her actions; and is also held to be a physical body, whose actions arise causally within the deterministic system of which it is a

part. Thus, a human being is deemed to operate both with free will *and* under the yoke of determinism. Society demands of *persons* accountability for their deeds, while their *bodies* may be dismissed (and forgiven) as little more than programmed machines. In both cases, people are idealized: as subjects *and* as objects. Their behavior is deemed both intended and causally determined, though perhaps not in equal measure in all instances. Both compassion and scientific understanding may serve as bridges between these disparate perspectives; but they can be abused as well, as when lawyers increasingly gain acquittals for their defendants on the grounds that criminal behavior was somehow caused by extraneous factors, not intended. Environmental forces outside one's moral responsibility have come to include everything from alcohol and drugs to a surfeit of sugar or a deficit of love—ultimately, perhaps, to bad genes! Excuses for individual lapses of intention give way to a general dismissal of intentionality—and therefore of accountability. Petty criminals can get off with a note from their psychiatrist, while corporations legally steal fortunes with a note from their accountant. Dualism is a political issue.

Intentionality is the key to any reasonable solution to the Mind-Body Problem—that is, one that does not take a dogmatically idealistic or materialistic stance, but which respects the indissoluble joint participation of self and world in determining experience and knowledge. Thus, descriptions of natural reality and questions about it should be acknowledged as posited by mind, while answers to such questions, and justification for them, must be provided by Nature itself. Human intentionality is the power behind intensional assertions, while Nature is presumed to be the power behind extensional facts. Since "facts," however, are inevitably human assertions, we come back to the co-participation of self and world in forming experience. These, then, must also be joint authors of assertions made on the basis of experience. To put this in philosophical jargon: we must be ontological realists while remaining epistemological idealists.

How might the peculiarities of language affect thought and perception of reality? Professor emeritus Leslie Dewart, of the University of Toronto, criticizes Chomskian linguistics for assuming that all languages possess a deep structure patterned on the Indo-European model—namely, that verbal expressions are composed of a verb phrase predicated of a noun phrase. In particular, the verb *to be* is a culprit, by ascribing to objects themselves (i.e., nouns) qualities (adjectives) that are in fact the speaker's assertions. Thus, to simply state that 'Jimmy *is* tall' is to make tallness an inherent property of his being, without reference to the speaker's perception or judgment, whereas Jimmy's size is actually a relative matter, involving at the very least the observer (speaker) and some comparative points of reference. According to Dewart, many non-Indo-European languages lack such a verb and are relatively free from this error.

Dewart begins by proposing that reality, for the modern Westerner, consists of the "other-than-self." (This, of course, makes it implicitly relative to a self.) As a *quality* added to experience by the organism, I have contended that realness implies a relationship between the organism and its environment that is established through an evolutionary history. Dewart puts it all the other way around: for the "ontic" mind, the realness of the other-than-self is added to experience not because of the crucial significance of the environment, but as a result of a "deficiency in self-presence," at least among some linguistic heritages. The reality of the self for itself, in other words, *ought* to be primary, that of the world, secondary; but for the Western mind this is reversed. Whether or not one agrees with this idealist perspective, there is something sound about the idea that "self-presence" arises with the assertiveness of speech. However, self-presence is the key not to simple awareness, but rather to subjective consciousness. It is this, I believe, that is also the key to epistemic responsibility.

In Dewart's terms, users of Indo-European languages do not easily grasp that the "otherness" (realness) of the external world is a *relationship*, and that the object can exist only in relation to a subject. They may tend to ignore, and so disown, participation in the thought processes leading to their statements about the world; and in that respect they resemble naive realists. They will simply point to their experience (of the world) as justification for their extensional statements, seeing little need to qualify the latter as reflecting *their* experience. He goes on to elucidate three consequences of this epistemic attitude. The first is that "substantiality becomes the paradigm of reality." In other words, only "substance" is real. Realness inheres in things themselves, not in the relationship between them or with the subject. The

second is that such a mind is alienated from itself and doubts its own reality as responsible agent. It will tend to view itself, if at all, as an ethereal substance (soul), because only substance is real. The third is that it will doubt the contents of its experience as well, since these are seen, in the subjective frame, to be insubstantial. It will instead be tempted to credit some realm beyond experience with ultimate (subtly substantial) reality.

Certainly, classical physics can be viewed as guilty of the first and third of Dewart's indictments. Mind-body dualism in general is guilty of the second point, and Cartesian skepticism and the abstractions of Hindu philosophy, of the third.

Dewart criticizes the divine copyright on reality, evident in the Semitic monotheism from which both Christianity and Islam are descended, as a sign of deficient thinking. Although the Israelites and other Semites were not of Indo-European stock, they comprised another group of warrior-herders who began to disrupt agricultural civilization in the neolithic, turning it toward patriarchy and mayhem. According to Dewart, the 'I am that I am' of Yahweh should be interpreted to mean that the very realness of things is God's doing and not an inherent property of the world; *everything*, including Nature and ones own private experience, belongs to God and only exists or happens by divine grace and decree. This is no less the understanding of Islam and Christianity. For a deficient mentality, Nature and the self can have only a deficient reality!

The key distinction of Semitic religions may involve this issue of fatalism. In an impersonal system (such as we also now have resorted to in mechanism) determinism is inescapable, automatic, impersonal. In contrast, the God of Moses and Muhammad is a personal, if invisible god—a father figure. And this may be the significance, psychoanalytically speaking, of the ascendancy of patriarchal over matriarchal religions: the father figure is more personal, because he comes into the child's awareness at a later stage of development, whereas the mother remains (even as the goddess) a less personal force, as she was during the child's infancy. In a sense, then, science reverts to a pre-patriarchal world-view in honoring the independent reality of Nature to the extent that it does, even though it was the child of idealistic medieval Christianity. The contradiction is still with us in the form of the debate between science and creationist fundamentalism, which holds (as the first scientists also did) that Nature is a divine artifact whose reality depends entirely upon its Creator. It is present also in the Platonism of some scientists, whereby Nature is assumed to be a deductive system, perhaps even a computer program.

In many non-Indo-European cultures, the natural world is thought to be the only real one and a seamless whole. Man is an integral part of this whole, which is his "true and only home." This belief is reflected in the absence of a platform for deliverance from mortality, in such cultures, or the goal of ascent to a "higher" reality. If some cultures are able to affirm that the earth is where mankind belongs, perhaps it is because their religions are not preoccupied with a concept of "salvation" or another world where their true destiny unfolds.

If only the "other than self" is real, then how can the self be real? One cannot then be satisfied with a sense of one's own presence, but is driven to imagine the reality of one's self as an *it*—for example, as a soul. The doubt of one's experience and of one's very being can certainly lead to doubt about the value of others' experience and being, as well as the being of Nature. The reduction of Nature to human concept, and the denial of its immanent reality, are ingrained in science. The reductionism of science stems from the propositional nature of thought or, as Dewart would say, from thematic language, which reduces experience to a concept or name. One evolutionary purpose of subjectivity, I would hold, is to question the mind's categorizations and judgments, allowing thought to step out of such self-imposed boxes.

Dewart argues that the characteristics of the "ontic" mind that have allowed it to dominate human civilization, nature, and (we must add) women, are of short-term value only. This begs the question of its ultimate value, and no more than restates his underlying position that the modern Western mentality is essentially defective, whatever its short-term advantages. He argues further that such a mentality is motivated *not* to take moral responsibility, either for its perceptions or its acts, but to overvalue external reality. It is this very overvaluation of external reality that has led to technological, economic, and reproductive success. If Man were morally autonomous or inner-directed, and could truly be the

originator of either good or evil, it would mean that good and evil were not objectively real but mere human whims. This the ontic mind cannot bear, with its outward focus that requires morality to be a matter of *obedience* rather than *initiative*. One sees this in the Christian dogma of salvation by grace rather than good works.

Edmund Wilson had a different problem in mind when he complained, concerning mankind's place in the amoral and directionless scheme of evolution, that in the absence of religious conviction we have no goals and no moral guidance but the outdated axioms of biological nature. But this is simply untrue. Mankind has *always* demonstrated the goal to remove from nature: to transcend all limitations, to be undetermined by nature's purposes and free to set its own goals—even when, ironically, these were projected as divine will. The decline of religion has simply invited political and economic ideologies to take its place as the great galvanizers of belief and action. Even the ontic mind of the Indo-European cultures eventually grasped its existential freedom; I do not think anyone seriously believes we can imagine only the premises that evolution built into us. The real question is how strongly we desire to live *intentionally,* with conscious free will, and how much we would prefer rather the convenience to sit back and allow outside forces to drive us—whether genes or gods, despots or destiny, corporations or climate change. Human destiny can be informed by biology and evolutionary history, but not fixed or defined by it. To expect the future to be determined by the past is merely self-fulfilling.

It may be true that what we currently think of as human nature—our genetic heritage—was largely formed long ago, in an Ice-Age environment. But the essence of the human spirit, which is self-consciousness, is of another order than evolutionary programming, and competes with it as well as complements it. It *is* a serious problem that the modern human environment bears little resemblance to the formative ancestral one, so that our "innate values" are mismatched to contemporary life; such values can lead us, in the modern context, to a far more dangerous tribalism and the destruction of the "cosmopolitan" values of a united humanity. Or they could lead us to a renewed localism and more humane scale of civilization. However, it is not true that we have only archaic values upon which to draw in order to choose consciously which values to hold presently. The fact that values and human nature itself may be consciously created is what makes us human. We are the creature that makes itself.

Part Two: CAUSE AND INTENTION

13. *Causality*

The term *extension* refers in logic to the set of objects that share some defining characteristics. It is distinguished from the complementary term *intension*, which is the set of defining characteristics those objects have in common. We saw in the previous two sections that extension refers to objects of thought or consciousness, including physical objects, and in particular to the extension in space of such objects. Intension refers rather to properties a thing must possess so that a particular term can be assigned to it. This implies an agent assigning such properties. (In logic, the related term *intention* has a similar technical meaning: an instrument, such as a concept, for knowing and referring to a thing as it exists in the mind.) We will use *intention* here in a broader sense that links intension to extension, organism to environment, mind to matter or body.

Objects and events in the environment have extension in space, and also in the logical sense of being the recipients of qualities assigned to them by an agent for whom they *mean* something. They have significance for that agent's well-being and survival, which is recognized and acted upon. Meaning cannot be accounted for by physical causes alone. It needs another form of description, in terms of logical relationships and connections, in addition to physical ones. Cause can account for the interactions

within a system extended in space (all from the point of view of an outside observer, whose agency is ignored). Only the *intentionality* of an agent, however, can confer meaning or explain what goes on within the observer as subject or agent. Intentionality is thus the power of such a system to relate to the extensional world in terms that are meaningful to itself. The difference between extension and intension (or intention, as it will be more broadly used here) is the difference of perspectives that is the essence of the Mind-Body Problem. One is the perspective *upon* objects and events—in physical space—that is presented in phenomenal experience. The other is the perspective *of* the cognitive agent in and from which phenomenal experience presents itself. Operations within a given brain can be described extensionally, since they take place in physical space and time; they can be explained as a causal sequence leading to behavior. But the *experience* of the subject and the willed actions of an agent can only be understood in terms of logical sequences the subject imbues with meaning.

Aristotle distinguished four types of cause, each serving a complementary aspect of what explanation meant to the Greeks. The *material* cause of a change of state was the presence of a medium in which the change takes place. An *efficient* cause was some agent or preceding event that brings about the change, whose *formal* cause was its tendency or end state. Its *final* cause was its purpose, use, or reason for being. The first three types of cause are extensional; they describe physical operations in space and time. But the notion of final cause implies logical operations and values; these involve intention.

According to Berkeley, Hume, and some later philosophers, causality amounts to nothing more than a succession in time. One event is regularly seen to follow another, and from this we understand that one is the cause of the other. But cause intuitively seems to involve a power of one thing over another as well as a mere succession of events. Hume criticizes this idea of "necessary connection":

> When we look about us toward external objects, and consider the operation of causes, we are never able... to discover any power or necessary connexion… which binds the effect to the cause, and renders the one an infallible consequence of the other... In reality, there is no part of matter, that does ever, by its sensible qualities... give us the grounds to imagine, that it could produce any thing, or be followed by any other object, which we could denominate its effect...[10]

Hume continues by arguing that we get this idea of power from our early experience of voluntary movement—an idea elaborated by Piaget. We learn to transfer the experience of power over our own bodies onto inanimate matter. Causality, in other words, is something projected into experience by the mind. It begins with the infant's noticing correlations among bodily sensations, visual impressions, etc. The baby learns it can "cause" its limbs to move by discovering that they *do* move in conjunction with certain sensations. Similarly, it learns it can indirectly cause the movements of other objects. But conscious volition involves more than the mere association of sensations. It entails a concept of self, and a concept of a world on which the self can act, and from which it is felt to be distinct. Cause and world arise together as notions, when objects are recognized as entities mutually interacting and distinct from self. Causality is intimately bound up with the notion of isolated objects or systems and with the separation of subject and object.

The process of perception evokes an interaction among external objects in order to organize the flux of sensory patterns; it has its intellectual counterpart in the scientific quest for fundamental entities and laws of Nature. The finding of regularities in Nature gives the power of prediction, especially if they can be quantified. Over and beyond that, we still want to know what sense can be made of these patterns. Must there not be an orderly reality underlying and giving rise to the flow of appearances? Explanation calls for a picture or story about reality beyond mere mathematical description. For this reason, continuity in space and time have been essential to causal explanation in physics. Intelligible pictures of how the world works are based on metaphors of everyday experience in which continuity plays a key role. What troubled Newton's contemporaries (and Newton himself), about gravity's apparent action at a distance, was the instantaneous leap across space and time—the discontinuity that seemed to defy a causal chain. Various theories of gravity waves, gravity particles, and fields have attempted to bridge the gap. What is troubling about implications of the quantum theory is a similar causal discontinuity apparent in

"non-local" effects.

Modeling the world in terms of causes is also, of course, an everyday experience. There are stories we tell ourselves, and inner "models" we consult, to explain the motivations and actions of people. One has theories about how the world works. We are not only interested in foreseeing how people will behave but we also want their behavior to make sense that can be assimilated to our own motivations. In daily living, we tend to demand certainty and sense without the painstaking evidence required in science.

14. *Intentionality*

The ordinary notion of intention, as a conscious purpose, may be expanded to include any internal connection made within a "self-organizing adaptive system." This expanded concept still remains consonant with its technical meaning in philosophy. Intentionality can then be defined as the ability to make such connections. In this sense, intentionality is a complement to causality, since both involve attributions of connection. For, cause invokes an observed connection between external events, which is believed to exist independent of the observer, while intention invokes an agent making internal connections in response to external events. Another way to put this is that cause is the connection *between* events external to the observer; meaning is the connection of events *to* the observer; intention is the connecting of events *within* the observer, which map external events with meaning.

A causal account of behavior describes a sequence of physical events *in time and space*, beginning with the input of a *physical* stimulus at the boundary of the organism, and moving through causal connections that are understood as electrochemical processes within it. An intentional description, by contrast, appears as a *logical* sequence independent of time or space, commencing with a sensory input of *information*, and moving through logical connections toward a conclusion that may be expressed in behavioral output. One description has physical laws and events as rules and operations, and the world of physical causes as field; the other consists of logical operations and information flow within a field of values, priorities, or intentions. Motor behavior, and neural processes involved in perception, can be described in causal terms, but this fails to provide a description in whose terms the life of the organism as agent can be properly understood. N*either* the meaning of an organism's behavior nor of its subjective experience can be understood exclusively in causal terms. Behavior cannot be explained causally beyond the reflex level, and what *can* thus be explained is merely the behavior of isolated neural circuits, not that of the organism as a whole. Experience cannot be explained causally either—which, of course, is the Mind-body Problem.

In short, scientific materialism, expressed in terms of causality, cannot account completely for the organism, let alone for the experience and activities of human beings. A song, a painting, a novel, or even a word cannot be adequately described in third-person terms only. All these involve meaning and intention, which cannot be elucidated merely in terms of physical processes involving neurons, sound waves, or molecules. Such processes must additionally be understood as units and elements of *logical* circuits—of a flow of information conveying meaning, as well as of energy or matter. Physiological explanation may be true, on its level, but is not exhaustive. An intentional system (which, if it is an organism, is also necessarily a physical system) has properties radically other than its properties as a physical system. This is another way to say that mind is categorically other than body. Cultural signs are visible expressions that carry public meaning and so belong to a distinct domain (for example, Popper's *third world*). But even the inner "language" of intentional connections, such as those within a brain, may be viewed as a non-physical domain of symbols. This should not be confused with the domain of the person's conscious experience. While my intentions must correspond with at least a subset of my brain processes, I have no privileged access to those processes, as considered from a third-person point of view. My speculations about what goes on inside the black box of my own skull are on the same footing as my speculations about other creatures.

All living things do appear to manifest intention. The bean pods of certain plants, for example, explode when they are ripe and dry, scattering the seeds far from the parent plant. Some seeds have burrs

that catch in the fur of passing animals. There seems to be a purpose at work, as the propagation of the plants is favored by this dissemination. In some sense the plants have learned their clever ways. But to call this a purpose or intention in the conscious human sense is to project our human experience into natural patterns. How then do we account for the seeming purposiveness of Nature's activities? One solution would be to call it by another name. To say, with Aristotle for instance, that the seed-disseminating mechanism demonstrates a 'reason', or 'final cause', as well as an efficient cause. While we do not impute any conscious purpose to the plant itself, we are left wondering *whose* reason? If we assume a Creator of the universe to have designed this ingenious stratagem for the plant, this merely transfers intentionality from the system in question into the hands of an intentional Agent outside the system. It does not explain intentionality itself. It is scarcely more explanatory to impute a purpose to Evolution, to Nature, or to the Cosmos. We are left begging the fundamental question, which is to understand what *purposing* is, and what sort of entities can do it and why. This mystery begins to unravel if we generalize the notion of intention to include a wider range of phenomena than the conscious purposes of persons. From this broader perspective, human volition, goals, intentions, and their meanings can be viewed as special cases of a category that also encompasses the adaptive and apparently purposive activities of plants, of robots, and of the parts of ourselves responsible for unconscious slips. To this end, we must find an understanding of intentionality that allows us to accept that organisms in general, not just conscious persons, hold intentions. Whether or not a robot can ever be said to be conscious will depend first on whether the robot is an organism, with its own intentionality, and only then on whether it is also a person.

So, an intentional connection is a logical operation or relation, within what we are calling an intentional system, that is meaningful to that system itself. This stands in contrast to causal connection, which is held to exist independent of any intelligent agent. Thus, a connection between billiard balls striking one another is causal; whereas, the connection in my mind that attributes the motion of one billiard ball to that of another is intentional.

An intentional system can be understood (at least metaphorically) as a formal system that is "interpreted" as referring to the real world, just as Euclid's geometry can be understood as an abstract system than can be applied to physical space. Intentional connections are "real" to the degree they consist also of physical relations, which instantiate them and can be described causally. The mental *is* the physical when intentional connections coincide with causal connections, which they do in any organism.

In the natural world, intentional connections are causal connections within an organism that map other connections external to it or representing some aspect of itself. Intentional connection extends and reflects, within the organism, the causal connectivity of the world, so that the organism (in particular the brain) is largely an image of the world. It could be argued that intentional representation within the organism is *caused* ultimately by the connections in the world it represents. If this connectivity, internal to the organism, is simply part of the causal continuity of the world, then why single it out by calling it intentional?

The distinction between cause and intention has to do with the organism as *agent*, as active cause, in contrast to the passive reception of causal influence in which inert matter is involved. In our commonsense notions (as opposed to pure behaviorism), causal processes may originate in the environment and *affect* the organism, but they are not held to cross the threshold of the creature's skin to *determine* its actions, except in the grossest reflex. Rather, the organism is perceived as a causal origin, an autonomous agent determining its own actions. According to this picture, what crosses the organism's boundary is information rather than physical cause. The information is carried, of course, on energy which has causal powers. But the organism *responds* to this information by acting to change not only itself but also the environment; it is not strictly driven by the information, but "processes" it to select among options. (Whether the options are pre-determined by the environment is another question, which we need not go into. There is some middle ground between the ideal of freedom and notion of determinism).

Still, we might ask, how is description in terms of information processing anything but causal? How is "information" (which after all is tied to energy) anything other than causal process in another

guise? However, we must not confuse the ordinary use of 'information', which refers to *meaning*, with the scientific (that is, Shannon's) definition. In the latter sense, there is no significant difference between transfers of energy and information, which has been used as a concept parallel to entropy. The notion of 'information processing', however, glosses over this distinction. The fact we should keep in mind is that the information in information processing is *for* an intentional agent, who evaluates it in terms of its own priorities before responding to it with its own energy. Otherwise we would be dealing simply with causal processes. This is a relative, and not an absolute distinction, which is why computers could one day be intentional, and even conscious, agents. At the present time, however, computers are not autonomous agents in the way that organisms are. Information processing within a computer is meaningful to the programmer, not to the computer itself. Information processing within an organism, in contrast, is the organism's own activity as an autonomous agent. In both the computer and the organism, physical connections or events are the carriers of an information flow and are *logical* connections and events as well. The intentionality in one case belongs to the human programmer or operator, in the other case to the creature itself. In order to have its own intentionality, a computer would effectively have to *be* an organism.

Cause and intention are complementary descriptions. The organism may be conceived as a continuous part of an undivided environment, with causal processes freely crossing the boundary of the organism both ways, so that the distinction between organism and environment becomes blurred and arbitrary. It is also possible to turn the distinction inside out. According to Maturana and Varela, the organism, as a closed system, does not act upon an environment at all, but only upon its own sensory inputs.[11] To the observer it is clear that an environment does exist, but there is no need for the organism to conceive of it or refer to it, only a need to act in such a way as to maintain itself within tolerable limits.

The concept of information, as intentional connection, is a middle ground between the material and the mental. An intentional connection, within a physical organism, is also a causal connection, and the flow of information is carried on energy processes in the physical world. The strategy, for bridging the explanatory gap of the MBP, then can be to concentrate on understanding how it is that a complex informational system (an organism) can be subject as well as object, agent as well as causal system.

Etymologically, the word *intend* means to "stretch out for, aim at." Its archaic sense means to direct or turn the mind, eyes, thoughts, etc. It is this directedness of mental states toward objects and events in the world that gives intentionality its meaning in philosophy as *aboutness*, usually in reference to objects or states of affairs in the extensional world. The traditional meaning in philosophy is closely bound to *assertions,* in language and logic. Beliefs and desires are intentional, since they have grammatical objects and are assertions of a sort. A belief is always a belief *about* something or *that* something is true. Desire is always desire *for* something or *that* something should be a certain way. Such a relationship between subject and object is acknowledged within some languages by the subjunctive case. In the broader sense we are using here, all cognitive judgments, including pain and pleasure, are intentional. In many instances, if not all, such feelings are clearly associated with events in the external world and can be formulated as statements about it.

Moods and some other subjective states, though consciously experienced, may not seem to involve assertions or to be about anything in the world. But if carefully introspected, moods and all emotional states are found to have definite, localizable referents in the body, if not in the external world. These are specific sensations to which the mind attributes a meaning, even when it cannot be verbalized. One can simply feel anxious or depressed, for example, without knowing why. On the other hand, one can feel anxious *about* something tangible, or depressed *that* something unpleasant has occured. The apparent lack of assertion or intention in the former case is due to the fact, that neither the mental imagery nor the real possible or actual events behind such bodily referents (to which the mind is actually responding and giving meaning), are present to consciousness, even though the referents themselves are. This is comparable to the effect of subliminal perception, which is known to be capable of producing anxiety and accompanying physiological changes. These bodily sensations are the tokens bearing the interpretive semantic content of emotional states. In other words, the somatic referents of feeling states

are connected to thoughts behind them— which thoughts may be assertions about the world. This connection between thought and somatic referent happens in just the way that words or mathematical symbols are connected to their meanings. Intentional connection is simply a mapping from one domain to another, regardless of its accessibility to consciousness.

As in spoken languages, there is latitude for play in the language of this mapping. Subjective consciousness, after all, is possible only because of the absence of a strict correlation between experience and world. Nevertheless, such correlation must be the rule of thumb if the creature is to survive. Consciousness generally *aims* at the world even if it sometimes misses. Intentionality, as the aboutness of experience, must correspond closely to the survival-oriented programs of the organism. Emotional states must refer ultimately to states in the world, if only the part of the world that is the body, or else to the organism's relationship to the world.

Thought and perception, of course, are intentional. To see is to recognize and classify, a process analogous to forming a theory about the nature of what is seen. Visual experience of an object is inseparable from the belief that the object is *there*, and there as seen. Similarly, pain and pleasure are also intentional, like other cognitive judgments. Physical pain is the "belief" of the organism that its tissues are being damaged. Emotional pain may be considered the belief that one's psyche or self or life is being damaged. The occurrence of pain or suffering is, in effect, a proposition about the world, if only because one's body and personality are parts of the public domain of causes and effects as well as apparently private theaters of experience. Anxiety and other vague feelings are equally intentional, since they have significance to the organism. If one has a vague feeling of malaise, it serves the same function as pain in raising an alarm. The belief is that *something* is wrong, even if we do not know precisely what or where.

Nevertheless, some philosophers insist that pain, moods, and other somatic states are not intentional because they do not refer to anything in the external world. This is simply a mistake, I believe, and probably a result of the historical development of the concept of intentionality within linguistic analysis, rather than as a biological or systems concept. Pain certainly *does* refer to something "beyond itself," namely some injury or disorder, whether the cause lies within or without the body. In the case of spontaneous pains (no evident injury or disorder), we could say that the nervous system is in error. The fact that a system is prone to errors does not negate the meaning of its proper functioning. That there are pains that are false, in the sense that they indicate no real damage or threat, does not mean that pain normally refers to nothing. The same is true of any sense modality. Visual and auditory errors and hallucinations are possible, but that hardly means that all sights and sounds refer to nothing external.

Sometimes one looks past the subjective experience as such, to the real event it betokens. And sometimes one looks at the experience itself, at the token. The difference between these perspectives is like the difference between what words convey and the words as sensual objects in their own right. It is a mistake to treat pain (or any other subjective experience that is not an error) *only* as the word itself, as though it conveyed nothing.

In the 19th century, Helmholtz proposed the metaphor of unconscious inference to explain the nature of cognition. What he was pointing to was what we are calling intentional connection. The idea is that perception mimics the conscious activity of theory-formation and logical inference. However, it is perceptual processes which are primary, not language, logic, or scientific thought. Helmholtz's metaphor should be turned the other way around: logical inference and conscious theorizing mimic and extend perception.

In language and perception alike, the mapping from one domain to another is intentional. Generally, we are unaware of the process of putting words together into sentences. Speech simply flows, we know not from where. We are similarly unaware of perceptual processes, simply experiencing what appears mysteriously as an end product. And it has been widely recognized since Freud that a person can act with unconscious intentions. Why, then, restrict intention to the narrow meaning as conscious purpose? In our daily involvement with creatures, we concede that they act with intention, though we may be unsure how far down the phylogenetic ladder they should be credited with consciousness. It simplifies matters to assume that every aspect of an organism's activity is intentional, quite apart from the

question of consciousness. This is but a way of saying that an organism comprises a system of connections through which it actively maintains itself.

If a beetle, upon being prodded from behind, opens its wings and takes to the air, there is an intention and a logic in its behavior. Being bumped *signifies* to it a dangerous contact, whether or not "sentience" is presumed. The reaction of flight *aims* at escape, whether or not "feeling" and "deliberation" are invoked. The difficulty is that we project our human experience of purpose, reasoning, feelings, and concepts (e.g. of "danger" and "escape") into the beetle's situation. We can only conceive its situation through our human representation of that, and so we are tempted to think that the beetle must similarly represent its situation to itself in some way that we can only imagine is like our own consciousness. We can deliberately refrain from this temptation to project our conscious meanings, however, and remind ourselves that it is *connectivity* which is the broader concept and primary. Representation is a form of connectivity, and consciousness is an even more particular form of representation.

Just as the notion of causal connection depends on the existence of discrete objects or isolated systems, admittedly the notion of intentional connection depends on the existence of self-contained agents. Such notions can be expanded—even to the point of dissolution. If the universe is one continuous and undivided process, it makes as little sense to think of individual agents as it does to think of individual objects or isolated causal processes. In such a universe there could be no distinction between cause and intention, but little basis for any distinctions at all.

15. *The "Mathematical" Meaning of Meaning*

From a point of view within the cognitive domain of science, we appear as organisms, which are a product of natural selection, and over which the grand game of survival holds the power of life and death. Our experience of the playing field of that game, as meaningful and real, is unconscious recognition of this fact. The perceived realness of the world and the significance of experience, in other words, arise ultimately from the fact of embodiment in an environment that does indeed matter. As far as the Mind-Body Problem is concerned, the germain issue concerning embodiment is not that we are carbon-based protoplasm (rather than silicon and copper, for instance) but that we are players in the game. The world in which we play is the primary domain of our experience, both the beginning point and end product of complex cognitive processes. Our brains are streamlined to exclude these processes themselves from cognition, just as the workings and controls of a well-designed aircraft are transparent, allowing the pilot's attention to dwell in air space. This is only possible because the aircraft itself embodies a reserve of knowledge about aerodynamics, mechanical engineering, and so forth, that may be taken for granted.

The organisms that exist are those that have survived by taking reality seriously. As life forms, we are committed to this game. Were it not so—if we led a ghostly disembodied life unaffecting and unaffected by the world—would we be compelled to regard experience as anything more than a dream? Indeed, in such a condition, could we experience anything at all? For consciousness itself is surely a motivated strategy in the game. Were we indifferent to the state of our tissues (or lacking them altogether), pains would not hurt nor would pleasures feel good. If we did not identify with the physicality, welfare, and purposes of the body, there would be little reason to take the external world seriously enough to bother to represent it in perception. There would be no *basis* for experience. Disembodied mind is an oxymoron.

An alien creature from another galaxy would probably share our most elementary response patterns merely by being, like us, an embodied product of natural selection. Such is not (yet) the case for machines. The issue of intentionality, and the role of embodiment that underlies it, is crucial to understanding what constitutes life and mind. In particular, it poses the question of whether, or under what circumstances, artificial life and artificial minds are feasible.

Perception always informs the creature of its relationship to the world. This is obvious in the case

of spacial relations: objects are perceived as distant or near in relation to the body. The fact that something tastes "good" is cellular knowledge, so to speak, of the body's chemical relation to that substance. Sometimes, of course, this knowledge is inaccurate or incomplete. (There are poisons that do not taste bad, or bad enough to steer us clear of them!) The fact that an object looks near or far, or tastes good or bad, is an aspect of the creature's knowledge of what to expect from it and how to conduct itself in regard to it. All perception, in other words, is imbued with the priorities of the organism, which must maintain itself within a specific range of conditions (temperature, water content, salinity, etc.). This range is narrow, selective, biased, related to the survival of the creature, whose picture of the world is also a portrait of its own needs. The ideal of a transparent window on the objective world is simply the wrong metaphor to explain the cognitive processing of information that takes place either within the brain or within society.

The information processing model of mind assumes an input and an output, the equivalent of a formal system linking them, and one or more programmer-observers who define the system in their own terms. The organic model of mind assumes a physical body, as part of a larger causal system, developing through natural selection in the course of long genetic histories, which thereby establish a cognitive domain of terms belonging to the creature itself. The information processing model may erroneously assume that an organism or its brain can be exhaustively codified (see Part 3). A formal (or deductive) system *can* be so codified because it is a finite intentional creation of the human mind in the first place. It contains just what was put into it by its creators. In contrast, a biological system is not the invention of its investigators, but part of the world-in-itself. Biological systems conform to our descriptions only in the way that the real business world conforms to the game of Monopoly.

The domain of information supplied to the computer, its operational guidelines, and its output alike have meaning only in the cognitive domain of the user. The computer has no cognitive domain of its own, because some other agent has determined the relation between input and output, between its states and the world to which those states refer. Nor has it any motivation or basis on which to specify such a relationship itself. The basis for an organism's intentionality is its embodiment in the game of life, which provides the values, axioms or guidelines on which it makes its decisions and constructs its representations of the world. In the film *2001*, the computer "Hal" *was*, in effect, embodied: the entire spaceship was Hal's body, and the humans on board were mere parasites initially tolerated because of their symbiotic relationship with the host. When the humans began to pose a threat, Hal acted in self-interest, and in what "he" believed to be the greater interest of the mission. The present day computer, we may be glad, has no body, no stake in the world, and no ideas of its own. The current generation of digital computer is a tool of human purposes, not an imitation organism. So it should remain.

It is sometimes said that the computer possesses only "derived" intentionality and not the "original" intentionality of organisms. This distinction parallels that between third and first person points of view, respectively. But original intentionality is the only kind there actually is, just as all points of view are actually first-person. The salient question always is: *whose* intentionality? The confusing notion of derived intentionality arises in linguistic analysis, where an intentional statement is a type of sentence rather than an action by a cognizing agent. But if intentionality is rather a defining property of an agent, then no linguistic expression has intentionality since it is not an agent. In this sense, a book (or a statement) could only be intentional if it were itself an agent—if, for example, it could write itself and read itself.

Descartes was probably the first to foresee the possibility of a mechanistic explanation of consciousness. By means of the coordinate system that bears his name, he discovered that geometrical figures could be expressed as equations, and vice-versa. The equivalence of geometry with algebra, of visible shapes with abstract formal operations, had escaped the Greeks in spite of their fascination with number and form. But this insight was possibly the clue that illumined Descartes' search for a mechanist explanation of mind. Just as geometrical figures are generated by algebraic operations, so are the shapes and colors we experience in vision, for instance, generated by logical operations represented by physical

events in the brain. The intuition of Descartes the mathematician was to see the relationship between brain and mind as like that between algebra and geometry—a relationship between two symbolic systems that is, above all, intentional.[12]

The language of experience, like the language of numbers or words, is conventional. In the meaning of each symbol we find nothing necessary in itself, but only whatever the mind has agreed to. The redness of the color red, the hurtfulness of pain, the spaciousness of space, the solidity of objects, the realness and externality of the world—all are conventions like the meanings of words. Cognition is like an interpreted formal system, one pressed into representational service. Other interpretations might be possible for a given system, and other systems could map the same territory. The subjectively experienced suchness and realness of experience derives from the mind's embrace of its cognitive premises, in much the way that logical conclusions derive from accepted axioms. The primacy of phenomenal experience, to which Descartes pointed, expresses the mind's irreducible intentionality, its "animal faith" in its own cognitive axioms. This faith is underwritten by, and ultimately refers to, the fact that organisms committed to such axioms survive to reproduce. The world is just as it is in our experience for the same reason that we are here by virtue of being just what we are.

Descartes sets out first to reduce physics to the mathematics of space. For him, extension in space is the only irreducible quality of matter, and space is equivalent to number through the coordinate system. Physical events in the brain could thus embody logical or numerical operations, which in turn could represent spacial relations—and, hence, external reality. The mechanisms which served as Descartes' model were unsophisticated: clocks, windmills, programmed waterworks, etc. But even lacking modern concepts of information processing, Descartes understood that these were devices capable of representing, transforming and transmitting numerical operations: rudimentary computers. He grasped that the brain is the organ of perception, and that the afferent nerves relay a pattern of signals to be interpreted, rather than a copy of the world.

Descartes had studied the optics of vision by dissecting an eye in such a way as to actually observe the image projected upon the retina. Long before the invention of photography, the camera metaphor was suggestive—and also misleading. For in no way does the mere transmission of an optical image in the eye explain the process of vision. (If it did, we might as well imagine that cameras and televisions can see!) Vision is rather an active process of interpreting, not the optical image itself, but the pattern of signals generated by it. This points to another way in which the camera metaphor is misleading. In looking at an object and at a picture or image of that object, we see that they are isomorphic and that there is an analogy between the lens of the camera and that of the eye. But this tells us nothing of intentional mappings, represented by neural transformations in the brain, as opposed to optical transformations taking place in the eye. We cannot conclude, on the grounds of the camera metaphor, that there is a spacial isomorphism between our experience of the object and the object itself. To do so is to confuse the optical image with the experience that is somehow generated in the brain on the basis of that image.

16. *What it is Like to Be an Intentional System*

While we may view the brain as a causal system whose output is behavior; but in order to determine the meaning of the behavior we must consider it to be the product of an intentional system as well. Whatever else it is, the brain must be the physical support of an intentional system.

Experience, like thought, involves *representation*, which is not any kind of *thing*, but a symbolic process of mapping. This raises three questions: 1) How does the cognitive system create this representation? 2) How does the representation come to be taken as real—that is, to be *experienced* at all? 3) What is the relation between the representation and the reality it represents? At first sight, the MBP is the task of accounting for experience in terms of material process: the relationship between mind and brain. But "mind," we have seen, may be considered as a system of operations. This is no less a third-person perspective. The mystery is still that there is such a thing as the first-person perspective at

all—any such thing as consciousness or experience! Even with intentional explanations, in other words, there remains an elusive and seemingly unbridgeable gulf between third and first person. We may never get used to seeing ourselves as the part of matter that sees!

If perception is an interpretive process, it cannot be a matter of copying what we perceive through an open window, so to speak; that would beg the very thing we are trying to explain, which is how the mind comes to have experience at all. We are not simply looking through an opening onto the world, nor merely through a lens, for this implies someone else who is looking out the "window"of the eye from inside. Then we would have to explain *that* person's experience—and so on in an infinite regression. Descartes recognized that the optics of the eye does not explain vision and pointed out the fallacy of assuming another pair of eyes inside the head to watch the scene on the retina. What is experienced is actually the final product of a complex interpretive process; but one is tempted to take it for the point of departure. Hence, we are caught in a circular reasoning. The interpretation that goes on in the brain, which is the essence of perception, is like a deductive process of reasoning from premise to conclusion. The problem, then, is that we have no idea what the "premise" is. Everything we experience and know is rather part of the conclusion—the brain's product, the *show* of experience. We can only talk about the premise in the familiar terms of the conclusion, and therefore we are sorely tempted to think of the world-in-itself as something accessible to experience. This creates the illusion of the open window—that seeing is just looking out through the holes in our eyes at a world which simply *is* as we see it.

But there is no window, and eyes are not just holes leading in to another pair of eyes, or to a soul that somehow sees on our behalf. And if there were such a pair of inner eyes to watch the retina, the little being they belong to would be sealed inside a windowless head. She or he would have no way to directly check the inner scene against an outside world. The buck must stop somewhere. This miniature person, or *homunculus*, is just a stand-in for the brain, for the process of interpretation we are trying to explain.

So, let us indulge this imaginary scenario, to explore what must be accomplished in seeing and how one might go about it. Imagine, in this thought-experiment, that you are this homunculus, sealed from birth inside a chamber with no windows or exits to the outside. In fact, at the beginning at least, you have no reason to think there *is* such a thing as "outside," and no way to imagine it. You are not provided with a closed circuit TV, but only with complex instrument panels busy with flashing lights. This corresponds to the situation of the brain, which is connected by nerves to the retina. The retina itself does not present an analog *image*, but rather a digitized array of scintillating signals from firing receptors. You also have at your disposal vast and complicated control panels of levers and buttons (you have no idea what *they're* for!) These correspond to the brain's motor pathways to the muscles and organs. In short, there is an input and an output, and your job is to figure out how to connect them appropriately. But appropriate to what? You don't even know what is supposed to be going on here! A word of advice: don't be discouraged—or in a hurry. It took the brain many millions of years to work all this out. So we won't put any time limit on this project...

Naturally, you are going to have to proceed by trial and error. Suppose you try that lever over *there*. (Suddenly there is a lurch. Apparently you are "inside" something that is "moving." But this is an aside from offstage, for it will be eons before you develop such "concepts".) You try another lever, and now there is a really serious lurch as the tray with morning coffee goes flying across the room and you with it. You pick yourself up, a little shaken, realizing there must be something you can do to avoid these mishaps. Remember the pattern of lights on the panels just before your mishap, you keep a sharp eye out. Oh, there it is again! What if you push those buttons over there? Go ahead, see what happens. Uh-oh. Sorry. Well, keep trying!

With patience and persistence you eventually get the hang of "flying blind" in this "vehicle" so as to avoid upsets. You make endless discoveries, taking notes on what works to avert catastrophe, and these you tabulate as a very long list of correlations between instrument readings and buttons to be pushed. When a particular command sequence is called for by a particular instrument display, a new instrument display results, calling for a new command sequence—and so on, in a perpetual readjustment. You begin to find ways to organize and structure all this data so as to reduce it to general principles. The structure at the highest level of all this "data processing" is a theoretical model or map, which begins to

take shape as a vision of a hypothetical world outside. The model represents the same information as your data file, but in a highly condensed form—much like a three-dimensional object contains and represents all the possible two-dimensional profiles it can generate. In essence, it is nothing but a program relating input and output. You test and refine your program, continually updating it through new experience "out there."

The model shows, moment to moment, what dangers and "objects" to look out for, where the clear paths are, which things are sources of fuel, how far away they are, etc. Imagine a sonar map in a submarine, but much more sophisticated and detailed: you have information about size, shape, wavelength of light reflected by objects, and minute details of their structure. In postulating an outside world, you have created a theory that helps make sense of the mass of data. And you even come to call this navigational activity "seeing." Of course, from your sealed control room you have no direct contact with these hypothetical objects or the space they occupy. These remain theoretical features of your map or model. Nevertheless, you can't help experiencing them as *real* and *external*. (This is the fundamental activity of the Realizing Principle at work.) Simply engaging in this project in the first place has led inevitably to the sense of a real world. All you've been doing, however, is playing a complicated game inside your sealed compartment, performing simple operations that are without intrinsic sense. But because your comfort (read: survival) seems to depend on playing the game in a certain way, you've lent your play the weight of reality and meaning.

Now that the novelty of the challenge has worn off, one day you get the idea that all these tedious operations could be performed as well by a machine. So you set to work in your spare time doing a little soldering to see if you can invent some circuitry that will mechanically execute the program you've created. It turns out to be an enormously complex and lengthy project. In the end you build a supercomputer so huge it fills the entire room. But it works—squeezing you out of a job and nearly out of a home!

A few things to note about this episode: First, "you" in this fantasy are in exactly the position of your own brain (the supercomputer)! You are sealed in a compartment without windows. So, the question, of whether the model you produce *resembles* a real outside world, has no meaning because there is no opportunity to directly compare. You were not creating a literal portrait of the outside world, nor a copy of anything at all. Second, the accuracy or truth of your model can only be defined in terms of how well this modeling activity *works*. If disasters are avoided, then by definition your model is "accurate" and "resembles the world." Third, accepting the challenge of the situation led to projecting the inner model as experience of a real outer world. You began to *see* an image in the patterns of tea leaves at the bottom of your cup: a real world in the patterns displayed on your panels. Fourth, all of this was accomplished with software that was eventually embodied in hardware: a "brain." Fifth, the key to *experience* is the interaction with an environment that *matters*. To be sentient—as to be alive—a system must have its own priorities and purposes, which matter to it with the urgency of life and death. This condition exists for natural organisms through long genetic histories. It might exist for machines in the right circumstances, under which they could come to possess their own intentionality.

17. *The Problem of Cognitive Domains*

The notion of objectivity implies that there is some way the world "really" is, apart from anyone's actual perceiving or conceiving it. This, presumably, was the original face of the world before there were any observers to observe it. What happens when the observer tries to give shape in imagination to this unimaginable countenance of the world-in-itself?

It may seem quite reasonable to think that the universe is simply the way it looks to *us*, and that (apart from its dynamic evolution in time) it would have appeared much the same before the arising of life and conscious observers. But there is a subtle error in this commonsense assumption. We need only ask: appeared to whom? What does it *mean*, in other words, for there to be an object without any possible subject? This question cannot arise for the pre-subjective mind, but as conscious subjects we are stuck

with it. It's the quandary of Berkeley's famous "If a tree falls in the forest..." We will refer to it here as the Problem of Cognitive Domains, because it involves mistaking the domain of one's personal experience for a domain of objective truth.

A mind has only the relative and subjective realm of its thought and experience in which to speculate about possible other realms such as an objective world or an absolute truth. It is likely to paint its picture of the latter in terms of the former. As Xenophenes remarked:

> Mortals suppose that the gods are born and have clothes and voices and shapes like their own. But if oxen, horses and lions had hands or could paint with their hands and fashion works as men do, horses would paint horselike images of gods and oxen oxlike ones, and each would fashion bodies like their own.

In other words, we can imagine the territory only as it is portrayed in our map. While we may have had a hand in making the map, certainly we did not make the territory. Nevertheless, unwittingly and circularly, we take this image on the map to *be* the territory it portrays.

In phenomenal terms, a cognitive domain is the range of experience of a mind—all that is "expressible," so to speak, in the language or system of that mind. In terms of behavioral description, it is the range of discriminations an organism can make.

Every creature has its cognitive domain, whether or not we think of it as conscious; for all creatures make discriminations. Many rely on sense modalities radically different from the five familiar to us. These include sensitivities to water pressure (depth), polarization of light, infrared and ultraviolet light, magnetic and electric fields, etc. Each sense modality also defines a separate cognitive domain. The visual appearance of a wound to one's body, for instance, occurs in a different domain from the felt sensation of pain associated with it. A distinct level of processing within the nervous system might also be considered a cognitive domain. These considerations could apply to artificial systems as well.

As an open system, an organism is immersed in an environment with which it exchanges information as well as energy. This is the situation as described by a human observer. It may be that the organism has no concept of its environment in the sense that we do, let alone concepts of information and energy. It may appear to us that the organism perceives and acts upon *the* environment, which is "really" as it is viewed and defined by humans, while its own representation of this environment is limited by its more restricted cognitive domain. But this is an anthropocentric prejudice, and the whole situation can be considered differently. A nervous system can be regarded as a closed loop. Changes on the sensory surfaces initiate activity in the brain, which (from an observer's point of view) results in some action on what the observer perceives as an environment. But this action can also be regarded as taking place exclusively upon the organism's sensory surfaces, since its ultimate effect is to bring about changes in them that the creature holds to be desirable. In this way, all that the organism is ever dealing with are transformations of its own sensory surfaces.[13] Its goal is to maintain the latter within certain limits, just as the homunculus within the sealed chamber maintains his instrument readings. The very idea of 'environment' is an intrusion from the human cognitive domain. The irony of this solipsistic analysis, of course, is that human observers making it are in the very same position. They too are organisms engaged in maintaining constant a set of internal conditions, and all their theories and grand truths may be regarded as nothing more than part of this process of self-regulation!

Conscious experience is the culmination of elaborate processes which themselves remain outside awareness. But every stage in a flow of information processing may be considered a distinct domain. Each domain, for example, may be searched for propositional knowledge about the presence of certain features. Out of this knowledge emerges a new domain, which can in turn be searched for features on that level. This bears on a long-standing debate in cognitive science concerning the role of images (and therefore of consciousness) in the forms of representation used by cognitive systems.

A representation is a mapping from one domain to another. It may be *propositional* in the way that a book, which consists of statements, represents or maps the subject it covers. While an image seems rather an analog representation, it too can be understood in propositional or digital terms. Whether or not a particular pixel on a television screen is illuminated constitutes a proposition, and from the ensemble of

such digital events a global analog representation emerges, which integrates the accumulated propositional knowledge of prior stages of information processing. We must bear in mind that it is the human nervous system which constructs the cognitive domain of the image from the domain of scintillating picture points. (The pattern of dots may mean something different to a cat or a moth.) The nervous system may then search the domain of the image for propositions in yet higher domains of meaning: esthetics, heroes, plots, etc. Each level is nested in a hierarchy of constructed domains. Analog and digital, image and proposition are terms relative to the level or domain in question. They are not mutually opposed, but have a dialectical relationship in the hierarchies of the nervous system. An image is a domain synthesized from propositional operations, and upon which further propositional operations may be performed.

A language one fluently understands is a different domain than the collection of sounds ones hears as gibberish before learning the language. Similarly, the babble of the senses, if it could somehow be intercepted before it has been processed in the nervous system, would not be the same domain as the sensory experience that results from that processing. Confusion arises around the idea that we can have some introspective access in awareness to stages of sensory processing prior to the end product that we experience. The hypothetical notion of 'sense data' was invoked by Locke and later philosophers as a kind of pseudo-object in the Cartesian style. Sense data, presumably, are what we would experience if we *could* experience the domains of sensory input and other preprocessing stages of perception. But, by definition, these domains are logically upstream of conscious experience. One is, however, able to exercise special attitudes toward phenomenal experience. Every painter, for instance, knows how to "flatten" visual space and see objects, not as three-dimensional things, but as shapes and areas of color bounded by lines. We fancy that such objects of introspection bear some resemblance to pre-conscious domains of sensory processing, but they constitute in fact a post-conscious domain, as subjective artifacts of conscious attention. It is this very ability to bracket aspects of phenomenal experience that is the basis of subjective consciousness.

The world has survived the comings and goings of generations of observers. It seems to persist during one's absence, inattention, or sleep. It must have existed before we were born and even before the arising of life. It is reasonable to think that it will continue to exist after one's personal death and after the passing of all life from the cosmos. While the existence of the universe may not depend on us as observers, certainly our knowledge of it does. The universe we know in experience is a product of our own mental activity in a specifically human cognitive domain. It is this acquaintance we mistake for the universe itself. But surely something must exist "out there," in its own right and not just as the experience of some creature? If cognition is a map, surely there is a territory! The problem is how exactly to conceive it. Every perception or conception concerning the territory is part of the map, which is the product of the mapmaking activity of a cognitive agent as well as of the territory. There is no access to the territory except through the map.

The subjectively conscious mind can grasp that it constructs experience of the world as a representation in a cognitive domain. It can, if barely, understand that its experience is not the world-in-itself, just as a film of a real scene is not the scene itself, but constitutes a separate domain. It can imagine the scene unfilmed, but cannot picture the world-in-itself unpictured. As though at once cinematographer and audience, we are eternally confined to watching our own documentary on the world outside. When the show is over and the lights go on, we discover that what we then witness is also a movie, a story nesting the previous story. We are trapped in the domain of film! In the same way, the proposition that there exists a reality untouched by mind is itself but a mind's assertion. Any conception of an absolute world, outside all cognitive domains, is nonetheless an idea in a cognitive domain. This is why the good Bishop Berkeley kept God around to hear the trees that no one else heard falling in the forest: to exist is to be perceived by a mind—so he said. The inescapable fact of being a cognizing organism is that we experience all within a finite cognitive domain. The specific achievement of being human is to have a highly flexible and relatively objective cognitive domain, with some conscious control over our point of view and where it is directed. But how can we pretend to have no point of view at all?

Every view is from somewhere. An individual subject is unique, if only because no other observer can share precisely that perspective in space-time. The observer in physics is an idealized standard, based on the relative interchangeability of subjects; but the point of view of the observer is nevertheless a first-person point of view. Many measurements are recorded automatically by instruments. Nevertheless the measurement or observation is only significant to some human being. At some point consciousness enters. What makes the standard possible is the agreement not to consider anecdotal experience in the field of scientific discussion. It is the physical world that is the focus of discussion and the object of investigation, not the personal experience of the observer as such, unless this qualifies the observation.

Points of view, particularly ones that pretend to objectivity, are usually identified with the visual sense. More than the other senses, vision lends itself to a confusion of cognitive domains. It is easy enough to conclude that the quality of pain does not exist "objectively," but only as a judgment made by a nervous system. It betokens a relationship of the organism to the environment. It is much harder to grasp that the same reasoning must apply to visual "appearances." The quality of a visual sensation has no more objective existence, outside a nervous system, than that of a pain. Classical science attempted to eliminate all subjective qualities from its descriptions, in theory leaving only properties accessible to any observer, regardless of subjective state. (Actually, this means any human observer trained in the methodology of Western science, provided that his or her results can be duplicated by others.) Vision is paramount for reasons related to the intersubjective, quantifiable, and "objective" nature of space, which make consensus feasible. In most instances we can discriminate with more precision visually than with other senses, which facilitates agreement about measurement. While *structures* (such as shapes) indicated through visual perception may be distinguished from visual *qualities* (such as colors), we must understand that the experience of color *is itself an appreciation of a structure in the real world*—namely: frequency of light.

Through instrumentation, and modelled on the objectifying power of vision, science has been able to describe the world beyond the limits of the senses. The ancient Greek thinkers, and their Renaissance followers, sought to distill out of the flux of phenomena the unalterable real aspects of matter. They set about to create a theoretical domain transcending sensory experience, yet this conceptual realm was implicitly based on the visual sense.

Scientific or objective description is supposed to take place from a "third-person" point of view. Yet all description is necessarily in the first person: from the point of view of "I." So-called objective description, which leaves out mention of the observer's experience (and that of any observed organism), merely evades acknowledgement of consciousness. What is meant by "third person" is simply that the observer is talking about the external world, not about his or her personal experience per se. It cannot be pretended, however, that the observer has some way of knowing about the external world other than through personal experience in some form.

All observers stand in a first-person epistemic relation to the world—either through the senses of the observer's own body or through some instrumentation that extends the observer's agency. The fact that I can estimate the frequency of light either through my sensory experience or with a spectrometer does not place these measures on a different epistemic footing. The spectrometer may resolve frequency better, and measure it more precisely, than human vision. But it is absurd to think that the measuring system is "objective" merely because it has no subjective viewpoint of its own; the observer, not the device, makes the measurement, which is a meaningful gesture and not merely an interaction of causal systems.

First- and third-person accounts are distinct domains of description. Though temperature and pressure are subjective, bodily, sensory, first-person experiences, they *refer* to macro properties which are presumed objective and which may be measured in other ways, describable in the third-person. To experience this same property "directly" or "subjectively" (in the first person) simply means to measure its effects using one's own sensory organs as measuring instruments.

Third-person observation isolates a system by separating the system-to-be-observed from the system-which-is-observing. If we try to expand the account to include the latter, we simply redraw the

line somewhere else, redefining the observer, who remains separate by definition. And so on, in an infinite regress! The deep truth revealed here is that the subject can never be phenomenal object: one can never see the point of view from which one is seeing. The problem posed by consciousness is not resolved through using third-person description, but is only deferred by removing consciousness from the system observed to an observer supposedly outside the system.

Nevertheless, the scientific picture—at least as portrayed in popular accounts—pretends to depict what the object looks like before it can look like anything at all: the *naked* object, the unmade-up face of the world, without reference to observers. For the sake of the ideal of objectivity, science rejects the cognitive domain of ordinary (first-person) experience as *merely* a cognitive domain, only to construct new cognitive domains that it does not acknowledge as such. Implicitly it holds its entities—conceptual, if no longer directly perceivable—to be the building blocks of the objective world-in-itself. But these conceptions often remain tied to acts of visual imagination. Recalling the homunculus in the sealed control room, the validity of scientific models may depend strictly on their mathematical power to predict correlations. Yet the huomunculus (and scientists) cannot help but experience the models, which summarize such correlations, as real. And more often than not, their intuitive appeal refers back to ordinary (subjective) experience, usually visual.

Cognitive science implicitly adopts the domain of physical description—the world portrayed by physics—as the causal substratum of the phenomena it researches. This gives rise to some confusion concerning the nature of the fields in which mind grazes. Merely to explain that the brain organizes a picture of the world, based on its interaction with the world it pictures, is too vague and obviously circular to be satisfying. One therefore has recourse to entities, held to be more fundamental than the objects of ordinary perception, as the starting points both for explanation of physical phenomena and of cognitive processes. But it is only a convenient manner of speaking. One is never aware, in any sense, of one's own nerve impulses, for instance; nor of the photons or molecules which trigger them; nor of any other such theoretical constructs in the domain of science. We are simply aware of the world as presented in our normal cognitive domain, which includes the objects encountered in daily life, but not the microscopic events of physics or neurology. To begin with the world of physics (or physiology) as the point of departure for explanations of mental activity must be done with circumspection. For the whole biological enterprise of cognition seems to leads to the physicist's constructed version of reality, which is then recycled as the starting point! Schopenhauer likened this bootstrap operation to the Baron von Munchausen, who—in order to cross the river without drowning—lifted himself and his horse into the air by pulling up on his own coiffeur! It may not be logical, but this is how we avoid drowning in solipsism!

The Mind-Body Problem is, above all, a problem of cognitive domains. The mental appears as a subjective domain, putatively caused by objective physical processes, which appear simply to *be* the world, rather than appearances in the human cognitive domain. The objective physical world, by definition, is held to exist independently of the perceiving or knowing mind. But, as we have seen, the subjective domain of experience is not unilaterally *caused* by neural processes. Rather "experience" and "neural process" are descriptions, in parallel cognitive domains, of the same events.

Consider the two cognitive domains of sight and sound. Seeing a starting pistol going off in the distance with a puff of smoke, and hearing the sound of the shot an instant later, are two experiences registering the same occurrence in the world in two cognitive domains. The visual impression is not held to cause the auditory one; neither is sound considered an "epiphenomenon" of vision, and there is no philosophical problem with their relationship. Rather, the two are experiences of the same event in different cognitive domains. The relationship between the domain of physical processes and the domain of experience *is* held to be problematic, however, which *is* the MBP.

In and of itself, the specific quality of each sense modality—what visual experience is *like* as opposed to, say, what auditory experience is like—does not constitute knowledge of the world but of the workings of that sense modality. There may be something that can be learned about the world through vision that cannot be learned through hearing, but it is not the *flavor*, so to speak, of visual experience as

opposed to auditory experience.

Objective knowledge of the world is abstracted from sensory experience, of course. It refers to *facts* derived from phenomenal experience but not to experience per se. It is knowledge of *differences* (structure), but does not refer to that which the differences are *of*. Such knowledge is objective in the sense that it consists of invariants accessible to standard observers. This is analogous to the situation in dynamics, where a difference of velocity is invariant for all observers moving uniformly, whereas absolute velocity may be meaningless. Similarly, I cannot know your experience of the color red, nor you mine, but we can both be aware of differences of color and communicate about them. The common perception of such differences is what we *mean* by objective knowledge or information. Since by definition it excludes sensory qualities themselves, we should not complain that objective knowledge is deprived of them. Nor should we be troubled to explain how the domain of subjective experience arises from domains of information or other such constructs. For they cannot logically arise *from* domains that are in fact afterthoughts. On the other hand, if all perception is perception of difference, then the experience of color (and of any phenomenal quality) must itself be a perception of differences in the objective world. The experience of color, for example, must then be a perception of the fine structure of light—and of the surfaces emitting or reflecting it—just as the experience of tone is a perception of the fine (vibratory) structure of sound and its sources.

The Problem of Cognitive Domains is that an objective domain is conceptually abstracted from the domain of experience, but then an attempt is made to derive the latter from the former. The error is that the objective domain is not recognized as a cognitive domain per se, but implicitly held to be the world-in-itself. Owing to the fact that these domains are not on the same logical level, there appears an inevitable "explanatory gap." We want to understand how conscious experience can arise through physical processes in the world, but it is naive to hold that the world-in-itself is the reality given in abstract concepts of physical science.

The error at the core of the Mind-Body Problem confuses first- and third-person perspectives. The world of physics is mistaken for the world-in-itself, as the starting point for a philosophical explanation of how "experience" arises from "material processes"—or how "qualities" arise from seemingly denatured information or structure. Yet differences and material processes are only known *through* experience, and qualities are a perception *of* structure.

The Problem of Cognitive Domains reflects the fact that we are not content, like the ancient Taoists, to let the world-in-itself be ineffable, unnamed. However, there is no access to the world-in-itself, independent of some cognitive domain, because "accessing" is an activity in a cognitive domain.

18. *Coloring It Real*

While not disputing that the world truly *is* real, we must consider that realness is also a *quality* projected upon it by the mind, a judgment concerning the status of experience and its particular contents.

Everyday experience, like scientific concepts, refers to structures or invariances in the world (ultimately to differences), which appear mysteriously "filled in" with the "qualities" we know in sensation—such as particular colors, smells, auditory tones, etc. If structure, difference, or information is considered the bare bones of reality, we are naturally led to wonder how the mind fleshes out this skeletal structure with the immanent qualities characteristic of phenomenal experience. But, as noted in the previous section, there is something fishy about this question. For, concepts of 'information' and 'structure', far from being raw data for sensation, are elaborated constructs in the cognitive domain of science. Structure is no more what remains, when the perceptual world is stripped of its qualities, than the outlines in a coloring book are what remain when the colors and shadings are removed from a photograph. An outline must be deliberately drawn in, structure must be intentionally imposed.

While it may be convenient to consider information or structure as the raw material of mind, such a starting point is at best arbitrary. Perception or sensation are sometimes treated as analog domains upon which cognitive judgments are enacted. The problem is that one never experiences such a thing as a pure

perception, sensation, or sense datum, free from cognitive judgment. Rather, perception or sensation *is* cognitive judgment. In looking at a source of blue light, for instance, one does not consciously detect "a vibration of such and such frequency" and then decide to add or project into that knowledge the quality of blueness. Rather, the experience of blueness *is* the brain's frequency estimation of that vibration, as assessed in the cognitive domain of vision. Similarly, to perceive the musical tone B-flat is not to fill an estimation of frequency with a particular auditory quality. Visual and auditory sensations are not epiphenomena. Rather, they are informationally rich propositions concerning the structure of their respective physical domains. Specifically, they are estimations of frequency. This is straightforwardly true in the case of sound: the experience of tone emerges as the wave fronts impinge too rapidly to be distinguished individually. But what is the sound of one wave lapping? It is certainly not a *tone*, which is the global effect of a series of wave fronts.

The case of color is more complex, apparently, since color is distinguished on the basis of the relative intensities of differing frequencies, for which there are separate sets of receptors. (Color vision does not operate, as once thought, through simple sensitivity to frequency. The fascinating experiments of Land, of Polaroid fame, are instructive.) While frequency is not straightforwardly the structure mapped by color experience, the latter nevertheless maps relations in the world that are relevant to the organism. Both the experience of color and of auditory tone are transforms of structures or textures in the world, which are *alternatively* described in domains posited by reductive science. They are the first-person experience, in the ordinary human cognitive domain, *of* those textures, which may also be described in the third-person cognitive domain of science.

In the case of sound, at least, the tone *heard* and the frequency *measured* by instrument are two equivalent expressions of the same cognitive judgment. A clear example is the organization of the perception of sound into octaves. Vibrations of frequencies x, 2x, 3x, etc., are perceived as *qualitatively* similar, the same note in different registers. The ear is responding to a quantitative congruency through its experience of qualitative similarity. The human eye perceives but one octave of the electromagnetic spectrum. Yet it is apparent that, if the eye were sensitive to a wider range of frequencies, perceived colors would repeat in a way similar to perceived octaves of sound; for, the farthest ends of the humanly visible spectrum (violet and red) begin to resemble each other qualitatively. In both examples this is due to the similar reaction, by the receiving sense organs, to wavelengths that are multiples of each other. An oscilloscope displays the same information: that a frequency and its multiples are congruent.

What we have said of visual and auditory sensation applies more generally to all experience. The apparently mysterious, projective, enhancing, filling-in capacity of mind is ubiquitous. And because it is universal, it is virtually transparent. Only in anomalous circumstances do we notice it at all. These include completion effects, habituation, adaptations of various sorts, and other phenomena of projection. Some examples are given in the Appendix. But the question suggested by these phenomena—of how the experience of tone or color *arises from* an analysis of frequencies (or some other construct in the domain of scientific description)—engages us in a subtle misunderstanding. The experience of sound or color does not "arise from"—but *is*—the subject's experience of frequency analysis. The question itself is an expression of the Mind-Body Problem. It suggests a homunculus within, computing answers to problems formulated in the very cognitive domain one is trying to explain. The homunculus has *his* experience—of symbols and quantities, dials and levers within the sealed compartment of the skull—not the experience of the brain that occupies that skull. Beginning with a demon in the machinery, we end up with the ghost in the machine. For no matter how good causal description gets at explaining the performance of perceptual systems, it never seems to demystify the miracle of experience.

Tones, colors, smells and other qualities may indeed be understood as completion effects, which are examples of this "miracle." If they appear as superfluously filled in around some skeletal structure of information, it is simply because such structure was abstracted from phenomenal qualities in the first place. (This is the subject-object split on an intimate scale.) The quality, for instance, of greenness is an estimation of frequency. As discussed above, it is not added the analysis of frequency but *is* that analysis in the first-person cognitive domain of visual experience. It is only "filled in" between wave fronts of

light in the way that the quality of being fifty years old is filled in between one's fiftieth and fifty-first birthdays: that is, by intention or convention. Apart from the information it conveys about structures in the world, *the irreducibly self-luminous quality of sensations is not a property of the world but an act of the intending mind*. It is a proposition expressed in the language of the senses. The mystery of how structure, mapped in neurophysiological processes, can result in experiences of color, tone, pain, or other qualities, is no more (or less!) mysterious than the process by which sound waves can carry meaning as words, or algebraic symbols come to be imbued with numerical significance.

Only by convention can information, nerve impulses, sense data, or other constructs in the domain of science, be taken as the raw data for the cognitive processes in question. So be it, so long as we realize that such starting points as 'the physical' are already mental constructs in the scientific formalism. They are not the world-in-itself. In order to procede with thought at all, we may have to live tolerantly with dualism; but we do not have to be hoodwinked by it.

Intentional states within an organism are physical connections that are also logical connections. The flow of behavior can be described logically *or* causally; in this sense, the mental *is* the physical. This identity may still appear to omit the fact of conscious experience, for both are third-person descriptions. How do we get, finally, from third to first person? How does an object, such as a human body, become subject? How does a system *have* experience, have a point of view? What *is* the solution to the Mind-Body Problem? We can recognize in this insistence the tail-biting stalemate I have called the Problem of Cognitive Domains. For, ideas of 'logical connection' and 'intentional states' are abstracted from phenomenal experience, while we want to know how experience of the extensional world arises in the first place from "mere" logical connection within an intentional system. If we are not willing to accept the answer, that it arises in the way that words or other symbols come to have meaning—that is, *by intention or convention*—then we are condemned to dualism and the inexplicable mystery of conscious experience. If intentionality is not acceptable as a solution, then we are forced simply to accept that consciousness *is*, just as we are forced to simply accept that there is a cosmos rather than nothing. Those two great mysteries remain insoluble for the same reason: because we cannot get outside of either existence or consciousness.

19. *De-realizing the World*

Though experience and knowledge are always a function of two variables—the world and the self, the object and the subject—there is a cultural tendency to ignore or discount the subjective aspects and origins of cognition. One likes to think that the world simply is as one believes or experiences it to be. It is tempting to disown epistemic participation by claiming that one's experience is driven by reality, through a direct access to the objective variable. One naturally eschews language smacking of subjectivity, or suggesting that one's claims are not the objective truth. Pointing to the world effectively directs our own and others' attention away from the mind's subjective contribution to experience. This may yield a short-term gain in hoodwinking others; but, if one believes it oneself, one forfeits the power to shape one's *own* experience and behavior.

Since it is the seeming independence of the world from mind that defines reality, we tend to objectify everything about which we need to feel assured and confident. We justify our personal values as objective truth. We justify our actions as compelled by objective necessity. We want to feel compelled to do what we do by forces or reasons outside us, beyond our control, more substantial than mere whim or subjective intention. If knowledge, attitudes, and beliefs are to be held as meaningful, they must refer to possible actions within domains of vital concern to the organism. Knowledge or belief about imaginary realms, or realms on which no action could be taken, is useless. If the perceptual world is to be experienced as real; if cognition is to have meaning; if belief is to be valid; and if behavior is to be justified, then the dependence of experience, knowledge, belief and action on subjective mind must be denied. And the denial itself must be denied. Reality must have no cause in the mind, but must spring immaculately into experience untainted by subjective participation and unhampered by acknowledgment

of mental lineage.

This parthenogenesis of reality permeates all levels of realism. We are, as it were, born realists; and realism as a philosophy recapitulates the fundamental outward orientation of the brain. The brain automatically factors out subjective aspects of its input, precisely in order to leave an objectified picture of the world. For instance, changes on the retina due to movements of the eyes or head are not credited to objects, which continue to appear as stable entities. This is because of the need to represent the environment in a way that facilitates choice and action within it. The focus of this representation is an unambiguous assessment of the options upon which the organism's well-being depends. And yet there is a point at which rigid formulas break down and are no longer adequate to the complexity of the situation. We may long for security, certainty and simplicity, and seek these within ideals of reality or truth. But merely embracing the *experience* of certainty, however comforting, may be illusory and dangerous. While the world has a real look to it, the subjectively conscious mind suspects that this appearance is somehow a product of its own looking, and that seeing through real-colored spectacles can be less than reliable as a guide to what is. Sages of both hemispheres have always warned us about the apparent solidity of things.

Most biologists today reject the notion that evolution has a direction. In the new fields of complexity and self-organizing adaptive systems, in contrast, some scientists believe that there is an inherent direction of evolution in such systems toward increasing complexity. I would add that an increase in complexity in a cognitive system implies a greater objectivity as well. Mind represents the world to itself, initially for limited and parochial purposes. Its primary responsibility is to map the world in whatever ways have selective value in the contest of life. Its cognitive programs are designed for advantage rather than truth. The mind therefore treats its cognitive beliefs as true and real, in order to reinforce its commitment to them. However, as the representation becomes increasingly adequate as a tool of survival, it may also become increasingly comprehensive, global, abstract, and disengaged from particular mechanisms and purposes. At this secondary level, mind becomes flexible enough to change its programs, allowing the organism to better adapt to changes in the environment. Increasing complexity does not guarantee freedom from bias or interest, yet it allows the relative possibility.

Subjective consciousness allows mind to transcend and change its own programming. One dimension of this power manifests through science, in the quest to discipline thought and imagination, turning it toward objectivity. Another is seen in the quest for spiritual freedom, from the contingency of experience and the programming of the biological self. In both cases, mind seeks to free itself from deception and to transcend identified self-interest. While truth begins as a strategy in the game of survival, life expands its capabilities of representation, responding with ever greater subtlety to the world. The movement toward objectivity is paid for, however, with the awkward overhead of self-reference. Mind is able to model its own situation and to grasp that it necessarily lives in a context of illusion. This sometimes painful awareness is the price to be paid for the privilege of self-consciousness. It is the precondition for touching truth, when mind was designed for advantage.

In de-realizing the world, as in splitting the atom, an energy or potential is released and at the disposal of the subject. This energy was formerly locked within instinctual and unconscious bonds between experience and behavior. As experience is subjectified, the power of the world to determine the flow of one's mental-emotional processes and behavior is taken back within the psyche and experienced as conscious choice. All experience then takes place "in quotes," so that it can be considered a subjective artifact rather than a reality "out there." In questioning our beliefs, we find the freedom to act rather than merely react.

20. ***The Evolution of Consciousness***

As brains evolve toward greater complexity, mind evolves toward two seemingly opposed capacities: subjectivity and objectivity. To subjectify is to bring experience in relation to the perceiving

subject. To objectify is to discern a reality transcending and independent of the subject. These two movements are reconciled as the polarities of a dialectic that is the motor of the evolution of mind. Subjectivity involves contextualizing a given model for the sake of creating a more objective one.

To understand the role of subjectivity in the evolution of consciousness as a biological phenomenon, however, one must be clear what is meant by consciousness. We sometimes entertain an image of the experience of other creatures as an interior life, whether or not we imagine that they are also self-conscious. But this image is a product of our own self-consciousness, which is distinct from the simple awareness we picture as the interior life of other creatures. While we may conceptualize our own subjective consciousness as a meta-awareness, transcending simple awareness or sentience, we must admit that this category of 'simple awareness' or 'sentience' is our own creation, not something we should impute to the cognitive domain of other creatures. We are stuck with self-consciousness and the sentience we impute is always some version of it.

Self-consciousness implies the emergence within a system of a 'subjective point of view.' But only a self-conscious system can ascribe to itself, or to any other system, a point of view. And any point of view we can imagine is inevitably viewed as a version of our own self-consciousness. This is the only sense that can be given to having a point of view at all; otherwise, the question of the "sentience" of a cognitive system remains indeterminate. For, unless a system is self-conscious and can claim its own point of view, any point of view ascribed to it is really that of the ascriber. This is the basic problem involved in imagining the experience of other creatures—let alone machines.

Clearly we are involved here in a Problem of Cognitive Domains. Yet it is tempting, and perhaps fruitful, to consider 'simple awareness' as the state of which subjective consciousness is a meta-state. And, if so, this suggests in turn that simple awareness itself could also be a meta-state of some yet more primary activity. This could be what we have been referring to as connectivity, representation, intentionality, etc. Then simple awareness would characteristically be a function whereby a system registers the entry of information within it. In other words, it would be a monitor for incoming data. It is possible, of course, for input to be stored directly in memory, bypassing such a monitor. There would be a housekeeping advantage, however, for the system to keep track of what data has entered it. Simple awareness would acknowledge information at the point of entry. There is anatomical evidence supporting the idea of awareness as a monitoring system, separate from a primary flow of information processing. The main pathway from the sensory surfaces to the cortex is not exclusively associated with awareness, which also depends on a separate structure, the Reticular Activating System.

The fact that humans can engage in functional acts of perception that do not involve awareness might appear to complicate a naturalistic explanation of consciousness. It is the basis for "zombie" arguments, that a person (or machine) could do all the things that a conscious person does, but without consciousness. It obliges us to account for other, factors besides intentionality, that might be required to explain conscious perception. This means showing why one internal organization rather than another should give rise to conscious experience, which is the very heart of the Mind-Body Problem.

For example, it has been demonstrated in numerous experiments that information can be retained in memory without entering awareness. A person can note some detail, for instance, without not knowing that she has noted it, and yet be able to act on the basis of the information noted, recovering it with at least a statistical certainty. However, this certainty belongs to the cognitive domain of the testing observer, while the subject is merely guessing. In awareness, by contrast, the information not only is known to the subject, but the *fact* that it is known is also definitely known. Then the information exists *for* the subject, rather than for the observer. There is a clear advantage in the ability to recall and use it "at will." Therefore, consciousness *is* functional. Zombies simply could *not* do all the things their conscious counterparts can do.

Whether or not information becomes consciously represented is generally not a conscious decision. Contents pre-consciously screened out may continue to be elaborated unconsciously and to maintain potential access to consciousness. The relation between conscious and unconscious is in constant flux. The Unconscious is not a fixed compartment of the mind, but a shifting relationship between potential and manifest.

In some sense, the "display" that conscious experience is involves an interface between internal subsystems, or between levels of processing, much as a computer monitor displays an interface between parts within the total system comprised by human user and computer. Note that this does *not* imply an inner someone looking at the display! An "image" in such an interface is simply a way of organizing propositional information within and for the system itself; it does not imply an *optical* image, but is rather what the system reads into a pattern of information. It is an analog synthesis of digital processes, a global representation on which various specializing subsystems can perform their appropriate operations, and from which they can draw information contributed by the other subsystems.[14] An image is a synopsis, so to speak, of the work performed in various departments, continually updated. The great corporate bureaucracy does the bulk of the routine legwork, which includes preparing such images, while the conscious system is reserved for novelty and urgency, for executive decisions it makes by reviewing them.

Attention, as the volitional aspect of awareness, is the ability to direct the flow of information through the awareness system. Attention is drawn by the unexpected, but may also create its own novelty through deliberate focus. Awareness is a strategy to deal with the novel or unexpected—that is, with those situations not automatically covered by fixed algorithms of the system. Routine situations, already mastered, may not require awareness, but new situations do.

It is an advantage to a system to be open to ongoing modification through interaction with its environment. The hallmark of awareness is precisely that it is a more flexible response to the world than reflex. The cognitive system transcends its fixed algorithms by creating an awareness system meta to them. This expanded system interprets aspects of its information flow as features of a real world. It then transcends this fixation on the world by creating subjective consciousness—a system which reinterprets aspects of its objectifying capacity as processes within itself. If these movements seem opposed, the essence and evolutionary advantage of each is the same. Both the realizing and subjectifying functions increase versatility and objectivity. If there is a need to relativize and subjectify experience and thought by bracketing its formal aspects, there is also a corresponding need to re-engage on higher levels, to assert more adequate models, to find better metaphors and explore wider worlds. To leave a nest is to find oneself nested in a larger matrix, with its own apparent reality, and the task of subjective consciousness is posed again on a new level. Neither creating nor deconstructing worlds alone converges on objectivity. But these seemingly antithetical movements of mind participate together in a dialectical evolution.

The limitation of a wired-in reaction pattern is its mechanical rigidity. In contrast, a global representation—an image monitored by subsystems capable of variable responses—is a synthesis upon which further analysis, as well as action, can be performed. The greater flexibility it affords depends on the gap it widens between stimulus and response. The course of action, though implicit in the image, is left relatively open to choice—and also lack of choice. At the extreme, the gap could be so large that one would see the world disinterestedly and have no response whatever. Obviously, to be viable for an organism, detachment must be relative and circumscribed. The advantage of an image of the world is that it can be studied, thereby postponing response, which is freed up to be more versatile and objective; but in fact this is a matter of reconnecting action to some higher domain or in some more effective way. Otherwise we are left with an organism that has no priorities, and is paralyzed in the indifference of total freedom.

The direction of the evolution of mind is toward an ever looser relationship to the axioms of organismic existence. Paradoxically, this relative freedom has evolved only because it serves the organism all the better in the game of life. The human "generalist" strategy for survival lies in the ability to transcend or override any particular algorithm. Such freedom has its dangers, like having too much "play" in a steering wheel. With our relatively free will, and enormous capabilities for transforming the environment, we humans are both free and capable to destroy the natural support for our very existence. From a long-range evolutionary point of view, this freedom to play God can only be justified for the purpose of serving the mandate of life on a higher level than we may already recognize. To transcend a truth is to find a larger truth, to open to a greater, more objective reality that is binding on a higher level.

One could say that the evolutionary purpose of subjective consciousness is to liberate the subject

from the object. In the extreme, this means an epistemic subject utterly free from the system of cognitive behaviors that is usually taken as the self. At the core of that self is perception, belief, feeling, and thought—all potentionally mistaken as reality. Before subjective consciousness, there is only the world. In subjective consciousness, there is both world and self at odds. *After* subjective consciousness—if that is even conceivable—perhaps there is once again no distinction between world and self. The evolution of complex adaptive systems proceeds through sensitivity to awareness to self-consciousness. If the organism has transcended its programs once in creating awareness out of sentience and instinct; twice in creating subjective consciousness out of awareness; perhaps a third movement toward objectivity lies in transcending the subject/object duality. Having played its role, perhaps the separate self with its imagined freedoms will ultimately be discarded as an institution that has outlived its usefulness.

Part Three: CONTENT AND FORM

21. *Analog and Digital*

Whatever is "out there," in the wild of the world-in-itself, must be tamed by mind in order to be perceived at all. Just as our immediate life is in the human world, and not in raw Nature, what we perceive is not chaos or wilderness but something pre-domesticated by mind. This partnership between world and mind can be elucidated through an exploration of the concepts of *analog* and *digital*.

An analog domain must be digitized in order to *read* it—that is, for it to have meaning that can be acted upon. A mercury thermometer, for instance, must be calibrated into units, which give a measurable significance to the level of the mercury. A thermometer without markings, like a clock without a face, is a perfect analogue but cannot be accurately read; in neither case is there a scale to define readings. A thermometer without a scale tells only "more or less" the temperature. A graduated scale, on the other hand, tells the temperature to the nearest marking. But the gap between markings is, similarly, an undefined analog space. The space between markings on a thermometer can only be estimated until a finer set of markings is applied to it in turn. Like a digital read-out, the scale—however fine—is defined in terms of discrete units. Instead of "more or less" it tells us "either/or": either the temperature is x or it is not, where x is a whole number. (If it is not x, the closest to x that it can be is either x+1 or x-1.) For a given scale, nothing is defined between its finest gradations. The digital is precise-by-definition, but is accurate only to the limit of definition provided by the scale. The analog is undivided, but ambiguous throughout until read by means of a digital scale imposed upon it. It is precise-in-fact, since it maps a function perfectly, but is uninterpretable since it cannot be read.

An analog domain is a continuum, infinitely dense and assumed to exist at all mathematical points (while strictly speaking it is defined at no point). A digital domain is defined only at specific points, which are the boundaries of undefined gaps within it. These gaps, upon closer examination, are analog domains in their own right. An analog domain is a territory, while a digital domain is a map-like grid employed for the purpose of relative comparisons. The analog is implicit content; the digital is explicit form. The analog is Nature, the thing-in-itself; the digital is artifact: measure and idea imposed upon the world.

A digital process involves binary decisions. Alternatives are systematically discarded until but one remains. A digital process rests upon the elimination of alternatives, and is implied whenever there is a mandate to decide, to choose, to discriminate, to reach a state of certainty. Since action must have a basis in certainty, the relationship of animals to their environment involves such binary processes. A plant can draw from the soil and air what it needs just as it needs it, on a continuous basis. It has no possibility and no need to act decisively. While it is true that grazing animals can browse as they require, they also enjoy the capability to mov to another location when a food supply is depleted, or to flee when threatened by a predator. Hunter and hunted alike are obligated to discrete actions requiring decisions.

Relations within the organism may be largely analog, but relations with the world have inevitable binary aspects, because the creature is itself a unit and must act in relation to other objects that are also units. Information concerning the world must be organized to facilitate choices required by the mandate to survive.

The concept of information is technically defined as the number of binary decisions required to completely specify a unique alternative out of many. But the amount of information presented by a situation depends on how the situation is defined. A day, for instance, can be divided in two parts (a.m. and p.m.), or into more parts represented by hours, minutes, seconds, etc. Only one bit of information (i.e., binary decision) is needed to specify the "time of day" if all that matters is whether it is morning or afternoon. More information is required to pinpoint the time more exactly. If time could be indefinitely divided, the amount of information required for infinite precision would be infinite. In practice, there is a tradeoff between speed and accuracy. There may be theoretical limits as well; time itself may be digitized at the quantum level.

Digital processes involve a reduction, or "chunking," of information. At any moment, a creature (or other system that makes choices) must reduce the plurality of options to a manageable set—ultimately to one. There are an infinite number of options that Hamlet could have pondered besides whether or not to *be*: whether to eat or shave, take up mathematics, change his name, speak in tongues, invent a character named Shakespeare, etc. All these alternatives were discarded in order to arrive at what he considered *the* pressing issue.

Information must be discarded in a cognitive system; things that are different in detail must be treated as the same. In order that choices can be made between broad categories, with irrelevant details ignored, cognitive processes bracket together masses of information, thus reducing the information load on the organism. In this way, we recognize the half-dozen official colors of the rainbow out of the continuum of visible wavelengths; we sort out domestic plants from a background of weeds; we distinguish good from evil.

What is gained through this chunking is a presentation of simplified alternatives with which the decision maker can cope. What is lost is knowledge of the finer points, as well as knowledge of the decision-making process itself. An election, for instance, passes through several stages of such chunking in which the candidates or issues are pre-selected. The electorate as a whole does not participate in this process, but decides only among the pre-selected choices.

Cognitive processes in general involve decision making among a predefined set of choices or categories, which are then accepted as given. This given appears as concrete and real, a natural feature of the world rather than a prior stage of discrimination. The decision maker recognizes responsibility for choices presented in consciousness, but not for those resulting from prior processing. The political issues appearing on the ballot (or in the news) are simply presented as *the* issues. We may be aware of shopping according to our tastes, but are rarely aware of how our tastes were formed and shaped to a limited range of products. A man may desire a beautiful woman without understanding the social context and meaning to him of feminine beauty.

As we have seen, by definition any cognitive system (including possible sentient machines) is "intentional." A cognitive domain involves a schema projected upon the world-in-itself by such a system. A true analog domain has no schema projected upon it, is not calibrated, chunked, categorized, digitized, etc. It cannot be read or interpreted, would fall outside any possible cognitive process, and so could not be experienced at all. It would be nothing other than the inscrutable world-in-itself. There is therefore a subtle contradiction, a problem of cognitive domains, inherent in the *concept* of the analog when this is conceived as an accessible domain. The analog is itself a product of definition, an idealization of digital processes carried out to an indefinite fineness; but then it is recycled as the aboriginal domain upon which cognitive operations are performed. Applied to cognition, the term 'analog' denotes nothing fixed, but a relationship in flux. There can be no true analog domain in perceptual processing, which always involves binary processes; there are only domains already chunked and read on differing scales, but in such a way that one may *appear* continuous and dense relative to another. In an absolute sense, all information is necessarily digital, just as perception necessarily involves acts of cognition, though it may seem to be

passive reception. Since information *is* binary choice, the alternatives chosen from must have been predefined.

Since analog and digital are relative to scale, there is a dialectical interplay of these aspects in any system. In particular, it is possible to represent an analog domain digitally to any desired approximation. A TV screen, for instance, gives a digitized representation of an (analog) optical image. The screen could be made up of more or less pixels, yielding a finer or cruder approximation. By the same token, a digital computer can simulate analog processes as closely as desired, but never perfectly.

A nervous system involves a complex interplay of stages of information processing. It is presented initially with an analog display—for example, an optical image projected onto the retina. (It could be argued, because of the quantum nature of light, that even the optical image is ultimately digital.) This image is digitized by the mosaic of the receptors, of which there are only a finite number. These, however, do not discharge in a simple digital mode like micro-switches, but involve analog aspects as well—like water waves sent down a canal. Overall patterns are conserved to some extent as analogues of each other in successive stages of processing, and chunking at each stage must be performed on domains that are relatively continuous and global. At the level of the organism as a whole, the creature makes grossly binary choices, but these rely on a global representation of the environment it inhabits. While digital throughout, its perceptual system begins and ends as an analogue of the world.

A mechanical clock with hands and face is an analog device with a digitized scale. If the face of the clock indicates hours only, there cannot be certainty or agreement about the time to the minute, since the position of the missing minute hand must be estimated by a judgment call. With a digital clock no estimation is called for—or possible—since a digital clock displays no visible analogue of the passage of time on which a judgment call can be made. From a digital clock that indicates hours only, we can only guess what time it "is" between hours; and this is only possible through a subjective sense of time provided by us, not by the clock. With a such digital clock, the entire duration of time between 3 o'clock and 4 o'clock is *defined* to be 3 o'clock. It is either *precisely* 3 o'clock or it is not 3 o'clock *at all*; there is nothing in between. In one sense, a digital readout eliminates uncertainty and regiments agreement about the time of day. But since we suspect that "time" really continues to pass during that hour, in another sense digitation increases uncertainty. If we had no intuitive notion of time independent of our definitions, we would always know exactly what time it is by our digital clock. There would simply *exist* no time between the defined gradations! But if we believe that digital time is merely a convention, and that all the while a real (analogic) time is continually passing, then we must view digital time as inexact in principle. The uncertainty is relative to the possibility of a finer-grained cognitive domain.

The analog/digital distinction parallels that between cause and intention. The causal world of (classical) physics appears to be an analog domain, infinitely dense. But Nature may or may not be unfathomably rich and deep. While we make a distinction between natural and artificial systems, it is conceivable that there *are* no natural systems. What we call Nature could turn out to be a deductive system: intentional, digital, artificial and finite, a mere product of definitions. This would be precisely the case if the universe *had been* created by God (or by vastly superior aliens)! If a bottom could be found to the complexity of Nature (e.g., truly elementary particles), would this constitute evidence that the universe is in fact an intentional construct, a deductive system, a play whose script is the laws of Nature? In this context, what are we to make, ultimately, of the quantum (digital) behavior of matter and energy, and perhaps of space and time?

The advantage of digital communication is that a message can be repeated indefinitely without accumulating error. Its drawback is the initial built-in error due to chunking. To put this another way, digitation has the advantage that only what is intended is transmitted. But this is also its disadvantage, because in one situation some information is gladly discarded as noise, while in another the same lost information may represent a dangerous insensitivity. Digital processes chunk information in order to minimize the complexity and uncertainty of the situation, and therefore the amount of information. The stress load of response is reduced by streamlining the representation of the world. But in the long run the stress load is increased if the oversimplified map corresponds too poorly to the territory.

An analogic mode of response may be adequate in equilibrium situations, when there is ample time to adjust perfectly to incremental changes in an environment that is basically stable. In preliterate societies, for example, cultural knowledge is distributed, global and ubiquitous. Everyone is a competent and relatively interchangeable member of the collective; change is normally very gradual. But in crises—or in rapidly changing societies, such as our own, which are virtually in perpetual crisis—the analogic mode of thought cannot deal with the flux of information. Cultural knowledge must be ever more specialized, chunked, processed in hierarchies. Specialists are needed to deal competently with increasingly delineated areas. The general public must be content with slogans and clichés, with highly simplified popularizations of the experts' models. This poses a great problem for democracies in the modern world, which are supposed to be directed by the will of the people and not by technocratic elites. Digitation is supposed to reduce problems of noise in information flows. Ironically, centralization, chunking, and digitation can give rise to social crises that are a further form of noise, and subject to runaway feedback.

22. *Time and Space*

Time is the means by which we first have our cake and then eat it. Space allows things the freedom not to be themselves. A physical world is an arrangement in which the logically impossible is permissible in fact. And what accommodates the difference between a physical system and a logical system is time and space.

An example of a logically impossible statement is "If X, then not X". This says that if something is true, then it is false. A logical system cannot be consistent if it contains such contradictions; there can be no reliable order in it. Similarly, if X is a physical object, rather than an assertion, X cannot both exist and not exist—that is, *not at the same time and place*. But it is entirely possible for X to first exist in a certain place and *then* not exist there at a later time—either because it is destroyed (transformed) or because it is has moved. Whereas logic is timeless, physical things move and change.

Time is a device with which contradiction can be lived out in an orderly way. Contradiction becomes oscillation when played out in time. "If day, then not day" expresses the diurnal cycle. "If up, then down" describes a bouncing ball. A logical contradiction is endlessly recursive but timeless: "If X then not X, then not (not X), then not (not (not X))...", etc. There is a static, timeless standoff between X and its opposite. They are two equal opponents in an eternal stalemate. A *physical* system, in contrast, accommodates self-contradiction by allowing one side and then the other, sequentially, to have its moment on the stage. Instead of cancelling each other out, one transforms into the other. In a physical system, implication—the *then* of logic—becomes cause, the *then* of time.

Indeed, time is measured by regular oscillation: the ticking of clocks, the planet orbiting the sun, atoms vibrating in crystals. How do we know how long the cycle of an oscillation is, or that a clock is regular? Solely by comparing it with other oscillators accepted as reliable. But what about *those*…? The problem is similar with space. How do we know how long a foot is, or that all the inches in a foot are equal? By comparing with other rulers or measures of space. Intuitively, we judge an interval of time or space by what other processes transpire during that interval of time, or what other distances fill up that interval of space. The interval has no intrinsic size apart from such references. Without them we can determine neither the magnitude nor even the nature of the interval (whether it is space or time)!

Just as time is the physical counterpart of logical implication, space is the physical counterpart of logical separability. If two things are logically, or "numerically," distinct, then by definition they are two and not one. (Two things *are* one when there are simply two references to the same item.) It must be assumed that any thing is identical to itself. Certainly there could not be any more blatant contradiction than the assertion "X is not X"—which says that a logical element is *not* itself. But in a physical system, it seems that two elements *can* be identical in every way except their spacial location. We are used to thinking that apparently identical objects are nevertheless subtly different, if examined closely enough. But it seems that, on the subatomic level, two particles may indeed be qualitatively identical, while

numerically distinct because of their separation in space. It is therefore space that allows for the multiplicity of particles, while in every qualitative respect there may be only one such particle. Like the answer to a riddle, X is not X *when it is somewhere else*.

Rather more poetically than scientifically, one could spin a story about the beginning of time and space, before which there *was* but one particle, a "singularity" infinitely dense (and apparently unstable). At the Big Bang its inner contradictions ceased to be potential (logical) and became actual (physical). The singularity became a multiplicity, and logical possibilities became physically distinct realities, creating the spacial extension and the mutual distances that their properties required. Space itself exploded under the pressure of logical contradiction, precipitating "particles" as fallout. The cosmos unwound from a very tense oneness, where everything was timelessly present together, to a structured world in which incompatible possibilities could systematically develop in time and space. Because the universe did not expand into a preexisting space, like a conventional explosion, all the matter and energy in cosmos is in some mysterious sense still together at one point, including our separate bodies.

23. *Formalism*

The fundamental human power is the power of the word. It is the power to create by decree. Nature is what it is; and since we did not make it, it remains ambiguous and we cannot know its truths with certainty. About things we have *defined*, however, we can have certain knowledge. This is the beauty of deductive systems. We invent some imaginary things and some rules specifying their behavior; then we can predict with absolute certainty how they will behave. The elements and rules of a deductive system, unlike natural things, have no reality of their own, no ability to behave other than how they are allowed or compelled by the rules. They are only what we have defined them to be.

Deductive systems contain explicitly only what has been defined for them and, implicitly, the logical consequences of these definitions (the rules of logic have also been defined). This gives them a self-contained, tautological, and "digital" character. Empirical facts (about the analogic world) do not have the formal validity that makes the truths of logic seem irresistibly necessary. This is why scientific theories are of a different nature than mathematical theorems; the latter can be formally proven whereas the former (being inductive generalizations) can only be *disproven*. Mathematics has the precision of definition, while science has only the precision of measurement. On the other hand, logically necessary truths are devoid of information for this very reason. There is no news in them about the world, because they concern what is necessarily timeless. Viewed as moves in a game, they concern not the real world, but the artificial 'world' of the game.

Like a parlor board game, a deductive system—or *formal axiomatic system*—has playing pieces and rules defining possible moves, as well as a game space within which the actions take place. It consists of: symbols, formation rules, axioms, and transformation rules. Formation rules are the basic conventions governing how the symbols may and may not be strung together. Axioms are basic assumptions—the input of the system. Transformation rules are the algorithms by which expressions may be handled in order to derive theorems. Theorems are the conclusions or output of the system.

Like the rules of a game, the elements of a formal or deductive system are well-defined; they are true and precise by definition—simply because we agree to make them so. For example, even though perfect right angles and dimensionless points do not exist in Nature and cannot actually be drawn, their idealized counterparts, which are matters of definition, can be manipulated in thought with complete precision.

A deductive system hangs together purely on the threads of logical necessity until it is "interpreted" as a mapping of some portion of the real world—as plane geometry can be applied to map the physical properties of the earth's surface. Then its premises may appear to coincide with truths about the world, its logical structure to mirror the organization of Nature. But a formal system can also be viewed as a self-contained game, played according to arbitrary rules, which have nothing to do with

reality.

The essence of the concept of formalization is that intuitive and vague areas of thought are replaced with precise and mechanical operations. A procedure is clearly defined for arriving at conclusions on the basis of prior specified assumptions and agreed-upon rules of reasoning. That is, a method of proof is defined. This precludes "just knowing" the truth, and also allows us to distinguish *truth* from either belief or *provability,* which is the relativized or subjectified version of truth. The concept of truth implies something that exists in itself absolutely and unquestionably; but a proposition or belief requires a mind to formulate it and a method to prove it. If proof is held to be merely an inconvenient formality on the way to a truth that exists prior to and independent of any methods of proof, then reasoning and provability are dispensable, since the mind can justify its leaps by direct appeal to "intuition." (Or, indeed, in any way it pleases.) Certainly the mind can conceive a theorem or proposition as true or false before a proof has been undertaken, just as it conceives that objects continue to exist when out of sight. Theorems may be *intuited* as true, even though unproven in any known system, just as some physical entities may be suspected to exist, though not yet observed. But suspicion is not certainty, and intuition must be followed by rigorous deductive proof. Whether in mathematics or science, intuition is a valuable and legitimate *phase* of thought, which should go on to seek further demonstration.

In contrast to 'truth', provability is always relative to some particular formal system. By definition, it is decidable for all mathematically "complete" systems. Since proof *means* derivability from axioms, no system can prove its *own* axioms, which can only be accepted by convention or justified by appeal to some outside reality. Proof is always proof-within-the-system, relying on an conventional acceptance of the premises of the system, while truth refers to such a reality beyond the pale of a given system. To lie beyond a *specific* system, however, is quite different than to lie beyond *all possible* systems. We can always find and specify a formal system larger than any given one. This constitutes an environing "reality" in which the propositions of the given system may be interpreted as true or false. But our intuitive idea of truth is that it lies beyond *all* formality, beyond all limitation, beyond all human involvement. Existing in and of itself, truth is transcendent reality. Such an ideal of truth can be approached but never reached, because it can never be fully specified. No sooner does one attempt to circle the world-beyond-all-worlds, in which this absolute truth is true, then in the same gesture a meta-system is defined outside the circle one has drawn.

Mind has the same dual nature as mathematics: it can indulge in gratuitous fancy or its creative efforts can be hitched to practical activity within the natural or social spheres. The map as self-contained formalism is a thing in its own right, perhaps a work of art; it can also refer to a real territory. The mind is a map that structures the organism's relationship with the world; and it is also an arbitrary construct, a formalism, a game.

Cognitive systems do not model the world in the way that a relief map models the topography of a landscape, but rather in the way that a road map does: that is, symbolically. Mind (as opposed to brain) can be formalized in symbolic systems; and human culture as a whole is a symbolic system or intentional construct—an artifact, a language, a map. What is *not* such a system is Nature, the world-in-itself, the transcendent analog domain, which can never be completely captured in any formal expression. What exists can thus in principle be divided in two categories: artifacts, which potentially *can* be exhaustively formalized, and Nature, which (unless it proves to be an artifact) cannot.

Cultural institutions generally, like formal systems, are embodied in their physical "interpretations"—for example: cities. Human culture mediates our relationship with the natural world, in much the way that experience is mediated by cognitive models. This is true also of particular institutions. A business or a corporation can be viewed as an interpretation of a formal system. So can a science or a school, a religion or a church, a state or a government, an art or an opus, a language or a library or a book. Cultural creations are externalized cognitive models, with a dual existence as self-contained games and as metaphors that point to reality.

24. *The Propositional Nature of Thought*

A proposition is an assertion, a statement. More broadly, I believe it should be viewed as any intentional connection. While a statement asserts a fact, the concept of proposition here proposed more generally includes assertions that are nonverbal—for instance, acts of cognition. Thus, "sugar tastes good" is a proposition, whether we mean a statement in language, a thought, or the neuro-chemical response to certain soluble crystalline carbohydrates. A proposition is "digital" because it asserts a judgment or binary decision, if only to assign something to a category.

Propositions are intentional because they are always proposed *by* an agent, reflecting that agent's purposes. Propositions about structure (of a system or organism, for instance) may reflect the interests, biases, and mentation of the agent proposing the parts involved.

Propositional thought facilitates decisive judgment, to mobilize behind action. Its disadvantage is that it leads to an artificial simplicity, which may be unable to model the complexity of the real world. A proposition not only proposes the fact it asserts, but also disposes of the alternative facts it ignores. This, of course, is its job. But there are situations in which it might be desirable to keep access to discarded information—as when minority dissenting opinions are kept on record. Each proposition, each decision, leaves out something that could have a ponderous effect if ignored.

Digitalness is a basic characteristic of defined constructs: something defined either exists or does not, either falls into a given category or does not, is either true or false. The physical, however, is analog and ambiguous (at least on the macroscopic scale), in the sense that nothing is truly either/or, but always in between or both. A proposition forces a maximized certainty, if we believe it can *only* be true or false. However, like the digital clock, it can *create* uncertainty if we believe the situation can be perceived more finely than the proposition allows, or be qualified by other relevant information.

Intentional constructs of all sorts leave a shadow of indeterminacy, in the sense that it is pointless to ask, for example, how many times Juliet sneezed in the year before she met Romeo; such details are not included in Shakespear's script. Juliet, being a fictional character, rather than a real person, is thinly defined. If she ever sneezes at all, it is only at the times indicated in the play. If she seems to come to life for us as more than a cardboard character or a mere outline, it is because her author, and the skilled actresses who render her part, have sensitively exploited the ways in which the audience fills in her silhouette with their own feelings and imagination.

Imagination plays a similar interpretive role in science. A mathematical model is a script for the behavior of matter. It cannot of itself promise to make the "characters" comprehensible in terms of everyday experience. It must be filled in with images and metaphors borrowed from a conceptual repertory of more familiar experience—for example, the image of the atom as a miniature solar system.

Any cognitive system, including science has its own properties as a self-contained deductive system. It rides on the analog reality of the world in the way that a drama uses the primary reality of material stage, real actors, and props. This metaphor ultimately fails, however, because within real life there is no closing of a curtain with a return to a primary reality—except perhaps in death. In the theater, the fiction is made plausible through the real experience of actors and audience jointly. In cognition, it is the other way around: reality is interpreted through schemata that are mental constructs.

Thought, and cognition in general, are *partial*—in the dual sense of the word: mind can never grasp the totality of a situation, and what it does grasp reflects its biases. Natural reality, on the other hand, appears in principle unfathomable, inexhaustible by mind, uncontainable by any formalism. The models of science bear a relationship to natural reality similar to perceptual models. As with perception, there is always something which thought ignores. The mind has a tendency to deal with the residue that does not fit into its representational framework as King Procrustes accommodated guests who were too long for his standard bed. It prefers to derive everything deductively from some basic source of axioms—whether in mathematics, in science, in economics, or in religion. Justifications then grow elaborate and spurious in the measure that more effort is required to force reality into the preconceived mold. But the remainder that does not fit returns to haunt its creators sooner or later in the ongoing confrontation with reality. It might be a falsified theory, a scientific revolution, a religious reformation, a personal crisis,

political action by the oppressed, or resistance from Nature to technological control.

History has long played leapfrog with dualities. Matter and mind, thing and thought, real and ideal, object and subject, cause and intention, absolute and relative, serious and playful: such antinomies dance a pas de deux through time. They are as bound together as the poles of a magnet. They alternate as undercurrents which rise to the surface as philosophical, social, political, esthetic, and scientific fashions. Their dynamic together drives consciousness toward greater objectivity. If there is a need to relativize and subjectify experience and thought, by disengaging its formal aspects, there is a corresponding need to reengage on higher levels, to assert more adequate models, to reinvest in better metaphors and explore wider horizons. To leave a nest is to find oneself nested in a larger world.

The limiting nature of any system of thought excludes whatever is founded on other than its own principles. Beyond the horizon of knowledge lies wilderness and chaos—the Unknown. The wider world is greater than any civilizing order. The map can never be perfect or complete, and what it leaves out returns somehow to haunt its makers, gaining momentum until the tables are turned and the current order is eclipsed by its own shadow. History breathes through a dialectic of complements, a dance of yin and yang. Every value, thought, or gesture has a darker side, a shadow cast by its own positive but limited presence. These shadow truths must in turn have their moments up stage, must come forth as dominant realities. The shadow solidifies as the solid dissolves into shadow.

The propositional nature of thought guarantees a complement to any assertion. There is always the other side of the coin. The shadow is a proposition in its own right, an anti-thesis. When value is defined in terms of the dominant system, the shadow may appear as repulsive, degenerate, wrong, frightening, evil, etc. If and when the shadow becomes dominant, this value system will be reversed, the world turned upside down. What was evil will be embraced as good, and vice-versa.

Shadows are cast by localized sources of light, not only by the solidity of objects, even objects of thought. Thought and perception, like light, always come *from* somewhere. Awareness is habitually identified with a perspective, with the premises of some system, beginning with the body, and focused upon some object. As such, it cannot contain opposites, cannot encompass both the object and its shadow in the same breath, so to speak. It cannot illuminate areas hiding in shade. Even so, there are no intrinsically dark regions. Darkness is always relative to the source of light, to the mind's assumptions and limitations, even when these are "ambient" because universally shared. Shift the light and the shadow moves.

There is an intimate relation between negation, or logical contradiction, and processes that oscillate in time. Once in motion, for instance, inertia carries a pendulum past its natural resting place, until its motion has been entirely converted to potential energy through the resistance of gravity. It then swings back the other way, gaining kinetic energy at the expense of its height, and so on. The logic of this behavior could be described as 'if yes, then no'. The same is true of an oscillating circuit: 'if on, then off'. Contradiction cannot be statically contained within the system, but gives rise to an alternation in time of one state with its opposite. Logical contradiction can be considered an abstraction, out of time, of physical oscillation. Conversely, dialectical cycles may be considered the acting out in time of logical contradictions.

But the notion of contradiction is relative to the system in which it is expressed. A proposition that is inconsistent with the premises of a given system may be compatible with an augmented version of that system. In other words, logical contradiction reflects as much a failure of the system to contain apparent opposites as it does an absolute fact. Expand the container and the contradiction may disappear.

Unresolved conflict moves in cycles; neither the thesis nor its shadow gains permanent ascendancy. The pendulum never escapes the restraining force of gravity, and gravity (in the ideal absence of friction) never finally overcomes the energy of movement. To the degree that historical forces similarly involve a mutually unstable equilibrium of figure and ground, they are also doomed to repeat in some sense. But to the degree that there is an overall net force acting on the system, it will be driven through a spiral-like evolution. History does not repeat exactly or literally, but thematically, on broad scales. There is a perennial resurgence of interests or points of view that are never completely exhausted

or definitively expressed.

The nature of thought renders it fundamentally incomplete. This generates an imbalance, lived out in time as dialectically opposed positions. In politics, for example, conservative and liberal views and governments alternate. In science, there is a shifting balance of idealist and realist currents; between mechanism and organism; between particle and wave or field. In the social sciences, there is an ongoing dialogue between behaviorism and humanism, nature and nurture. There are platonist and formalist mathematicians, intuitionists and logicists. In religion, there are fundamentalists and ecumenists, revivalism and orthodoxy. Art too has its polarized movements. Nature itself resolves "contradictions" by allowing one factor and then another to play out in time. There are cycles in weather, in species proliferation, in sunspots, in biorhythms, etc. It would be difficult to point to a natural phenomenon that is not subject to cyclical variation. And in human affairs there are little understood cycles of business and war, for example. There is a season to all things.

The key to dialectics lies in the synthesis of opposites, which depends on the transcendence of opposition. If there is an overall progress in human thinking, it is through growth in the container of thought itself—allowing seeming incompatibles to be assimilated to more powerful models. In this process, the system is freed from some particular oscillation when specific opposing forces are reconciled. The system can then go on to engage in a new or higher-level dialectic.

Throughout history there has been a dialectical relationship between the playful, inventive, *ironic* side of the human spirit and the *heroic*, serious, goal-oriented, realist side. The heroic is straightforward, straitlaced and straight-lined, earnest, passionately simplistic, concerned with content over form, tending toward militarism, conservatism and even fascism in its monolithic march toward a monumental and purportedly objective ideal. The ironic is subjective, liberal, witty, tongue-in-cheek, curvy and ornate, permissive, self-indulgent, and even aimless, tending toward vapidity or formalism empty of content. Subjectivity is the ironic foil of realism, the world put in quotes, bracketed as an esthetic object to play with in mental fields. We need a balance of the heroic and the ironic, and both are hazardous in isolation. We need the ability to make things real—to be earnest—but also the ability to play, to question, to laugh at our absurdities.

25. Language, Artifice, and Ideality

Technology expresses the ethos of its time and place, but something timeless as well. It has never simply responded to social need. Other motifs underlie the promise of abundance and mastery: social status and worldly power over people, god-like powers of creation, and, perhaps above all, the intention to separate from Nature. Whatever else it is, the cultural impulse, like the religious impulse it includes, reflects a desire to escape from the limitations of matter and the mad embroilments biology to a world conforming more to human liking and subject to human control. Man is the creature that strives to create an idealized (if far from ideal) world of its own, defined on its terms, a streamlined environment in which it is master. The very import of culture is to make an other-worldly, if man-made, sanctuary away from the horrifying contingencies of Nature and physical embodiment.

Culture is the human answer to mortality, disease, pain, helplessness, vulnerability, and biological conditioning. Mankind has staked out for itself its own territory, inwardly in imagination and mental or spiritual life, and outwardly in the world of artifacts, ritual and convention. Having long ago taken up residence in the inner realm—the mind, the spirit, *the ideal*—we have found our ways to have in eternity the cake we can eat only briefly and unreliably in the flesh. The significance of this territory is that it is *man-made*. It is a pro-active *story,* told by people, even when attributed to the divine. Homo sapiens is specifically the creature that builds a deliberate realm of artifice upon the platform of Nature. As invention, the essence of artifice is that it is ideal, simple, and controllable—unlike the natural world for which it substitutes. Culture not only removes us from the stigma of animal origins but also creates a proper habitat in which we are both author and master. (This is the deeper meaning of the ever-spreading urban technological environment.) The great evolutionary significance of language, culture, and

technology is not simply their utility for survival, nor their relative edge over animal vulnerability, but their potential to found an absolute transcendent realm, a divine status and home established in an artificial world.

We have seen that self-consciousness confers the ability to qualify perception through the awareness of one's participation: to look before leaping. It also gives access to an inner domain of thought and imagination in which actions can be abstracted, simulated, tested. All artifice begins in imagination, which is an inner workspace, a research and development laboratory, where elements originally drawn from perceptual experience can be combined in new and arbitrary ways—as happens in language and mathematics. These elements are idealized symbolic forms, in the way that words and geometric figures are. Both language and cultural creation reflect the prodigious inner world as a non-physical space, in which elements adopted from the external world can be arbitrarily re-combined. The inner world is plastic in ways that the outer world is not. In your mind's eye you can move a mountain or grow a horn on a horse. The mental world is hardly subject to physical laws, even if modeled on the physical world. (If energy is spent in thinking, in mentally moving the mountain, you are not obliged to *imagine* any energy at all required to move it). Imagination is a magical realm of freedom, where new things are possible. Magical thinking reflects and projects the freedom and plasticity of this inner world. Fantasy and magic are the forerunners of technology as means to bend external reality to human will. Thought and imagination, whether magical or technological, are means to manage the contingencies of the external world by asserting power over objects of thought in the inner world. Such magic is essentially the power of the word, the power to define.

The ability to consciously self-program, so to speak, means some direct control over experience. It also implies greater flexibility and a more adequate perception of the world. Obviously, these abilities alone would favor survival. More than that, power over an inner domain is the basis of our power to manipulate the external world and, so, to change experience indirectly as well. In this way Nature is redesigned by human consciousness, at least and at first in imagination, and eventually in fact, through technology. That the image, the idea, can be altered consciously before it leads to action allows action upon the environment to be far greater in range and sophistication. It also allows the world to be used as an extra-somatic memory, an external working display of thought, a blackboard. With imagination you can turn a forest into pencils. And you can use the pencil to do calculations that would send you to the moon.

Though physical and external, the human world is always and primarily symbolic—a world of signs, meanings, ideas—the same milieu that is created within language. Language too is artifice and idealizes. Consciousness, language, and material culture are three aspects of a single impulse to create an alternative, intentional, ideal world. While the organism lives in Nature, the conscious human subject lives primarily in this human realm. This means first of all in the subjective domain of perception, feeling, imagination and thought; secondly, in the social world facilitated by language; and finally, in the constructed environment of artifacts. All of these constitute the human world, as opposed to Nature; mind, as opposed to body; city, as opposed to wilderness.

While our first home away from the home of Nature is the inner subjective space, this would never have developed had we not been highly social, language-using creatures. Right from the start, mankind's immediate environment was composed of other people and their communications, and very likely subjectivity was highly favored by the need to simulate or second-guess the motivations of others. With self-consciousness behind it, it is language that has made the human species and the human world. Ideas and ideals found their first and immediate expression in language. The process of abstraction involved in idealizing is essentially linguistic; through it words acquire a flexible autonomy independent of concrete referents or contexts. The grammatical "space" of a language is analogous to the interior mental space. Language is an instrument, and perhaps the model, for the creation of an inner world of possibilities, freedoms, and idealized entities unfettered by the constraints and complexities of the external world. It creates the possibility, and the precedent, for unlimited arbitrary combinations of elements, which is the key to culture, technics, and mathematics.

While the subject-predicate distinction is fundamental in language, it transcends what we see in Nature. While creatures move and act, and actions seem to have agents, actions and things in language are detached from each other as abstractions. Though sometimes violating common sense, any action can be ascribed to any agent without violating grammatical sense. Grammar suggests the ability to combine abstracted elements in arbitrary ways. A basic feature of the inner world generally, this ability is a prerequisite for technology, which rearranges matter in ways that are arbitrary as far as the design of Nature is concerned.

Language is a major source of thought as well as its product. It is probably useless to ask which came first, thought or language, for they are mutually implicating and practically inseparable. The *formalization* of thought, of course, follows language, both logically and historically. Syntax is the archetype of formalization. The ability to manipulate objects of thought in formalisms, such as geometry, and in quasi-axiomatic systems—like music, scientific theory, and magic—mimics the manipulation of words in language far more than it does the manipulation of physical objects in space. Language is the prototype for all forms of creative expression, including production of material artifacts, because all creative expression involves the manipulation of formal elements and is intersubjective in the way that language is, translating inner image into a publicly accessible carrier of meaning.

26. *Simulation and Replication*

Representations of natural systems—for example scientific models—are products of definition, of imposing formal structure on the world-in-itself. Any definition is selective and intentional, an outline. Our *concepts* of physical reality, however, are so imbued with the ideal of objective realness that scientific *description* of Nature can only with effort be distinguished from the reality it describes. We tend to confuse real and deductive systems. What is meant by physical reality (or Nature) is therefore often ambiguous. Nature, as the world-in-itself, must be distinguished from the theories and whole sciences that describe it, and which are human creations.

Science, as an intentional construct, may or may not be a completely formalizable system; but Nature *cannot* be unless it is also an intentional construct. Physical processes appear to be well-defined and rule-governed. But it is a mere assumption—and wishful thinking—that they can be exhaustively codified or completely represented in some deductive system. Moreover, physical processes may *appea*r rule-governed for the same reason that they are recognized at all: because they have already been confused with intentional constructs, which are finite, simple, rule-governed, and well-defined. Nature, on the other hand, may not be well-defined, and may be unfathomable if it is infinitely complex.

Einstein marvelled that Nature is comprehensible at all and saw therein a mystery. Perhaps he viewed Nature as ultimately inscrutable. Ironically, however, the deductive/reductive approach views Nature not as a mystery but as an unnatural thing, an artifact, a machine that is knowable *in principle*. Mechanism is still the dominant metaphor in Western thinking because, like all myth, it assimilates the world to human definitions and artifacts: we can only truly understand (and control) what we have made ourselves. Mechanism is idealism at work, because it re-creates everything on conceptual ground, mirroring the mechanics of language and the power to define. The mechanist view of Nature imposes the grossly simplified ideal of the isolated system, with its artificially defined parts, upon the complexity of the world. Dominion over Nature is achieved first by defining its parts, just as the creatures were named by Adam in the biblical paradise.

This aspect of mechanism reflects the religious view of the cosmos as an intentional creation. In 1802 theologian William Paley proposed an argument for the existence of God that has since been known as the "watchmaker argument." If you were out on a stroll in Nature and you came upon a watch lying on the ground, he reasoned, you would think of it very differently than of a stone lying there. It stands out as having a rational, intentional design, and must therefore have had an intelligent maker and a date of manufacture. In other words, it is clearly an artifact. This is perfectly reasonable, so far. He then goes on to assert that the works of Nature similarly bear the marks of design that would indicate an intelligent

creator. The stone and the ground itself are viewed as artifacts—a non-sequitor that *un*reasonably projects human intentionality back onto Nature.

The Creation, inherited by science from religion, was by definition an intentional act. Arising by divine decree or fiat, the world is therefore an artifact invented by the Creator, embodying His thought and will, just as a machine embodies the thought and will of its designers. This radical change from the cosmology of the middle ages hinged on the notion of *law*, a word with two quite distinct meanings. These meanings, of "regularity" and of "decree," are merged in the Renaissance conception of the laws of Nature as divine decrees. Nature was God's invention, and its regularities manifest his will. The goal of science was initially to discover and decipher divine law and purpose.

Now, a machine is a paradigm example of law-as-regularity. Thus the idea of mechanism had a divine imprimatur. What was sacred was no longer the physical creation itself, as it had been in pagan religions and in early Christianity, but the divine blueprints, so to speak, as revealed in scientific law. This shift in values toward the mental and abstract had the major consequence of freeing up the materials of Nature for human exploitation. It is no coincidence that mechanism emerged at once as scientific model of the universe and as inspiration for the Industrial Revolution, alike from the womb of Christian idealism, which considered the natural world to be an intentional construct. The Newtonian world machine is rationally comprehensible because it is conceived as an artifact in the first place, inspired by literal machines, which can be understood because they were fabricated. However, what constitutes natural reality is precisely that it was *not* created by anyone at all, but is found in an incompletely known state.

I would argue that *only* if philosophical idealism were literally true could Nature indeed be completely modeled by deductive systems, scientific theories, or computer programs. This would be the case, for example, if Nature turned out to be nothing other than a simulation (in a brain? a computer?), or if it were literally a divine creation, as Man supposed for ages. In other words, unless the cosmos itself happens to be artificial, it *cannot be exhausted by thought nor thereby reconstructed as artifact.* Only if reality is in fact ideal, and not real, is the faith justified that it *must* be possible to map it thoroughly with other idealizations. Only if Nature is *unnatural,* in other words, can we assume that it can be perfectly comprehended, let alone perfectly simulated! Nature is comprehensible when it is viewed as an artifact, because only that which has been *made* is guaranteed to be commensurable with reason. Only artifact is truly acceptable to a creature who derives its identity from making things.

What escapes deductive systems and the machine metaphor—what might be called the 'mechanist stance'—is Nature's rich and complex interconnectedness, its totality, its very realness. Mechanisms, and the mechanical laws governing them, are simple, idealized abstractions. In contrast to the short while humans have been attempting to fathom Nature, Nature herself has had eons to spin forth intricate and inscrutable ways. And nothing is more elusive than the living organism. Machines are human artifacts (or constructions of other machines), whereas organisms grow by creating themselves. Machines have well-defined parts, whereas the "parts" of organisms are fictions. Machines are designed to produce an output for a given input; organisms define their own functioning. Machines obey linear chains of cause and effect; organisms use circular feedback and their functions are to some extent distributed over the whole creature, so that one part may take over for another. Organisms are self-organizing, self-repairing, self-reproducing and self-defining systems. Machines (so far at least) are designed and constructed for human purposes, to produce something other than themselves. The identity of the machine depends on the categories of the observer, whereas the organism maintains its identity (so long as it is alive), however it is perceived by humans. To view Nature reductively is to abstract, simplify, define, streamline, and idealize reality; to reduce it to an essence, formula, or schematic outline; in effect, to view it as a machine. The concept of the well-defined, closed, reversible system, which is the cornerstone of 19th century science, is the conceptual equivalent of a machine. It is a deductive system, translated into physical terms.

Two systems are believed structurally or functionally equivalent when they are considered to be structured alike or to behave alike. However, likeness is inevitably selective. A set of characteristics is

abstracted to express a formalism of which they are both examples. They may differ, however, in a potential infinity of other details, which are treated as irrelevant and indeterminate with regard to the definitions of the formalism, supplied by some agent. This applies as well to the definition of equivalence itself. It is an error to regard two *natural* systems as structurally or functionally *identical* simply because they are indistinguishable when viewed as examples of a common deductive. It is an unwarranted assumption that structure can be exhaustively specified at all, and a further assumption that function resides in structure.

These considerations should lead us to regard the concept of scientific modeling with some reserve. In particular, *simulation*—when it implies only selective functional equivalence—must not be confused with *replication*, which means producing an identical structural copy. Moreover, if the artificial replication of a natural system is to proceed by first exhausting the being of the original in some formalism—that is, by fully axiomatizing it in some "blueprint"—then it is in principle doomed to failure. Assuming that Nature is real, and not an artifact, it simply resists being caged in such definitions.

Simulation probably originated with mimicking other creatures. It is an aspect of the human appropriation of natural phenomena. Tools such as stone blades imitated specialized adaptations such as teeth and claws; spears imitated horns, etc. In modern times, simulation is the general idea that a natural process can be reverse-engineered for human use. By taking it apart, at least in thought, it can be reconstructed as an artifact analogous to the original process.

While the flight of airplanes, for example, imitates the flight of birds only in a metaphorical sense, a model airplane that flies does qualitatively simulate a real airplane because it *is* a real airplane, though reduced in size. A model bird or insect, however, even should it fly, is not a real one; it is not a natural organism but an artifact like the airplane. The defining principle of the Universal Machine (computer) is that it can simulate any other *machine* exhaustively. But whether a machine, program, artifact, or any other intentional construct can exhaustively simulate an organism, or *any* aspect of analog reality, is quite another question.

It is the characteristic "chunking" involved in language and thought alike, whereby a rose is a rose is a rose, that deceptively makes perfect simulation seem plausible. But there are many varieties of rose and every individual flower is unique. The baseball player and the pitching machine are both called pitchers, but the machine only crudely and metaphorically imitates the man, no matter how accurately it hurls the ball over the plate. It is a travesty to call these the "same" action. When a "piece" of behavior seems to resemble another action, it is not being compared to the physical reality itself, in its potentially infinite detail, but implicitly to a common defined protocol ("pitching"). Similarly, the essence of the being or structure of a natural object may be falsely thought to be captured in an algorithm or blueprint for the construction of that object as an artifact. This is the delusion on which standard top-down Artificial Intelligence has foundered.

Common sense recalls the differences between the living person who throws the ball and the automatic device, and the differences between the intricate human action of pitching and the mechanical hurling of the ball. The concept of simulation, however, rests on obscuring such distinctions by conflating all that can pass semantically under a given rubric. The algorithm, program, blueprint or formalism is the bottleneck through which the whole being of the phenomenon must pass in order to be simulated. One thing is said to simulate another when they both embody a common formalism. This can work perfectly for two artifacts such as the airplane and its model. They are in principle two alternative constructions from the same design. It is a profound error, however, to think that the being of a natural object or phenomenon is exhausted in a "design," which is then mistakenly believed to be a blueprint for its production, just as the aeronautical engineer's design captures the essence of the flying machine and becomes the blueprint for its production. The natural object is a found object, not an invention constructed from design. It does not come with a blueprint, which is imposed after the fact through a structural analysis that can never be guaranteed complete. Such analysis proposes, instead, its own completeness as a deductive system. The mechanist fallacy involved in reverse engineering of natural systems and organisms is the belief that it is possible to perfectly replicate a natural thing by first formalizing its structure and behavior and then constructing an artifact from that design. The artifact *will*

instantiate the design, of course. But it will *not* duplicate the natural object, any more than an airplane duplicates a bird. This is why "beaming up" à la Star Trek will never work as transportation. At best it might be a sinister way to replace a human being with a simulation!

A computer program (animating an android, for instance) might simulate human behavior in a highly detailed way, even at the level of involuntary physiological response. But the basic premise behind simulation involves a funnel through which all behavior must pass: the digital program that is its formalization. In other words, it is assumed that all aspects of an organism's behavior as a physical system can be reduced to the equivalent of a message in Morse Code. While this is merely an assumption, it is the working hypothesis of artificial intelligence. The heuristic of the Turing Test—that if we can't tell the difference then it makes no difference—becomes a technological fudge factor, in which subtle differences and their potential consequences are ignored for practical payoff. Subtle differences, however, may be ultimately crucial. Even when they are not catastrophic, they may lead to false expectations, as in the pipe dream of "uploading" human minds into computers, where they can supposedly live in simulated "realities" or, alternatively, be "downloaded" into cloned or artificial bodies.

Computer simulation rests on the principle that any analog reality can be approximated by a digital program to an indefinite degree of perfection. This is a useful principle for technologies whose intended product is not a truly exhaustive replication of an analog original but, rather, a noiseless flow of information, providing a subjectively satisfactory representation in the consciousness of the human user. The fact that digital sound subjectively passes as an analogue of the original does not mean it exhaustively replicates it. In other words, simulation is useful to create entertainments and virtual realities for eyes and ears that cannot (or do not want to) tell the difference, but not for exhaustively replicating reality. This distinction, unimportant to most audiophiles, becomes essential in considering futurist technology. Only laxness about it permits cyberspace fantasies of reverse-engineering brains, downloading minds, or fabricating artificial creatures, where the motive to create a useful artifact is confused with the unconscious drive to create life and reality itself. More broadly, this confusion serves the idealist mandate to transcribe reality as story, thought, image, simulacrum. A simulacrum, however, is not a duplicate, except where artifacts are concerned: a digital file may be duplicated exactly, but it does not duplicate the analog original. This may be relatively unimportant for sound recordings. But when the purpose is to duplicate natural entities, it makes all the difference in the world.

The discovery of the Mandlebrot and similar sets has stimulated a renewed enthusiasm for the promise to capture in simple algorithms the essential complexity of Nature. The patterns generated by computer from these formulae, however, are but relatively sophisticated simulations. If Nature is indeed infinitely complex, there is no hope of capturing it in any algorithm whatever, even one capable of generating endless detail of the self-similar sort. This does not prevent such simulations from being highly useful, as they have already proven to be in computer-generated films and image enhancement or information retrieval. There is a danger, however, when *generated* images are passed off as real. Already, for example, popular awareness of astronomical objects has been altered by computer images. There is no longer a clear distinction between literal photography and computer generated "artist's conception." The scientific and legal danger is no less. The validity of image enhancement lies in the assumption that its interpolations are veridical. They may be valid enough for entertainment purposes, while perhaps not for scientific or legal purposes.

The humanly created object is in principle predictable and exhaustible, while the natural found object is in principle unpredictable and unfathomable. The apparent predictability of the deterministic universe derives from the success of thought at simulating reality. A deterministic system is actually a deductive system, superposed upon Nature. It is an idealization, with no guarantee of perfect fit. While products of definition have an inescapable order, Nature may have an inescapable chaos, a residue of randomness that eludes any degree of imposed order. At the deepest level, this is the dilemma of the finite part trying to grasp the infinite whole. Any isolated system is artificial in principle. The real and rich interconnectedness of all things, on the other hand, may imply that ultimately the cosmos can only be

considered as a unified system, a whole.

The problem of exhaustively modeling real systems amounts to finding an algorithm to express a random sequence. A machine (that is, a deterministic system) can neither generate a true random sequence nor be counted upon to find an expression for one that is shorter than the random sequence itself. This is known as the Halting Problem, which states mathematically that there can be no finite general program to examine any other program, in combination with any input, and thereby decide whether the result would be an infinite loop, such that it would never stop if started. It describes, in effect, the relation between formalizations and the real systems they model. Related to Gödel's theorems, it sets limits upon the ability of mind (or computer) to formally describe reality. In contrast, the Halting Problem *can* be solved for machines with finite memory; machines, though not Nature, *can* be perfectly modeled. This is the real motivation behind the mechanist philosophy, behind the idealist faith that Nature should be approached as a deductive system. The only certainty lies in deductive systems, in the truths of logic. This brings to mind a Sufi story, about the good Mulla Nasrudin, who was found on hands and knees searching for his house keys beneath a street lamp. When a friend shrewdly asked him where he had *lost* the keys, he replied: "I have no idea; but here, at least, there is some light."

The implicit idealist faith behind the notion of perfect simulation is that each and every property of a thing can be represented with precision. A *property*, however, is a human construct, an assertion that disregards a potential infinity of other assertions. The collected properties of a natural thing do not constitute the thing itself, although they do constitute an artifact. Any list of we could assign to an object is finite. It cannot exhaust the being of a real thing, which has infinite—or at least indefinite—properties, but it could exhaustively describe an artificial thing. Using Shakespeare's folios as the authoritative reference, we *can* know all that there is to know about Juliet, or her fictional balcony, or her thoughts about Romeo. But that is not very much.

When from the outset the object is *taken* to be a collection of properties, however, it is mistakenly assumed to be exhaustible. This is the fallacy of the mechanist stance. Once the circularity of this reasoning slips by, one predictably fails to see any difference between reality and its "perfect" simulations. We see only, then, what is visible in the circle of light provided by deductive systems. In that idealist frame of mind, it is tempting to see the whole universe as a simulation (of what, one wonders?), and one's own existence as nothing but numbers being crunched in a vast computer (like thoughts in the mind of God?). This is nothing other than a technologically updated version of Descartes' forced solipsism, in which the "evil genius" who counterfeits your sensory input is replaced by a computer. The answer is the same as to Descartes' skepticism: there *is* a way to tell the difference between reality and simulation, for no simulation *can* be perfect. The meaning of realness lies, in part, in externality and independence of mind; but it lies also in Nature's unfathomable richness of detail.

27. *Artificial Intelligence, Artificial Life, and Nanotechnology*

In 1936 Alan Turing conjectured that any effective procedure that can be accomplished by human problem solvers can be formalized, and therefore duplicated by a machine. Hence the idea of the universal machine: the computer. But this conjecture seems to involve some circularity, since an "effective procedure" is already a member of the restricted class of human activities that can be formalized. It may not be true, however, that *all* human thought and behavior, including cognition, is formalizable and therefore replicable by a machine. The Turing hypothesis is nevertheless the cornerstone of artificial intelligence. Ultimately, it is the notion that the functioning of a human brain can be understood and abstracted in such detail as to be exhaustively simulated by a computer.
But the grand dreams of classical AI have essentially met with failure, and we should be skeptical about the idea that intelligent behavior can be completely formalized, abstracted out of its embodied context.

In addition to the fallacy of perfect simulation, there is the danger of indulging a pathetic fallacy: by judging the purposiveness or intelligence of a system's behavior in human terms. For an action to be considered genuinely intelligent, it would have to be evaluated in the terms of the intelligent system itself

(which, of course, must *have* such terms in the first place. An organism is the cause of its own behavior, a machine is not. A computer program can indeed simulate selective human mental processes and actions, and the optimism of classical AI theorists was founded on their success at simulating formalizable systems such as logic. But such success can mislead us to the conclusion that an information processing system necessarily processes "the same" information as a brain, embodies the same structure, performs the same operations. In some cases this assumption may be justified. These seem to be cases which involve aspects of human intelligence that have already been formalized, or for which effective procedures have been found.

Information is often defined as the number of binary decisions needed to reduce alternatives to one. This presumes a decision procedure for a finite set of defined alternatives. We must distinguish accuracy from precision, however. To be accurate, a model should be only as unambiguous as the situation it models. But this definition of information parallels that of proof of a theorem within a deductive system: each binary decision (bit of information) is analogous to a step in a proof, which narrows the margin between premise and conclusion, just as each binary decision reduces the alternatives by half. While precise, it is not necessarily accuracte, since it presupposes a perfectly defined situation that the information is about. Beginning with such a concept of information, it is easy to suppose that the world can be informationally exhausted in formalisms.

The same considerations apply to behavior as to structure, function, and information. Qualitatively simulating behavior is not the same as fully replicating it. Formal models, and mechanisms constructed to embody them, capture only aspects of the behavior of natural systems. Behaviorism *assumes* an organism to be a mechanism and, hence, the embodiment of a deductive system. It selects aspects of the totality of the creature's activity to characterize as its "behavior."

The limits of simulation hinge on just how much of human (or any) activity can be formalized. We might attempt to draw a distinction between the organism's activity as a causal system and as an intentional system. However, such a distinction will ultimately be arbitrary, since an organism is necessarily both an intentional and a causal system, which are but two ways of viewing the same thing. It might be argued that an organism, though an artifact *of itself*, is nevertheless an artifact—and, so, potentially formalizable as such. I believe, however, that there is an essential difference between the autopoiesis of "self-organizing" systems and a treatment of the organism as divided within itself. To regard of the organism as an agent that creates *itself* as a product somehow separate from its own agency, is simply a projection of the self-conscious dualism of human beings. No such hierarchical division occurs within the organism, or anywhere in Nature, where everything is mutually interacting in multiple feedback loops. An intentional system is "carried" on a richer physiological substrate, just as a message is carried on a physical signal. But it would be a mistake to separate the message from the medium, or to think of the organism as somehow a message sent unilaterally to itself, or as a program written by itself to direct its own activities. There remains in such ideas a residue of the top-down thniking that has plagued Artificial Intelligence from the outset. In any case, the information contained in the structure and functioning of the organism is information "for" the organism, not for an outside observer.

Even though the organism must be regarded as an intentional system that is a product of itself, it is also a causal system that is a product of the entire universe; as such it will in principle elude complete formalization. It may even elude adequate definition, since "the organism" is only separable from embedding environments by convention. It is merely an assumption that the intentional system, supposed to characterize the organism, can be correctly identified, much less codified, by an outside observer. The elusive wealth of the physical organism as an analog substrate must be discarded as essentially irrelevant to such a project. But this amounts to disregarding the reality of the organism except as it conforms to our preconceptions. An organism is its own author. Any intentional system proposed to account for the organism's behavior will in truth be authored by the observer. It will be some theoretical construct—a deductive system, a human artifact, imposed upon the organism.

Some critics of AI believe that intentionality is an inherently biological phenomenon. Perhaps they mean that it is an inherently *embodied* phenomenon, a product of natural selection—which I certainly believe is true. No artificial system has, as of today, its own intentionality—no doubt because

no artificial system has come close to duplicating the conditions of embodiment that provide the intentionality of organisms. Little effort in this direction has been made in traditional AI, work proceeding instead on the simulation of fragments of behavior.

Developments in the new field of Artificial Life (AL) may change all or some of this. If physical, rather than virtual, artificial systems can be evolved through circumstances equivalent to natural selection, then the conditions of embodiment could be met, resulting in systems possessing their own intentionality and hence their own intelligence. However, by definition, this could *not* be a process entirely within human control. Human purposes would have the same relationship to artificial evolution as they do to natural evolution. In other words, the process could at best be guided, as it presently is with natural organisms, through applied selective breeding or genetic engineering.

At present, the conditions of embodiment are simulated in computers, through programs that simulate both the defining properties of life and the process of natural selection. So far, these artificial organisms are *only* simulations, principally because embodiment itself is only simulated. These are virtual, not physical, creatures. The selection rules are arbitrary inventions of programmers, not conditions in the real world, nor resulting from a competitive environment of other self-defining entities. Fortunately, we do not yet have the capacity to construct self-replicating *physical* machines that could be turned loose to compete with each other (and with us!) for survival. But this capability may not be out of reach. Developments in microtechnology may soon render it possible to physically embody the virtual simulations of AL, resulting in true artificial organisms, and a new chapter in the history of the biosphere. Commentators on this scene already warn of the inevitability of superintelligent machines and a "singularity" in the exponential development of future of technology, a point of no return where technology is no longer within human control.[15]

The mechanistic view of Nature applies the machine metaphor to the cosmos at large as well as to identifiable subsystems. It is sometimes assumed, in other words, that the universe itself is a machine that can be analyzed into its "true" parts in the same unambiguous and reversible way that an engine can be assembled and disassembled. I believe that this assumption arises from the wishful thought that confuses natural systems with deductive systems.

Natural systems might inspire the design of a machine, but someone designs and builds it. It has a finitely delimitable set of well-defined parts intended by its fabricators. It can be dismantled into these same parts by reversing the process of construction. In general, the machine will function properly only if assembled in one specific way. We can hardly say any of this with certitude about Nature, nor do we know what are the "proper" parts or design of an organism. Whether a design is sound may be analogous to whether a theorem can be proven in mathematics, and this should be a different situation from a scientific experiment, which is supposed to query Nature, as an unknown outside the defined system of the experimental set-up, for the answers. Testing the performance of a scientific theory, however, may resemble testing that of a machine more than superficially, if a theory is essentially a deductive system and the experimental set-up is what one philosopher has called a 'nomological machine.'[16]

What are the limits to mechanism as a metaphor of biological reality? Can a living cell, for example, be considered a self-replicating machine? If so, we are faced with the astonishing possibility of constructing or evolving microscopic self-replicating machines modeled on living cells.[17] If these little "factories" were controllable from the macro level, they might in their ensemble constitute *programmable matter* that can reassemble itself into any desired shape and function. They might also become a rogue new form of quasi-life, a new plague dangerously beyond human control.

The program of instructions contained within the cell for replication may appear to be an exhaustible formalism, since it consists of a finite structure. But can it be assumed that *all* information governing the processes of replication and development is contained strictly in the genetic "code"? Even in the case of a computer, by analogy, only part of the relevant information is contained in the program. The functioning of a computer is an interaction of software, hardware, and the human user's inputs. There is a similar figure-background relationship between DNA and its context, which includes the materials from which the new organism forms and chemical instructions in the environment or in the

body as a whole.

There is an essential difference between an organism and a machine, however sophisticated, that makes products specified by human intent and under human control. It is a misleading metaphor to consider a cell a "factory," for instance, since the principle product of an organism is *itself*. This is not only because it self-replicates, but because it self-defines; its whole endeavor is to maintain its existence and identity. As an intentional system with its own purposes, it is only incidentally an instrument of human intention. It is not a tool, except in the incidental way that Nature is presently exploited for human purposes. Aristotle and creationists notwithstanding, it was not *designed* for human use.

Nanotechnology proposes the complete control and redesign of matter at the molecular level. It assumes that living organisms can serve as model and inspiration for the creation of microscopic "factories." As long as the process of fabrication is understood in the conventional sense, we are at least on familiar ground. We already know, for instance, that manufactured products are subject to various imperfections, and that industry is wasteful and makes pollution. The biosphere, on the other hand, is a self-maintaining and self-policing system. By definition, there is no such thing as pollution in Nature as a whole; moreover, "imperfection" is genetically weeded out. There can always be unforeseen and perhaps undesirable byproducts of human enterprise, simply because thought is simplistic, while Nature is wholistic. But when we turn the issue around and *assume* the perfectability of thought, we assume that Nature owes it to us to conform to our ideas! If we wish to make factories that operate as instruments under human control—but *tiny and numerous*—then we can expect of nanotechnology more or less what we get from present technology, multiplied incredibly by numbers, speed and capacity. But any technology that truly imitates the autonomy of organisms or the wholism of Nature will necessarily be beyond human control.

How *would* molecular factories be controlled from the macroscopic level? It is one thing to imagine nano-factories that are autonomous, self-programming, self-modifying and evolving like microbes. Like microbes, these could only be controlled indirectly. It is quite another thing to imagine tiny entities fully programmable, top-down from the macro level. How would communication take place? Through radio transmissions, each nanite on a different frequency? Practical difficulties aside, might there be theoretical limits as well? Perhaps it could be argued that such communication between hierarchical levels already exists within organisms, so why not in artificial organisms? But this natural "communication" is no one-way directive from a central or external source. The cells of the body are not "operated by" the brain, let alone by the person. Rather, the organism, like Nature as a whole, is a continuous cycling of mutual influences, through many channels, bottom up and top down.

From the beginning of time, the vulnerabilities of the body and the fact of mortality have contributed to the desire to control Nature and even to create artificial versions of life. While science has proven a powerful instrument of human will, it has added to the injury of mortality a succession of insults to human significance. Perhaps the final nail in the coffin of human superiority will be driven home, ironically, by the advent of superintelligent artificial organisms. Whatever consciousness is, it will then be finally clear that it is not a uniquely human prerogative.

28. *Extraterrestrial Intelligence*

The intelligence with which we are familiar is limited to the life forms we know on this planet. The search for extraterrestrial intelligence (S.E.T.I.) has been called a "science without a subject," because by definition it is about something we have not yet encountered (flying saucer cults notwithstanding!) Outside the human cognitive domain—and the terrestrial creatures to which we consider ourselves superior—we have no reference to serve as a context within which to understand and situate ourselves as members of a cosmic community. We are looking for our counterparts somewhere else, but lack an external perspective from which to grasp even what we ourselves are. Perhaps our concept of intelligence is entirely provincial, in spite of the drive to abstract intelligence from its biological roots. Popular notions about ETI tend to imagine extraterrestrial beings in our image, partly

because we lack a general perspective or theory of embodied intelligence within which to imagine anything else. In this circumstance, speculation cannot advance much further than the science of Aristotle or the medieval scholastics. Even so, pondering the possibility of extraterrestrial intelligence—even when it seems to be stating the obvious—helps at least to organize thought about life and intelligence on *this* planet, and is a step toward a needed cosmic perspective.

S.E.T.I. concentrates on the earthbound detection of intelligent signals from deep space. Exobiology, on the other hand, consists mostly of the search for organic chemistry within our solar system, and particularly on Mars. A huge gap remains between these domains for want of a common theoretical basis. And that would be the entirely speculative study of what possible alternative forms of conscious intelligence *could* exist, and their alternative evolutionary histories. I can only offer some rudimentary thoughts about some constraints on this *terra infirma*. Admittedly, such speculation merges with the realm of science fiction.

To begin with, the human form and intelligence has evolved, and is maintained, as an integral part of a biosphere. Though obvious, this has several implications. First, intelligence is the dependent and fortuitous result of long evolutionary histories; it must remain compatible with the biosphere that gave rise to it. (This is relevant to the last factor in the "Drake Equation" for estimating the number of detectable galactic civilizations.) *If technological civilization is inherently inimical to the biosphere that supports it, it could only be short-lived*. The only way around this plain truth is the dubious assumption that intelligence can take over and manage the functions of the biosphere itself. In the preceding sections, we have encountered reasons why such a project would not be possible. The presumptuous idea of controlling Nature is tantamount to re-creating it, which would depend on the possibility to exhaustively codify it. This, we have seen, is not possible unless Nature happens already to be an artifact. This is not to say that an ecology of artificial organisms would not be possible. But such would evolve on its own, and not under human control, even if initiated by human beings. As an extreme example, organic life might be entirely or partially supplanted by an ecology of artificial life forms, perhaps more robust than their natural counterparts. While this ecology would have its own rules, technological intelligence within an artificial biosphere would undoubtedly face the same problems as it does within the natural one. One way or another, therefore, the advanced civilizations of sci-fi imagination would have to have resolved basic dilemmas currently facing us regarding their impact on their "biosphere," whether natural or artificial. These would be inseparable from what we identify as moral, economic, military and social issues. (Hence the significance of the question Dr. Arroway would pose to extraterrestrials in Carl Sagan's *Contact:* "How did you do it?")

Second, because evolutionary advancement means rising in trophic pyramids, an advanced life form will probably, like us, constitute a relatively minor part—at least in terms of biomass—of a very large self-sustaining ecological system. This is what renders unrealistic a barren landscape, populated only by intelligent humanoids, such as depicted in some science fiction. The biosphere has followed its own evolutionary parameters, in terms of which ascendancy toward the human form could only be permitted as the exception and never the rule. The great mass of life must remain lowly, for otherwise it could not support the higher forms. This parallels the probability that life only exists in isolated localities within vast reaches of a universe that is on the whole empty.

What does it mean to be at the peak of a trophic pyramid or food chain? Above all, it means playing by the prevailing rules of the biosphere: the "game of life." This means that organic intelligence is highly conditional. As part of the genetic heritage of the biosphere itself, we are here as *consciousness* only by virtue of our programming as *life*, standing on the shoulders of a vast heritage of present and past forms. The nature of our intelligence is conditioned, if not strictly determined, by our participation in that game. This heritage is animal: we exist as flesh that lives by consuming other flesh. Our minds themselves are made of flesh and we perceive and think with the intelligence of flesh. Only by recognizing the embodied and parochial (if not beastly) nature of our own intelligence will we succeed in creating a *general* concept of intelligence that is capable of embracing the varied constraints upon alternative embodiments and varying levels of freedom from those constraints. The flight from embodiment and the denial of animality that is such a general feature of human culture, more particularly

of our current technological civilization, has only impeded a truly general concept of intelligence.

Conditionality means living with priorities, the rules governing the body in the game of survival. The mind, as an extension of the body, has internalized these rules as *values*: good, bad, pain, pleasure, etc. A stone has a less conditioned existence than a plant, and a plant than an animal, because the conditions required for its existence are less stringent, simpler and fewer. It has no use for values. We can only speak of values, preferences and the qualities of experience where a being has internalized the conditions for its existence in order to be able to act on the environment so as to bring about and maintain those conditions.

Now, conditional existence is a different notion than physical embodiment. We might imagine forms of existence that are non-physical yet conditional: angels, for example, who might be obliged to obey spiritual rather than physical laws. Conditionality implies a basis for an intentional structure, as it does for organisms, reflecting structure in the worlds in which such beings exis. It implies mind—thought, perception, value, etc. Physical embodiment guarantees conditionality, but conditionality is the more general concept, which does not in itself guarantee embodiment.

In talking about extraterrestrial life, we are not considering spirits or beings from other dimensions, but possible *embodied* consciousness, which, like ours, has evolved within the physical parameters of a "game of life"—even if not based on the same, or any, organic chemistry. The emphasis here is upon the evolution through some series of lawful physical interactions. Such life, whether satisfying definitions of Earth biology or not, would have a history and a logical development analogous, if dissimilar, to our own. This would seem to imply evolution of the species form through natural selection operating upon individual units that reproduce and die. Note that mortality and natural selection over long periods are necessary features of sexual reproduction, without which the natural evolution of advanced intelligence seems impossible. Though mortality is the platform for the evolution of intelligence, such intelligence becomes (in the human case, at least) capable of contemplating alternatives. Moreover, it can substitute cultural for natural selection.

But we might still ask whether evolution by natural selection is a universal precondition for life, or merely a terrestrial fact. In other words, can there be a basis for the self-elaborating intentionality of a complex adaptive system other than its participation in a game of selection over generations? Must it be self-reproducing, which implies generation and death? Must it be part of a chain of being, which implies competing with, exploiting, and consuming other creatures? Perhaps an intelligent system could evolve merely through a process of self-organization that did not involve survival of competing genes, nor even the existence of multiple individuals. After all, the chemical elements, and compounds we are familiar with, themselves "evolved" through a non-biological process of self-organization. Could a self-organizing principle bring itself to a state in which it can consciously act upon itself? Perhaps life as a whole should be viewed as such a system, the coming and going of individuals and even species and genera being irrelevant on that scale.

Human beings are self-reflective. We must assume this also about extraterrestrials with whom we might expect to communicate, wherever they might be. Reflexive consciousness is a concommittant of (grammatical) language. We are compelled to assume that any race capable of physical space travel or sending signals into space will be reflexively self-conscious (even if they happen to be "artificial") and will use some language. If so, then we can also assume they will have similarly confronted issues generated by self-consciousness: mortality, ethics, individual vs. collective, transcendence of embodiment and its programming, and the footprint of their civilization upon their biosphere. To what extent will they have resolved such dilemmas? They will no doubt have a philosophy and perhaps spiritual aspirations, and may to some degree, like us, be befuddled by contradictions and regressive tendencies within a consciousness at once transcendent but still conditioned by the path of their evolution. They will experience the suffering of (bodily) existence, and may manifest inconsistencies in their behavior associated with the split between embodied programming and reflexive consciousness.

In order to produce a technological civilization capable of space travel, they will necessarily possess some appropriate way of manipulating their environment. For humans, it is the hand, in

conjunction with eyes that focus and an upright posture that frees the hand; but it could be alternative organs. While cetaceans may have an orally transmitted "culture," their intelligence does not manifest as material technology because of the lack of appropriate manipulative organs.

A technological civilization requires the cumulative effort of many individuals over time. This presupposes stable sociality, communication and information storage, language and planning. Information storage would not necessarily have to be external to the memory of the individual, especially if individuals somehow had super memory capacity or telepathic access to each other's knowledge. Of course, language, reflexive consciousness and (non-material) culture are possible without a technological orientation. Again, the cetaceans and elephants may be cases in point, with their elaborate communication skills. We should also bear in mind human cultures and subcultures that do not value technological development. Yet, if signals are to be sent across space, some means of directing electromagnetic energy seems to be implied, at least within the framework of classical and relativistic assumptions of physics. If a species communicated using electromagnetic signals (rather than sound), perhaps some natural means of amplification could evolve, without externally fabricated amplification and transmission devices.

The general considerations of biologists Maturana and Varela would apply to all cognitive systems satisfying their definition ("a system whose organization defines a domain of interactions in which it can act with relevance to the maintenance of itself"). Hence, an alien cognitive system would have an input from sensory surfaces and an output to effectors. Whatever else may be said about this system, it would act to maintain *itself*, through maintaining its inputs within certain bounds by means of its output. The action is effectively upon the system's own internal states, while the question of a real world, upon which the system acts, need not enter and is a concept in the cognitive domain of the outside observer. However, since *our* experience is of a real external world, we can only imagine that an alien mind would also project its own cognitive domain as "real." Like us, moreover, its self-consciousness would be a means toward self-transcendence and a more effective grasp of the world.

Cognitive domains can be ordered in terms of their comprehensive "objectivity." A creature is "superior" which has a more "adequate" model of the world. It is difficult to formulate a concept of objectivity without reference to an external world. Otherwise, all that can be said is that a cognitive system either works or doesn't. (If a species persists through geologic time, then by definition its cognitive system works!) Intuitively, we have a concept of intelligence that allows us to rank other creatures in the terms of our own cognitive domain, on a scale on which we place ourselves at the upper limit. Contact with a superior intelligence would force us to revise this scale, the upper limit of which would no longer remain within our comprehension. A superior alien consciousness might view things similarly, except that we would rank among the lower intelligent forms. It would measure our limitations in much the way that we assess the limited intelligence of animals.

What then is objectivity, and what are its *imaginable* extensions beyond our present cognitive domain? Objectivity is widely held to be the perception of the world itself, as it is, independent of the medium of perception. We have seen, however, that for any embodied creature this is a contradiction, and therefore a myth. We percieve *through* the body and its mind and personality, supported by the entire genetic heritage of the game of life, and by the entire cultural edifice that is its extension. Objectivity for us therefore means a relative freedom from the constraints and programming of this entire structure: Nature/body/mind/culture. Since we cannot disengage in any absolute way from these contexts, objectivity is at best an ideal, and at worst an absurd pretention.

Reflexive consciousness allows the transcendence of limitations. In our case, the value of this has been supported by biological evolution long enough for conscious self-direction to be conceived if not achieved. We might then expect advanced alien life forms to have taken charge of their destiny in some measure, and in some way to have transcended the social conflicts engendered by individual embodiment while retaining the genetic advantages of diversity. We could expect a transcendence of the entire survival-ethos we associate with embodiment and biological programming. There would be resolution of the existential problem of individual separateness—that is, of how to live in the material world in a way

compatible with individual freedom, but also to have achieved a state of complete integration of the individual unit of consciousness within the collectivity. This might imply an objective and benevolent consciousness, with an absence of conflict and ego based on competition for resources needed to sustain bodies. It might also imply a self-lessness that is foreign to our "individualistic" mentality. These have been ideals of human spirituality, and also of the humanistic intent behind technology. Some spiritual ideas may be regarded as precursors of ideals pursued through technological development; conversely, some technological ideas may be viewed as fulfillments of spiritual yearnings. (Indeed, in the West there is a close association between science fiction and esoteric spiritual teachings, both of which often concern themselves with the conscious evolution of humanity.) As cultural evolution has built upon the platform of biological evolution, so a consciously directed technological evolution could build upon the platform of culture as it has "naturally" occurred.

Some aspirations of science and technology—and of AI and AL in particular—resemble those of spiritual idealism: namely, freedom from the vulnerabilities, limitations, and biases of embodiment. In both cases, identity is viewed as independent of infrastructure. Liberation from biological programming is a spiritual ideal: freedom from identification with the body and the parochial quality of thought that permeates embodiment. A logical consequence of this transcendence is the archetype of objective loving kindness touchingly portrayed by an alien intelligence in the sci-fi film *Contact*. We have imaginatively projected this ideal, first onto a fatherly God and more recently onto superior alien beings. But perhaps it is also our own destiny.

As a product of the game of life, consciousness owes allegiance less to truth than to genes. Nevertheless, human cognition has transcended this allegiance, at least to conceive the possibility of reality or truth, independent of human cognition, which is the meaning of objectivity. We might hope for the same and more from a superior consciousness, which would have achieved yet greater independence of its own evolutionary origins. If the notion of objective truth is not self-contradictory, then advanced intelligence throughout the universe ought to converge upon it, simply because those conscious species *not* in harmony with reality would not be expected to survive the crises entrained by their own development. Meanwhile, on this planet, we appear to be engaged in an experiment to determine how far technology can outstrip wisdom.

Part Four: WORLD AND SELF

29. *A Conscious Relationship to Experience*

In a conscious relationship to experience one is aware of participating in the creation of that experience. To be aware of experience as such implies a subject/object relationship. In being aware of the object, one is simultaneously aware of one's presence as subject and of the indeterminate nature of the object. In a non-conscious, or pre-subjective, relationship to experience, there is awareness of the object alone, but not *as* an object, nor in a way that implies a subject. There is no awareness of the subject's contribution to the experience. Rather, there is total absorption in the experience, without awareness of *having* experience or of being present to one's own consciousness. In some spiritual traditions, this state of absorption is lauded as non-dual. In others (for example, the Gurdjieffian tradition), self-awareness is the ideal. In any case, it is my belief that conscious subjectivity has a crucial evolutionary role to play. It frees the subject from rigid programming, through the reflexive ability to transcend given boundaries. It also creates an inner domain in which imagination and creativity can freely play, in which ideas can be tested at little cost, and which serves as prototype or laboratory for creation in the external world.

While the subject has no place in the naive or pre-subjective mind, where only *the world* exists, as the subject comes further into being, so does the object. Our view of the world becomes more

objective as we become more consciously subjective. Claiming our share of responsibility for experience allows us to separate our contribution from that of the world. More of the cognitive apparatus comes into view; less of it is projected into the world as "reality."

The power to consciously choose a course of action requires a self to choose it. In the absence of a self, there is only the programming of the organism or of society, which cannot be distinguished from the reality of the world. In subjective consciousness, the point of view of the subject is increasingly disengaged from that of the organism and its social context. The subject is less obliged to see and respond to the world through conditioning, and more able to see the world for what it is and also to see the conditioning for what it is. The subject becomes more cognizant of its epistemic situation. But, along with this recognition comes a problematic relationship of self to the external world itself, which includes the body.

The evolving subject is potentially free from its identification with all contents of experience. What invests experience with urgency is the sense of reality with which it is imbued. This reflects the organism's survival programming. If that sense of reality is questioned, there may be little reason to feel compelled by it, which can be both a liberating and a dangerous situation for the organism. When experience loses its compulsory import, action may acquire a greater range and become a more voluntary play. On the other hand, there may appear to be no basis for action at all. One may literally lose all meaning and bearings. The ultimate implication of subjective consciousness is that the subject is freed even from its own apparent reality. The subject grasps that it is nothing but a fleeting glimpse of a shadow show. Even its own drive toward subjective freedom and empowerment is seen as part of the machinery entrapping it: another ephemeral object with which it can decline to identify.

Experience is the panorama, so to speak, of the subject. This perspective is literally a dimensionless point, interior to all experience. The pure subject is nothing other than this point. It is therefore not an object of any sort, nor any part of the panorama. Subject and object are totally disjunct. The object is seen from a perspective; the subject is the perspective itself. The perspective can never see itself, yet it may come to mistake some portion of its panorama for its own being.

In subjective consciousness there is a shift, so that the panorama is no longer naively interpreted as the world, but is also seen as an inner view that includes the workings of the cognitive system. The subjectively conscious mind recognizes all experience to be a product of both world and self. But this "self" is implicitly the cognitive system: part of the organism, part of the world. It is not the idealized and transcendent point of view, which remains interior to all objects of consciousness.

Not itself a thing, the subject is not any object of its own or others' experience. It is a point of view, never what is seen from a point of view. Awareness points *through* the "lens" of mind, which has its own focusing and distorting properties but is normally (unconsciously) transparent. It is possible, however, to look *at* the lens of mind, provided we understand that what we take to be the lens is itself a product of the lens, an image. In other words, the lens can be known through inference from its distorting effects upon experience and behavior. (We are not speaking of the brain, which is a physical object, but of the mind, which is not.) An analogy might be the black hole, which acts as a gravitational lens. It cannot itself be seen, but produces visible effects on the background of more distant objects, whose images are distorted or multiply refracted around it.

Though motor agent, the subject is not the body which performs actions; nor is it the thoughts or feelings that animate them. Though experiential witness, it is not the brain nor any entity or homunculus such as the soul. And since it is a dimensionless point of view, there is nothing but location in space and time to distinguish one subject from another, except through the (intentional) act of identifying with a particular cognitive system. Therefore, in an absolute sense, there is no plurality of subjects, which are qualitatively one. The idealist truth, of the unity of subjects, parallels the materialist unity of the cosmos as one undivided object or process. Just as objects, agents, and physical systems are recognized as artifacts of thought, so are subjects recognized as convenient fictions.

What is the role of this idealized conscious subject in the behavior of a cognitive system? I have

described "identification" as adopting the constricted viewpoint of a playing piece within a game. But the player as conscious subject, with a life outside the game, is always free to take the game seriously or not. And just as the cognitive witness may remain unidentified with its perceptions, so the agent may remain unidentified with its actions. Indeed, from such a perspective of utter detachment, actions are not the player's but those of the game!

From a materialist point of view, information and causal forces in the physical world impinge upon the organism, driving its responses. But from an idealist point of view, awareness goes out from the subject, through the organizing lens of mind, which imbues the objects of consciousness with its own structure. Idealism should not take an ontological posture, to pronounce what exists (for example, mind); for, any concept of existence is subtly modeled on *physical* being. Rather, it should proclaim responsibility for experience and action. And materialism should not deny subjectivity, but foster a deeper understanding of the subject-object relationship.

The vulnerability of the organism is the root of both pain and pleasure, because the organism must evaluate stimuli in terms of their significance for its well-being. But this mandate—which is the reality principle itself—may ironically be a greater source of human suffering and delusion than the external world itself. Treating its experience as real *is* the organism's primary adaptation to the world upon which survival depends. But this capacity may extend reality, or "objectness," into areas that are non-physical or not literal. In such situations it is, in effect, a form of hallucination. The solution becomes a problem when reification runs amok, seeing reality where there is none.

Therefore, the ability to question one's experience, to treat it not as reality but as a creation of the nervous system or the self, can be a great boon. This ability to bracket experience as subjective is the hallmark of modern consciousness, the achievement of eons of biological and cultural evolution. It can be intentionally cultivated as a discipline. It is the source of the inner freedom promised in the great spiritual traditions. But, ironically, this answer to the mind's vulnerabilities is also the source of further vulnerabilities. Subjective sophistication can lead to paralyzing doubt, alienation, pathological detachment and inaction. Moreover, thoughts are far easier to manipulate than matter. The physical world resists human whim, but thoughts can be molded to desires. The way to control bodies is by controlling minds. The subjective mind, which lives in an inner world of images, can be prey to those who would manipulate its images. Deception and rationalization are possible, after all, only in a mental, subjectified world. The very cues the mind uses to assure itself of valid perception are just those that can be carefully and cleverly forged to produce the impression of authenticity.

The most effective way to control people is to define the games they will play and the goals they will pursue, and to control the cognitive channels through which they experience the world and themselves. Disinformation is communication given out by those in power to control the models held by society—either by distortion or by flooding the communication channel with irrelevant information. Facts are disarmed, or become outright lies, when isolated from their proper context, history, or integration with other facts in a comprehensive picture.

As a discipline or practice, a conscious relationship to experience is simply the willingness, each and every moment, to meet what is, however wonderful or unsavory. Spiritual freedom is not freedom from pain and suffering, and does not guarantee the ego's concepts of happiness; quite the contrary, it is rather freedom from the *compulsion* toward happiness or away from suffering. Ultimately, and paradoxically, it must be freedom from the compulsion to transcend.

Freedom is always contextual. In a "free" society, one is at liberty to drive along the street of one's choice, but not on the wrong side of the road. One can vote among several political candidates, but these options are preselected by processes from which most people are excluded. Even the elite, we must presume, have only relative freedom, which is the possibility of playing in a limited range of limiting games. The question is: who defines the game; and is it the largest, least limiting, and most engaging game one can play?

30. Objectivity

'Objective reality' is not the world-in-itself, but an intentional construct, an image and an ideal in the human cognitive domain. Evaluating the adequacy of this image, or of any cognitive model, cannot be a question of holding it up alongside the world-in-itself for an inspection of the fit. Our intuitive understanding of the adequacy of cognition is tacitly based on the assumption that the world simply and really *is* as we personally see it. This permits us, for example, to compare the cognitive domains of other creatures to our own standard. When we attempt to evaluate the *human* cognitive domain, however, we have no other standard against which to measure it.

A creature can get along quite well without what a human would regard as a global or objective image of the world, because its response to the environment happens to work—in the sense that the continuity of the genome is ensured even if that of the individual is not. Instinct is adequate in many situations, and Nature wastes little effort on programming that is more sophisticated than required.

A creature's representation of the world, which is thus locally adapted, may seem to us sparse and selective, like a digitized image that only indicates gross features. These might be, for example, large areas of light and dark, sudden movement, straight lines, etc. This is far from what appears to the human eye as an image of the scene. However, such an impoverished representation might be adequate to a given situation. By human standards, the creature's cognition would appear to be a simple rule-governed system based upon a few principles.

The downside of locally adapted systems is that there can be changes in the conditions on which the rules are based. The effectiveness of simple systems depends on the stability of their environment. This general arrangement works well for species with massive numbers of progeny (the species, if not the individual, survives) in a stable environment; whereas, to be globally adapted is more important for species of larger creatures of fewer numbers, which place heavier reliance on the individual's capabilities in a variable environment. For instance, mammals, which maintain an internal temperature in spite of ambient changes, are more globally adapted than cold-blooded creatures.

Relative adequacy means that the cognitive model takes more into account. The creature's knowledge is tuned to a wider range of contingencies. It employs greater sensitivity and flexibility in its dealings with the world in order to preserve constant its own identity. Its cognitive program relies more on learning and less on instinct. The simpler and more rigid the cognitive program, the more efficient it may be in the context to which it is tuned, but the more the organism depends on the environment. The more complex and adaptable the program, the more freedom the organism has from external contingencies, while internal consistency is a greater challenge.

The adequacy of a model is a function of how much the model encompasses, and the creature's range of actions depends on the contingencies it is equipped to deal with. Perfect adequacy could mean a model that mapped every possible contingency to effective responses. Apart from the neurological unfeasibility of such a brain, this creature would still not be *impartial*—and therefore not objective. Adequacy is a measure of how a mind responds to the complexity and subtlety of the world, but still in terms of given priorities. Impartiality is a measure of a mind's ability to engage multiple viewpoints, to transcend its own priorities. The former addresses possible states of external reality, the latter possible states of self-organization.

Linear or propositional thinking connects one point to another through rules of logic. Like a blind mountain climber, one moves through mental space step by step, unable to see the desitnation and sure only of the next foothold, in the faith that the top will eventually be reached if each step takes leads further "up." If many ascents are made by different paths, one finishes with a global knowledge of the mountain. Linear or propositional thinking may be a strategy pursued in the absence of a global picture, or it may be knowledge abstracted *from* such a picture. Either way, it is also the beginning of a *new* picture. And what is asserted at one level of cognition must be bracketed in order to create the next level. Individual pixels cannot retain their significance as isolated propositions if an image is to emerge from

their collective ensemble. The opposing propositions that are the horns of a dilemma, for example, cannot stand as independent assertions if the paradox is to be resolved. Thesis and antithesis must give way to synthesis. What is required for this to happen is a shift away from the given level, in which attention vacillates between two contradictory propositions or impulses, toward a meta level in which the contradiction disappears in an emerging larger picture.

The conflicts of life should force us to higher ground, revising our models to become more encompassing, adequate, and even impartial. The more holistic thinking becomes, the more convergence there should be of individual minds upon a commonly conceived reality—just as differences of visual perspective tend to diminish with distance from a scene. The maturing human consciousness recapitulates the child's development of an objective perception of space. It is an evolution of toward a more objective representation of the world. The key to this process, we have seen, is the development of the inner space of subjective consciousness, the very purpose of which is to create a subject-object relationship with the contents of experience.

Both objectivity and naïve pre-subjective awareness face toward external reality. But objective consciousness is a dialectical synthesis of the pre-subjective state with the qualifying skepticism of subjectivity. What distinguishes objectivity from the pre-subjective attitude is the willingness to take epistemic responsibility for the subject's role. There can be no true objectivity without responsibility. Objectivity and subjectivity are not opposites but partners in this dialectic, which appears to have functional value, through the increasing adequacy and impartiality of cognitive systems. Objectification is an evolutionary strategy, therefore, and not a fixed endpoint or state. As such, it is an article of faith and no guarantee of evolutionary success.

Objectivity does not confuse map with territory. While it may rely on particular perceptions, experiences, memories or states in order to access knowledge, it does not idolize these. Such experiences are *referents* for knowledge; the mind depends on them to know *that* it knows what it knows, but does not mistake them for the world that is supposed to be objectively known. The ideal of objectivity is never final, but always a relative state emerging in a cycle of learning; for, the adequacy of any system of thought is relative and temporary. Mind cannot finally attain to the absolute reality of the territory in the singular ultimate sense of truth it is so fond of imagining. Though it may be indefinitely refined, the map is never complete.

31. *Personal Identity*

The absolute nature of the epistemic subject is that it is a dimensionless point of view, interior to all experience. This can be pictured as a point in geometric space. We could invent a map of the panorama from this viewpoint, to classify phenomenal experience. On this map, immediately surrounding this point is located a zone of thoughts; surrounding that is a zone of emotions; surrounding that, a zone of somatic sensations. Beyond lies the zone of the external world, outside the skin. The entire "space" between the central point and the external world constitutes the "lens" of the mind-body, through which the world is experienced. But this lens may focus on anything in the space as well. In practice, some portion of the map is appropriated as "self." As the subject identifies with more and more of experience, this expanding shell is mistaken for the central dimensionless point.

To be a particular self is to have a special relationship to that part of the world that is one's own body/mind. This has two aspects: particular dedicated channels of information (the hardware of nerves), and particular response patterns (software). Both concern one's own body in a different way than other bodies in the universe. One's identity is based upon identification with the particular body/mind that is "one's own." While this relationship is not epistemically privileged, it is physically and affectively special. If I *can* know as much about the status of your body's tissue as I can about my own, it is through different channels, and with a different attitude. While I may not care as much about your tissues as about mine, in fact I can entertain these differing relationships to my own body as well: observing one's experience involves a different relationship than believing it. The *sense* of being myself, therefore, is

more a matter of attitude than a fact of being one body, or history of experience, rather than another.

Suppose a copy was made of the physical you, identical down to the molecular level. In what sense would this copy be "you"? It would be conscious, would think and react emotionally just as you do, and would—at the moment of creation, at least—be endowed with identical memories. It would be numerically different from you, in occupying a different physical space, but would otherwise be qualitatively identical. This numerical difference would give it *henceforth* a different experience. Even if it remained forever near to you (like a Siamese twin), it would occupy at least a slightly different perspective. With the possibility to move about separately in the world, its experience and store of memories could eventually diverge greatly from yours. Its identity would come to be that of someone very like you but clearly not you: like an identical twin.

Suppose instead that an identical copy is made of you, but at the same time (heaven forbid) your own body is destroyed.[18] With the copy substituting instantaneously for the body that was yours—and presumably complete with all your thought patterns, feelings and memories—it now seems that there would be temporal continuity between your consciousness and its. This makes it far more plausible to think that the copy would actually "be" you. A variation of this would be the basis of teleportation ("beaming up" à la *Star Trek*), where the copy appears at a different place and slightly different time, replacing the you that is here and now. There would presumably be an interval in which you were discarnate, a message traveling across space.

In the former case, we said the copy is *not* you and in the present case we say it *is* you—depending solely on whether the original survives! But how can the existence of another body affect your identification with this one? In either case—of teletransportation or of duplication—it is assumed (erroneously, I believe) that the process of copying is feasible and that only a finite amount of information is required to define "you." It is assumed, in other words, that "you" are virtually a program running on a digital computer. The identity concerned is actually that of an artifact, an intentional construct, and not that of an analog reality such as your body really is. If the reality of "you" happens to be infinitely complex and unfathomable, then the copy is merely a cartoon sketch of you. Simulation is not duplication.

But what are we to make of the subjective *sense* of being, experienced by myself as "I"? Since I can recognize memories as my own, just as I can recognize this body as my own, my consciousness might appear to me seamlessly continuous if the copying or teletransporting were perfect. But what if it was not? Might I then have the experience of waking up in someone else's body, if the new body seemed strange or the memories were unfamiliar or imperfectly matched? Even so I could perhaps be convinced by others that I was suffering from amnesia and that I must really be who they insist I am. Merely waking up in a different place and time, but as my familiar body and self, could be assimilated either as a lapse of memory or as evidence of teletransportation.

In the case where both of us survive the duplication, my copy would surely have his own sense of being, belonging to his numerically distinct body. From my point of view, at least, we shared the same qualitative identity at the moment of his inception. But ever since that moment he has been free to create an identity of his own through experience diverging from mine. We each continue with a distinct sense of selfhood and now a separate history as well. Moreover, *he* could argue—perhaps just as convincingly—that it is *I* who am the copy!

In the case where my body doesn't survive, my absence would not affect the copy's claim to selfhood. But it might affect his claim to my identity, at least in the view of others who might be inclined to doubt that he is me; or who might claim that there was always but one person, who was teletransported or bodily renewed through replication. Identity is a public fact to which he would be forced to reconcile.

It seems, then, that the subjective sense of "I" is a different matter than public identity: the particulars of what sort of person one is, with what history. These are matters of record, accessible to others. Memories contributing to one's personal sense of that identity may be regarded as personal experiences, but they are memories of events in the world. What differentiates me as an individual is a record of events. The experience of a numerically distinct version of myself could consist of (very nearly) those same events, creating very nearly the same identity. But the *sense* of selfhood is distinct

from such identity and does not depend exclusively on it. It must, therefore, be a sheer function of self-awareness, of the potential for self-reference.

We have looked abstractly at the game as a formalism that structures play. What of the competitive aspect, the play with other players? To view the game *as* a formalism is to see it with detachment, from the outside. To view the game from the inside, however, is to adopt a single point of view that identifies with the interests of a particular playing piece in competition with others. In the game of life, I believe myself to be *this* body, not that one, and certainly not all bodies. I agree to enter the game not only as a player but also as a playing *piece*. This is a process of "incarnation"—as though into a prosthetic body that mediates one's interaction with what is believed to be a world "outside." I then *become* this prosthetic body that fits like a glove or a space suit. To have a body is to take on the point of view, the purposes, values, and judgments of an organism: a discrete playing piece in the game of life. Without thus incarnation, one is free to look upon the game as a mere curiosity, a fascinating structure of which the body is an integral part. This is the consciousness reported by people who have near-death or out-of-body experiences, of serene detachment from the physical body. It is also the sort of experience reported both by sages and psychotics, who have lost the sense of self while still living in the body. One forfeits this detachment through "incarnation," which is a relationship of the subject to experience, not a metaphysical state nor an event in time. The notion of the transmigration of souls (reincarnation) is a misinterpretation of the timeless *logical* relationship of subject to object—another version of the MBP.

The consciousness associated with the body adopts the body's point of view and interests as a player in the game. The name given to this consciousness in various traditions is *ego*. The sense is different than in western psychology, where the ego is considered the part of the individual psyche whose job it is to juggle the internal demands of the organism with those of the socially constructed conscience and those of the environment. In other words, in western psychology the ego is part of the plumbing of a human being, who is already assumed to be a mind-body. But we are here using ego to mean the sense of particularity, separateness, and self-centeredness which accompany identification with the playing piece and its priorities within the game. It is a normative as well as descriptive concept. Ego is identification with the "lens" of the mind-body and all its programming and concerns, and with the view of the material world as seen through it. Ego, in other words, believes in the reality of the game. A disembodied observer would have no reason to take the game seriously. Playing in earnest, however, one is set against other players, with a narrowed focus on the goals of the game and on winning.

The mind that identifies with the game can only look outward to a reality it experiences as external and separate from itself. Since all that exists for it is what is defined in the game, its attention is glued to the world of the game, which it is committed to experience as real. It only judges real what it perceives as not itself, and can only see what is inside itself by projecting that into the world. It has the nasty habit of disowning responsibility through projection. Believing that the external world holds all the cards, it sees itself not as actor but as reactor, not as free agent but as victim. A victim can only blame and be blamed by other victims, believing that the good and the bad lie with other players. Ego presumes to judge, therefore, but objectifies both good and evil, projecting it into the world. It can tolerate neither the thought of *being* evil nor the *responsibility* to be good. Rather, it disowns its judgments as realities *out there,* over which it has no control.

If the ego is but an empty shell in a shell game that is itself a mere form, where is reality? We have seen that the *sense* of reality comes from embracing the terms of the game. Ironically, it is this very sense or reality that can be illusory: what the Easterners call maya or samsara. From an absolute point of view, the game is an arbitrary and meaningless recreation, an exercise of inventive imagination. The importance we attach to it is ultimately gratuitous, intentional, not logically compelled. In the gaming metaphor, a player "incarnates" by embracing the world and body of a playing piece, provisionally fogetting the fact of having done so. We are in the dream trying to remember the waking state, in the game trying to recall a life outside the game.

Ego, however, does not particularly want to awaken from this lapsus, to be unseated from its kingdom, and correctly views inquiry into identity as a threat to its entrenched power. Its resistance is

enormously resourceful. In fact, ego will do anything to keep its position, which, like that of the Wizard of Oz, is tenuously maintained with smoke and mirrors. Specifically, it is committed to the illusion of a separate and private life with its personal entitlements, a body implicitly opposed to other bodies, and control of its surroundings (and hence of others) in order to maintain its position. Like any tyrant, it justifies a heavy hand. Stories of actual tyrants over the ages reflect the qualities of ego: self-centered, self-serving, self-indulgent, greedy, ruthless, brutal, heartless, cunning, malicious, underhanded, manipulating, power-hungry, arrogant—in short, the very profile of evil. The evil in the world is ego unchecked by mitigating forces such as Nature, other egos, social convention, morality, law, conscience, love, understanding, spirituality. The Devil is an apt metaphor to personify egocentricity.

32. Social Responsibility

The drama of history focuses on great names, but power requires an enormous supporting cast, who double as audience for this drama, and who allow their lives to be directed by others.

Perhaps the world has always been dominated by a small elite. Evil in modern times, however, has been democratized; it has systemic through the willingness of society to be co-opted into the schemes and institutions of elites. Modern evil is automated, impersonal, out of sight. It is at once the invisible machinations of the corporate world and the faceless *others* who are emerging as the new slave labor for the West. The globalist economic system, which consumers accept and help to maintain, should be called what it is: a return to pillaging, in lieu of genuine productivity and the global advance of civil society that is promised in the myth of "progress."

By selling out to the company store, one forfeits the autonomy that is the nominal benefit of democracy. By allowing others to act immorally in our name, we become the secret accomplices of power and the co-creators of an immoral society, which pretends to have found the key to abundance and humanitarian values through righteousness, industriousness, and technology. In truth, it maintains a privileged position in the world through insidious new forms of domination, which it must deny to keep faith with itself. Evil deeds become the normal workings of society, compounded with hypocrisy, self-deceit, propaganda, and outright lies, in a society that must carefully manage its self-image.

The middle class—traditional bastion of human rights, equality, and democratic institutions—shrinks in retreat, because the present economic system is a contest for dwindling resources that can only increase the disparity between winners and losers, eliminating those in the middle. Manufacturing is no longer a mainstay of the overdeveloped world, which has shifted to an investment economy dependent on overseas factories, with virtual slave labor. The social struggles of the Industrial Revolution in Europe and America have simply been pushed offshore, out of sight and mind. Management and labor live no longer in separate parts of town, across the tracks, but in entirely different societies, isolated by thousands of miles, national borders, and economic treaties. Far from being a collective effort to produce the commonwealth, an investment economy is less an economy at all than a system for distributing plunder. But in spite of harnessing the rest of the world to produce its cheap goods, the consumer society bankrupts itself with debt and loses the class equilibrium that originally made consumerism possible. The debt and unemployment enslaving the Third World has become a domestic crisis. The greedy corporations we invest in to bring us the good life gobble up our own governments, public property, and private pocketbooks, along with those of foreign nations and the planet's natural riches.

Rule by corporation is a new form of government that has slid into place almost unnoticed in the last century and silently taken over the world. It is a mutant cleverly adapting to overcome society's resistance to manipulation and exploitation. It mimics democracy (after all, shareholders can vote) and mouths the rhetoric of freedom (of the flow of capital), while destroying democratic institutions. It impersonates responsibility and good will while impoverishing millions in the obsession for profit. It is the modern monarchy and nobility, in a new feudal order, but lacking *noblesse* and social obligation. Like so many portrayals of machines run amok in science fiction, corporate behavior is crude, automatic, heartless, and totally unconcerned for the human future.

As a legal entity, with the rights of a person, the corporation is virtually an immortal artificial life form, designed to shield and magnify the power of the real people behind it. It is a virtual man-made pathogen that has escaped legal containment, just as genetically modified organisms and artificial life forms risk to escape biological containments. Like a cancer or viral plague in the body politick, its sole purpose is growth without restraint. No longer under public control, it now controls the public. Like some parasites in the natural world, it takes over its host's behavior and thoughts. No matter if the victim dies, so long as another can be found.

Global capitalism is a virtual machine, an artificial parasite, a world-wide pandemic that feeds on slave labor and sucks wealth as a nonrenewable resource wherever it can. While the rhetoric of globalism is that all countries and all segments of society will ultimately benefit from free markets, in plain truth the rich get richer and everyone else poorer. Wealth does *not* trickle down in this economic pump, but always *up*. And this is not only a quantitative impoverishment; as the world becomes a monoculture, its diversity is reduced to a few standard monopolies, franchises, and products in look-alike cities around the world. This is the actual *intention,* moreover, behind corporate globalism, and the ultimate economic fulfillment of an ancient dream of power through mechanism: to re-make the whole world as a virtual engine of profit. The monoculture itself has become a factory for mass-producing consumer clones and values, whose function is to be an economic host to a handful of ultra-powerful parasites. Just as pre-industrial societies milked and bled their animal stocks for sustenance and slave labor; and just as industrial magnates herded peasants into a poor urban labor class; so corporate capitalism today manages its consumer-investors, on the one hand, and its dehumanized offshore labor, on the other. The metaphorical lesson of *The Matrix* is but thinly disguised. What consolation can be taken in the fact that our overlords are only human and mortal like ourselves?

The paradox of corporate globalism is that it is at once a world machine gone berserk and the very epitome of control: the ultimate refinement of power, representing the promise of human mastery and unity. The contradiction is only apparent, for both are illusions. Corporatism appears out of society's control because society refuses to exert the control it does still possess. *No company could exist but for those who buy its products or services or invest in its stocks*. A company grows large and powerful only because so many people patronize it, both as consumers and investors. The counterparts of the corporate bottom line are the consumer's "right" to a bargain and the shareholder's "right" to an automated income that is not a result of productive labor. These are the twin pillars of our society, the dual seductions that lead us into evil, in spite of our religious pretensions.

We pursue of the "good life," as though it were a right and a natural inheritance, rather than a spoil of conquest and a lethal imposition on nature. This is the fundamental delusion through which consumer/investor society dissociates from the consequences of "progress." Malls and walled communities provide our false security, the media our bubble of delusion. Literally on drugs—legal and otherwise—we prefer Hollywood to reality.

External censorship and top-heavy propaganda, in the style of Soviet Russia, are laughably passé. Here in the West, propaganda and censorship are immanent in the system, not externally imposed by dictatorial governments. This is not to say that there are not powerful elites who pull the strings. Only that we are happy enough to be puppets. The issues in the twenty-first century are far more sophisticated than old-style totalitarianism: *self*-censorship, commercial propaganda, fear mongering, corporate media domination, internet control, political-military spin-doctoring and, generally, apathy and willful ignorance. In concert, these promulgate the "necessary illusions" that effectively constitute a paradoxical new form of repression. We agree not to speak, listen to, print, read, nor *think* the simple obvious evils of our moral complicity in the "New World Order," because we do not care to forfeit the lifestyle founded upon it. Any fool has the freedom to get on a soapbox in the seedy part of town; no one will hear him from the suburban shopping mall. The much-vaunted freedom of Western society offers little more than trivial shopping and entertainment choices squeezed from the same tube. Globalism *is* the elimination of genuine choice and autonomy, whether in the political arena, on the big screen, in universities, or in the marketplace. That is its intent. We pay for consumer convenience with a homogenized world of glitz and falsehood. But convenience can never relieve us of the responsibility to create a more genuine and moral

world.

APPENDIX to Section 18 (*Coloring it Real*)

The retinal blind spot is a small region in the visual field corresponding to the area on the retina where the receptor nerves are bundled to form the optic fiber, thus excluding the presence of receptors themselves. This hole in the visual field is simply *ignored* by the mind. Indeed, awareness of it would be redundant and a nuisance, since the blind spot is standard equipment in the healthy eye. However, destructive lesions in the retina or optic nerve give rise to an actual experience of a blank area in the visual field. It makes sense that the abnormal condition of lesion at this primary level must be noticeable, since it could represent an absence of light in the real world. It is the job of the retina to pass on this information to the cortex, which is able to adapt with time to such lower-level injuries.[19] Lesions in the visual cortex itself can cause similar defects in sensitivity, without ever being perceived. The patient may deduce the presence of a defect, but is unable to see it directly, since (so to speak) there is no one higher up to countermand the orders of the cortex.

Neurophysiologist Karl Lashley noted that the transitory scotoma he experienced due to migraines were filled in under certain conditions. Specifically, when he gazed at checkerboard patterns, or other regular patterns of wallpaper, these appeared subjectively completed over the missing parts of his visual field.[20]

A person dressed in a black costume, so as to be invisible in the dark room, except for small white dots strategically placed at the body's joints, is seen only as a pattern of dots until the person begins to move. Then the invisible outline of the human form appears subjectively as a moving shadowy figure.[21] The tribulations of proofreaders also demonstrate a form of pattern completion, since there is a tendency for errors to be seen as correct in their context. And short deletions in the spoken word may also go unnoticed, the normal flow of the sentence being subjectively assumed.

Phenomena of apparent change are further examples. One of these is the characteristic of human vision responsible for the ability to see a motion picture as a continuous changing image rather than a series of distinct stills. When a spot of light is briefly flashed against a darker background, followed within a specific interval by a similar spot flashed a short distance away, the subjective effect is of one spot moving from the first position to the second. If the time interval is shorter, two simultaneous flashes are perceived; if longer, two successive flashes. This is the deceptively familiar principle of the “moving” lights of the illuminated marquee. Such effects raise an interesting question: how does the mind know where the next spot will occur, in order to fill in between the two? One investigator has speculated that the intervening apparent motion is projected backward in time.[22] Other kinds of apparent changes include transformations of size and shape. For instance, if instead of a dot a square is flashed, followed by a triangle and a circle, what will be seen is one figure changing into another in a smooth transition from square to circle.[23]

The projection of experience backward in time is implied by other experimental evidence as well. The paradox of habituation is demonstrated, for instance, by the familiar experience of becoming accustomed to a sound, so that it is only noticed when it stops. Neurologically, an orienting response in the pattern of brain waves begins when the habituated signal ends. [24] One's attention is suddenly drawn to the “deafening silence”. Similarly, the phenomenon of backward masking indicates that awareness of a brief, near-threshold visual stimulation can be obliterated by another, longer stimulus occurring shortly afterward.[25] Other experiments compare the timing of experiences due to stimulation of the skin, with that of similar experiences due to direct stimulation of the part of the cortex receiving impulses from the skin. These demonstrate that the skin stimulus is subjectively referred backward in time about half a second-presumably to compensate the time it takes for the nerve impulse to reach the cortex from the sensory surface. Strictly speaking, it seems that sensations are always replays of past events, for which we nevertheless have the subjective sense of real-time participation.[26]

Sensations of vision and hearing are projected outward in space as characteristics of objects in an external world. This capacity of spacial projection can be acquired by the skin surfaces as well. Such projection is an everyday tactile experience. Anyone who has had to loosen a screw in a visually inaccessible place is familiar with the extension of one's sensitivity of touch to the tip of the screwdriver. No doubt blind people experience the tapping of their canes as *out there* as well as being sensations in the hand.

Daniel Dennett has criticized the interpretation of a range of completion effects that are by now quite familiar.[27] I believe his point is well-taken, that the apparent experience of completion in the visual field is in truth a (mis)judgment-something like seeing a snake when there is actually a rope-rather than a veridical perception of some sensory quality imputed with some kind of existence (a sense datum). It is the projective capacity of the mind (the ability to even make such mistakes) that I wish to underline with these examples.

The projective adaptability of the brain has been dramatically studied by a type of experiment first performed around the turn of the century and subsequently repeated in several variations. First, let us note that in normal vision the optical image on the retina is upside down, since it is inverted by the lens of the eye. In other words, the brain *normally* adapts itself to this inversion, so that the world is subjectively experienced right side up. What would happen if the optical image itself were re-inverted? In the experiment performed, a subject wore a special lens over each eye which inverted the retinal image-and therefore re-inverted the subjectively experienced visual field. Needless to say, this resulted in a very disorienting experience. Over several days, however, the subject was able to adapt to this condition-if not perfectly and consistently-so as to experience the world once more as normal and right side up! Though the results of this drastic experiment were equivocal, less extreme distortions of the visual field have proven conclusively the adaptability of the human visual cortex. The effects of lenses which systematically curve the visual field, of prisms which displace it by a few degrees to one side or the other, and of special glasses which color half the visual field red and half green, have all been compensated after a few days or even hours of getting used to them, so that normal experience is restored. In one instance, where prismatic spectacles were worn to displace the visual input by a few degrees, subjects would at first consistently mis-reach for objects about them. But after only one hour of *actively walking about*, the subjects completely and exactly compensated the displacement. When the subject was *passively* moved through the same environment in a wheel chair, little compensation took place.[28] This suggests that we do not see the world as passively presented by the senses, but rather as we come to know it through active interaction. We are reminded of the homunculus who, solely through interaction with changes on his display panels, pieces together a functional knowledge of an outside world-a knowledge that is experienced as *seeing*. Just as the spectacled subjects in these experiments learn to see in their retinal displays the world implied through their active explorations, the little man comes to see, in the patterns of his instrument readings, a real external realm.

Not every creature has such adaptive flexibility. If the eyes of the goldfish are surgically inverted, it never manages to compensate for this disruption, always swimming away from the true location of its food. Experiments similar to those described with human subjects were performed with monkeys suffering from various cerebral lesions. Frontal lobotomy effectively prevented the hapless creatures from compensating a displaced visual input, whereas other cortical removals did not. Apparently, the adaptivity concerned in such experiments is a high-level function of the most complex nervous systems.

---

[1] Leibniz *The Monodology*, sec. 17, from *Philosophical Classics* ed. by W. Kaufmann, Prentice-Hall 1961, p227-8

[2] Experiments of von Bekesy, cited in Karl Pribram *Languages of the Brain* Prentice-Hall 1971 p167-71

[3] Plato *The Republic*,VII 514, transl. by Francis M. Cornford, Oxford UP 1945 p227-8

[4] A.J. Ayer *The Problem of Knowledge*, Penguin 1956, p37 (quoting an analogy of Ryle's)

[5] Plato op cit p229

[6] Gordon R. Taylor *The Natural History of the Mind* Secker & Warburg 1979 p193-5

[7] Robert Rosen *Life Itself* Columbia UP 1991, p47-8

[8] Sigmund Freud "Formulations Regarding the Two Principles of Mental Functioning" (1911), sec. 4, in *A General Selection of the Works of Sigmund Freud* ed. by John Rickman, Liveright 1957, p43

[9] In the same vein, 'This sentence is not green' *could* be "formally" expanded as "'This sentence is not green' is green"-- in contrast to its "unrelativized" expansion as "'This sentence is not green' is true". The statement is subjectified or relativized when the predicate ("greening") replaces the truth function, so that the sentence does not refer beyond its own terms (i.e. to a larger world in which it could be true or false).

[10] David Hume *An Enquiry into Human Understanding*, sec. 7 part 1, in Kaufmann op cit p352-3

[11] Humberto Maturana and Francisco Varela *Autopoeisis and Cognition* D. Reidel 1980.

[12] Hiram Caton *The Origins of Subjectivity: an Essay on Descartes* Yale UP 1973, p87-8

[13] Maturana and Varela op cit

[14] Bernard J. Baars "Conscious Contents Provide the Nervous System with Coherent, Global Information" in *Consciousness and Self-Regulation* vol 3, ed. by R.J. Davidson et. al., Plenum Press 1983

[15] Vernor Vinge "Technological Singularity" Whole Earth Review #81, Winter 1993, p88ff. The term "singularity" was used earlier by Von Neuman in regard to runaway technology.

[16] Nancy Cartwright

[17] Eric Drexler *Engines of Creation* Anchor/Doubleday 1986

[18] D. Parfitt *Reasons and Persons*, 1984

[19] cf. G.R. Taylor op cit p209: the experience of the pioneer psychologist Kenneth Craik who (deliberately) burned a hole in his retina through exposure to the sun. He reported two days later that the missing part of the visual field was subjectively filled in.

[20] Annecdote reported by H.L. Teuber in *Brain and Conscious Experience* ed. by John Eccles 1966, p193

[21] Keith Oatley *Perceptions and Representations* Methuen 1978, p3-4 and 190-1

[22] Nelson Goodman *Ways of Worldmaking* Hackett 1978, p72ff (discussing Paul Koler's experiments in *Aspects of Motion Perception* Pergamon 1972)

[23] ibid

[24] Karl Pribram op cit p256

[25] Eccles (ed.) op cit p175

[26] Eric Harth *Windows on the Mind* p199-202

[27] Daniel Dennett *Consciousness Explained* 1991

[28] Early experiments of Stratton, embellished by others in several variations.