Chapter Six

Learning to Think Differently About Knowledge and Mind

> [F]olk psychology has already been pushed way past its limit.
>
> —E. O. Wilson (1998, p. 202)

The purpose of this chapter is to pull together ideas from the preceding five, in preparation for shifting the focus of discussion in the remainder of the book to education. What I have set out so far is a way of thinking about the mind that is different from the folk way and that I have argued has important advantages for dealing with knowledge. It is not what I would call a theory. Possibly it constitutes some elements of a theory, but that depends on what we intend a theory of mind to be. In any event, the point of my argument has not been to establish matters of truth but rather to argue pragmatically.

In order to get you to adopt a new way of thinking about knowledge and mind (assuming you are not already there ahead of me), I do not have to convince you of theoretical propositions. In fact, I must try to get you off the theoretical track, because it is a track that can lead to endless quibbling about definitions, demands of "how do you explain such-and-such," and counter-arguments of greater or lesser theoretical weight. Instead, I have to convince you that there is a pay-off in making the conceptual shift—that it will help you in your work, for instance—and then I have to help you make the shift. For the shift is not simply a matter of understanding and agreeing with an argument, it is a matter of learning a new way of interpreting and ultimately of perceiving events. In this chapter I take a stab at both, first trying to convince and then doing a walk-through of several cases that show how the alternative way of thinking can work.

How Far Can Folk Theory Take Us?

Besides providing the commonsense psychology and epistemology to get us through our daily rounds, folk theory also forms the basis of more advanced theories. Until the rise of

connectionism, most of cognitive psychology and artificial intelligence was based on it. And as E. O. Wilson (1998) notes, the social sciences, including economics, generally assume it without question. These more advanced disciplines share the same basic limitations as the naive psychology and epistemology we picked up as children, but this is not always obvious, because of the overlay of sophisticated concepts. Yet the underlying folk theory hampers the social sciences, both in their scientific progress—which is Wilson's concern but not the concern I am pressing here—and in their practical value.

To illustrate the latter point, I will examine Nonaka and Takeuchi's model of knowledge creation (1995). Why pick out this theory among the many epistemological, cognitive, and sociocultural theories that have a more established scholarly reputation? Firstly, because unlike almost all the others its approach is pragmatic. It is intended to help businesses become more successful at making use of the knowledge resources they have and at generating new products and strategies. And unlike most of the business literature, which consists of aphorisms larded with examples, it tries to explain and not merely to instruct. There is another strain of pragmatic theorizing about knowledge that is gaining attention in both the business world and in education. This is a strain that runs from John Dewey on in to contemporary work on situated cognition (Brown, Collins, & Duguid, 1989). However, as we saw in Chapter 3, situated cognition represents a rejection of folk epistemology and folk psychology; Nonaka and Takeuchi's model, in contrast, remains solidly rooted in it.[1] Accordingly, their model provides a highly relevant test case for seeing how far folk theory can go in dealing with knowledge creation. It goes pretty far. Among other things, it incorporates ideas of situated cognition to the extent that this is possible within a folk theoretic framework. Finally, it merits critical scrutiny because it has received wide acceptance within one substantial community concerned with knowledge practices;

---

[1] This is curious, inasmuch as Nonaka and Takeuchi maintain that the Japanese have profoundly different conceptions of knowledge and mind from those that characterize the West—more organic, free of the Cartesian dualism, and in fact free of epistemology in the classical sense. But when it comes to theorizing, Nonaka and Takeuchi are right in there with their western counterparts, positing individual minds that are full of unformulated knowledge and an external world that they must somehow project that knowledge into.

in the field of knowledge management it is practically the only theory going.

Nonaka and Takeuchi's model of knowledge creation recognizes two kinds of knowledge, tacit and explicit. Although its components are not spelled out, tacit knowledge would appear to include five of the six kinds of personal knowledge discussed in Chapter 5, the exception being statable knowledge. Explicit knowledge comprises statable knowledge and conceptual artifacts, thus, as is characteristic of folk theory of mind, making no distinction between them. The process of knowledge creation consists of four subprocesses that transform knowledge. These are summarized in the following table:

Input	Process	Output
Tacit knowledge	Socialization	Shared tacit knowledge
Tacit knowledge	Externalization	Explicit knowledge
Explicit knowledge	Combination	Synthesized explicit knowledge
Explicit knowledge	Internalization	Tacit knowledge

Nonaka and Takeuchi are clearly trying to coordinate the two folk conceptions of knowledge—knowledge as things in the mind and knowledge as something out in the world. In keeping with the traditional individualistic view, however, they treat knowledge in the individual mind as primary. Public knowledge is private knowledge brought out into the open—rather like what happens on talk shows. However, Nonaka and Takeuchi recognize that this externalized knowledge can take on properties of its own, much as Popper asserted. So what is missing? We have World 2, represented by tacit knowledge, and World 3, represented by public knowledge. Presumably there is also World 1, the physical world.

Inasmuch as theirs is a pragmatic theory, the way to identify what is missing is to see what their theory cannot do or cannot do well. Nonaka and Takeuchi's model falls short on four counts:

1. *Creativity.* Although it holds that new knowledge is always created in individual minds, it does not explain how minds produce original ideas and novel solutions. The book is full of examples of people doing that and suggestions for stimulating mental activity, but the model does not suggest what could be done about a group that is highly productive of unoriginal ideas and ineffective solutions. This is a common failing of the knowledge management genre, but it is worth noting here because of the authors' claim to be uncovering the secrets of knowledge *creation*.
2. *Understanding.* Although the model deals with ways that knowledge gets from person to person, it offers nothing about understanding and depth of understanding. That immediately disqualifies it as a model for education, but it is a serious weakness in a model for knowledge management in business as well. Depth of understanding is a distinguishing characteristic of expertise in knowledge-based fields, and productive creativity presupposes expertise (Bereiter & Scardamalia, 1993). Accordingly, a model that offers nothing about deepening of understanding is severely limited in what it can offer about the creation of knowledge.[2]
3. *Knowledge work.* Although Nonaka and Takeuchi appreciate the importance of knowledge abstracted from practice, their model has little to say about the production, management, improvement, or application of such knowledge—in short, about knowledge work. As interpreted by their model, what all those scholars and scientists who populate the great research universities are

---

[2] One of the most interesting stories in Nonaka and Takeuchi (1995) would seem to argue against the importance of understanding. It is about the design of a bread-making machine. An engineer apprenticed herself to an expert baker, observed that he twisted the dough in kneading it, and designed this twisting action into the machine with profitable results. Presumably neither she nor the baker *understood* why twisting produced better bread. She simply observed the practice and translated it into another medium. I am inclined to call this an example of dumb luck. In the absence of understanding, there was no reason to believe that twisting the dough was an important part of baking. It just happened that it was. Furthermore, if one understood how twisting made a difference, it might be possible to achieve the effect in ways much more efficient than building a machine to mimic the baker's motion. Centuries of observing birds and trying to mimic them did not produce a flying machine. Success came with growth in understanding of lift.

busy doing is ‘combining’ pieces of each other’s knowledge. Yes, you can interpret it this way, but is it useful to do so? You can also interpret making a shirt as combining pieces of cloth, and although this might provide an interesting slant it does not seem like an interpretation that will carry you very far in trying to improve manufacture. Don’t we need a model that treats the production of cognitive artifacts in at least as sophisticated a way as we treat the production of garments?

4. *Collaborative knowledge building.* Although cooperation and teamwork are praised, the idea of cooperating in the creation of knowledge never comes to life in Nonaka and Takeuchi’s theorizing. There is a lot about knowledge moving from one mind to another, but the rather straightforward business of a group of people working together to produce a design or plan or to solve a problem seems to become mysterious in their conceptual framework. Because public knowledge in their model is merely an outward extension of personal knowledge, it is hard to deal with the everyday fact of people jointly producing a piece of knowledge that is neither the product of one individual’s knowledge nor a combination of several individuals’ knowledge. Such knowledge is typically an emergent of discourse and cannot be understood at the level of individual interacting minds. To promote it, you have to know how to promote progressive discourse.

The problem in Nonaka and Takeuchi’s model is not so much missing concepts as missing perspectives on those concepts. This comes from flattening everything into the two-dimensional space of folk theory. In this two-dimensional space, where public knowledge is externalized personal knowledge, several very important relational properties are lost. One is the relation of persons to conceptual artifacts: the person’s familiarity with, understanding of, and attitude toward a particular theory, technology, problem formulation, or interpretation. Another is the relations of conceptual artifacts to one another and to the the tasks or problems to which they may be applied. One theory or design may be destined to win out over others regardless of the vagaries of personal opinion, because it can do things that the others cannot (Dennett, 1995; Thagard, 1989). Together, these

represent the major distinction that both folk and sociocognitive theories have difficulty making: the distinction between the knowledge involved in productive work and knowledge that is the product of that work. Without this distinction, you cannot formulate the essential *educational* challenge that a knowledge-creating organization needs to face: Promoting learning that will increase the organization's ability to create knowledge.

An Alternative Approach

The failings of folk theory, I have argued, can be traced to its way of treating knowledge. Under the influence of the mind-as-container metaphor, knowledge is treated as consisting of objects contained in individual minds, something like the contents of mental filing cabinets. Although that way of treating knowledge is convenient for many purposes, it severely limits and distorts our dealing with issues of understanding, mastery, and knowledge advancement. To overcome its weaknesses, I have proposed that folk theory of mind needs to undergo two radical changes:

1. The idea of mental content should be reduced to the status of a metaphor, useful for some purposes and not for others, but not to be taken literally. Connectionism provides an alternative metaphor, which enables us to conceive of a mind that can act knowledgeably without containing propositions or other knowledge objects. In order to gain benefit from the connectionist metaphor, we must find ways to construct mentalistic accounts that do not refer to things residing, being searched for, or undergoing changes in the mind.
2. Abstract knowledge objects, such as theories, numbers, and designs, should be accepted as real things outside the mind—as conceptual artifacts—with which people may develop relationships, much as they do with animate and inanimate material things. Understanding and mastery may then be treated as characteristics of such relationships, and the advancement of knowledge as the creating and improvement of conceptual artifacts.

Let us see how such an altered approach to knowledge and mind can deal with the four points on which I criticized Nonaka and Takeuchi:

1. *An alternative view of creativity*. Folk theory mystifies creativity. This is because ordinary rational thought is believed to proceed in a step-by-step manner, with each step following in some logical (or illogical) way from its predecessor. Creative thinking obviously doesn't work that way. But, according to the connectionist view, ordinary thought doesn't work that way either. Ordinary thought is, in fact, relentlessly creative. Ordinary speech is not anything like an orderly logical process. It is full of invention. Logic comes in after the fact as a way of *justifying* what we have done or said. Ordinary acts of thought can be justified by logic that is simple enough that we can easily imagine it having been run off in the mind; creative acts cannot. But we can also logically justify the actions of fish and insects, without having to suppose that a reasoning process is actually carried out by the creatures in question. What is remarkable about human cognition, from the connectionist point of view, is not creativity but the ability—achieved with considerable effort and practice and quite fallible for all that—to sometimes control our thinking in such a way that it does proceed in a step-by-step logical manner.

It is not creativity per se that needs explaining and promoting, then; creativity may be taken as a given. It is big-bang type creativity that produces major leaps ahead and also sustained creativity, which can be relied upon to produce advances year after year and to carry them through until they amount to something. Folk psychology directs all our attention to the first and ignores the second, which in most walks of life is more important. Present a problem to any group of children or to a group of mentally active adults and you will get plenty of ideas. Ideas are a dime a dozen. The children's ideas will mostly be useless because they fail to meet the constraints of the problem. The adults' ideas will mostly be useless because they honor the wrong constraints—constraints arising from convention, habit, and surface appearance rather than constraints of a deeper kind. Useful creativity depends on a deep understanding of the constraints and a fund of relevant impressionistic knowledge (what folk psychology mystifies with the word 'intuition') that serves as a basis for recognizing promising ideas. Carrying a promising idea to fruition depends even more on knowledge of many kinds. Accordingly, if an organization wants to increase

the creative output of its employees, the starting point ought to be enriching and deepening their knowledge. But this is not a simple matter. The knowledge that makes creative advances possible is knowledge that arises from trying to make such advances.

Thinking of creativity in this way does not lead forthwith to a recipe for promoting it. What it does is naturalize the concept so that it takes its place among other concepts relevant to boosting an individual's or an organization's intellectual capital. In particular, it ties creative capacity to the growth of understanding.

2. *A relational conception of understanding.* Everyone is in favor of understanding, but folk theory diminishes and isolates it, so that in discussions of innovation and intellectual capital (e.g., Nonaka and Takeuchi, 1995; Stewart, 1997) it is barely mentioned. Understanding is reduced to holding true beliefs and being able to explain, which are anemic virtues compared to having the red-blooded know-how that gets things done and produces innovations. There is a kind of understanding that is more fully appreciated, however. It is the understanding salespeople have of their clients, the understanding craftspeople have of their tools and the materials they work with (cf. Harper, 1987), the understanding athletes have of their own bodies and what they can do. Understanding of these kinds is not usefully thought of as consisting of things in the mind. It is a sort of intimate familiarity arising from and closely tied to intelligent action. In Chapter 4, I tried to show that theoretical understanding is best thought of in the same way. Understanding is the totality of one's knowledge of an object (material or abstract), considered from the standpoint of the ability of that knowledge to support intelligent action. Thus the development of understanding has to be at the heart of any educational effort or effort to enhance the innovative capacity of a person or organization.

3. *Knowledge work as work that creates or adds value to conceptual artifacts.* As I noted in Chapter 3, the term 'knowledge work' as it is currently used is little more than an honorific. Knowledge work is simply whitecollar work that for whatever reason commands a relatively high salary. Until you can conceive of knowledge as real stuff that it is possible to do

work on, you cannot pin down what distinguishes knowledge work from lowly paper-pushing on one hand and from knowledge-demanding manual occupations like brain surgery and safe-cracking on the other. Unless they start licensing knowledge workers, a precise definition is not likely to become a priority. However, if we want to consider educating people to become knowledge workers, or retraining manual workers as knowledge workers, or making the work environment more conducive to knowledge work, then it would be helpful to have a clearer idea of what we are talking about. The second of the radical conceptual changes indicated at the beginning of this section leads to a very natural definition of knowledge work. It belongs to the same class as metal work, woodworking, leather work, and personnel work except that the objects worked with are abstract: they are conceptual artifacts. Knowledge workers create, improve, find new uses for, or otherwise add value to conceptual artifacts.

We may now ask what skills and other kinds of knowledge are needed for knowledge work and what kinds of conditions facilitate it. Our very abstract definition of knowledge work ought to suggest that the answers to these questions will depend a great deal on what particular kinds of conceptual artifacts we are talking about. If your job is calculating the heating and air conditioning needs of office towers and mine is allocating the resources of a large social service charity, it is not immediately evident that our skills have anything in common worth mentioning. Or, to put it differently, our knowledge would overlap more with that of *non*knowledge workers in our respective fields—yours with the knowledge of the people who install air ducts, mine with that of the social workers, health care providers and others who deliver the services I allocate funds for—than it does with each other's.

Are there any special competencies that distinguish knowledge work of all kinds? What about thinking skills and skills in locating and organizing information? To anticipate the conclusions of Chapter 10, there is much less to general cognitive skills than folk theory leads one to suppose. Furthermore, although you and I might have certain planning and information-handling skills that go by the same name, we developed those skills through work that got us deeply into the

problems, constraints, and know-how of our respective fields. If we exchanged jobs, the amount of intellectual competence that would transfer from one job to the other would be quite limited and would owe more to our genes and to how we have lived our lives than to any identifiable learning experiences. This does not mean that nothing can be done to prepare people for knowledge work, but it means that we must be wary of facile prescriptions. Like any kind of job competence it is mostly acquired on the job; what is done in advance—through schooling and other measures—must be aimed at increasing the likelihood that that learning, which is mostly acquiring job-specific understandings, will actually take place.

4. *Knowledge as a social product* . "In a strict sense," Nonaka and Takeuchi say (p. 59), "knowledge is created only by individuals. An organization cannot create knowledge without individuals." The second statement, which is obviously true, is presented as if it constituted warrant for the first statement, which is a questionable one. The argument does not hold together. A company cannot produce automobiles without individuals, either, but it does not follow that automobiles can only be created by individuals. In fact, no individual could build a modern automobile. It is doubtful if any individual knows enough, has all the necessary skills, and would have enough hours in the day just to produce the *design* of a modern automobile, let alone produce the physical product. Nonaka and Takeuchi manifestly know this, so why do they treat knowledge so differently, as something that can only come about within the individual mind? Well, because that is where folk theory locates knowledge. Until there is a way for your co-workers to get inside your brain and fiddle with the synapses, there is never going to be such a thing as the collaborative creation of knowledge, according to folk understanding. But if you can conceive of knowledge as consisting of conceptual artifacts, then you can imagine something like a knowledge assembly line, with theories, designs, and so on moving along it, being worked on by various people according to their various skills, and coming off the end as finished knowledge products. Does that sound ridiculous? It is not far off from the way news stories are put together for news magazines. It is not the way masterpieces of journalism are produced, but that is not the point. The point is

that something can be produced through collaborative knowledge work. Most of it is not and need not be original or brilliant.

Having established that knowledge can be produced as a social product, however, and that many complex cognitive artifacts can only be produced through collaborative effort, we can then turn our attention to problems of quality. That is what Nonaka and Takeuchi aim to do, but they are hobbled from the beginning by their belief that everything originates in the individual mind. The mark of a really successful design or problem-solving meeting is that something brilliant comes out of it that cannot be attributed to an individual or to a combination of individual contributions. It is an emergent, which means that if you look at a transcript of the meeting you can see the conceptual object taking shape but you cannot find it in the bits and pieces making up the discourse. There are, of course, instances where the design or solution does come from one person, but then you have a different kind of meeting, one that is devoted to grasping, accepting, and elaborating an idea. The result is still a social product, no matter how much it may bear the stamp of an individual. A 'knowing organization' (Choo, 1998)—whether it is a business or a school—needs to find ways to foster individual knowledge creation (I will say more about this in Chapter 10) but it also needs to find ways to foster the kind of knowledge that emerges from discourse, broadly conceived. Without declaring one to be more important than the other, we can reasonably conjecture that progressive knowledge-creating discourse falls more within the scope of things that organizations can do something about than does individual genius.

Nonaka and Takeuchi offer their model as an attempt to convey to Westerners a Japanese view of knowledge. They see Westerners putting all the emphasis on explicit knowledge and placing it in competition with fuzzy tacit knowledge, whereas the Japanese see tacit knowledge as primary and complementary to explicit knowledge (1995, p. 61). If there ever was validity to that comparison, it is fast vanishing. In *Intellectual Capital,* Stewart (1997) places enormous emphasis on the tacit knowledge of individuals and groups and cites scores of supportive examples

and testimonials, mostly drawn from American business.[3] Nonaka and Takeuchi, I suspect, have been overly influenced by Western philosophers and their concern with the validity of explicit propositions, and have failed to appreciate the extent to which folk epistemology recognizes the virtues of intuition, faith, gut feelings, and knacks. The problem is how to do anything about these—how to promote, guide, and use them effectively. Educators probably recognize them more fully than most people, because of their close engagement with students' thinking and learning, and yet they count for next to nothing in educational policy and curriculum planning.

Folk epistemology, I have said, makes tacit knowledge mysterious. Nonaka and Takeuchi do not make it any less mysterious. They have practical advice about how to disseminate tacit knowledge and how to covert between tacit and explicit knowledge, but they do not leave us any wiser about what tacit knowledge is or how to advance it to qualitatively higher levels. The gaping hole in their model is where we would look to find an answer to the main problem facing a knowledge-based organization: How do you develop the tacit knowledge that enables the creation of explicit knowledge? This is not the same question as how you acquire the tacit knowledge that is converted into explicit knowledge, which is a question Nonaka and Takeuchi do address. It is a question that only makes sense if you can conceive of conceptual artifacts as real things that people produce, drawing on varieties of tacit knowledge and using other conceptual artifacts in the process.

Of the two radical conceptual changes proposed at the beginning of this section, the shift to treating abstract knowledge objects as real things should not be difficult for people in knowledge-based businesses. They are familiar with the idea of intellectual property, and the idea of conceptual artifacts is only a somewhat broader version of that idea. That idea seems to be much harder for educators to hold on to. It keeps slipping over into learning, to changes taking place in individual minds, because that is what education is mainly about. The shift to a connectionist view of mind, however, is difficult for everyone. To make it comprehensible and functional takes considerable

---

[3] And by the way never cites Nonaka.

concentration and practice. In the next section I work through three examples that I hope will be helpful in making sense of the idea of knowledgeability without mental content.

Mind Without Mental Content

Folk theory never tells us what the mind is. That seems to me a sensible evasion and one to be maintained. Mind is just what things come to when they are said to 'come to mind.' Mental content becomes a mischievous idea when it refers to things that we do not have 'in mind' but that are nevertheless assumed to be back in the mind somewhere, ready to be put to use or already covertly at work. Folk theory of mind assumes an infinitude of unattended beliefs which guide action. You are said to believe that your coat is on the hook in the closet, through all the hours in which you do not think about your coat at all. The evidence for this is that when you want it that is where you look, and if it should not be there you will say something like "I was sure this is where I left it." The reason you look for the coat is that you believe it is cold out and that the coat will keep you warm, even though neither of these beliefs actually comes to mind. Behind these are unattended beliefs in object permanence, gravitation, heat (or cold) transfer, and everything that one would have to explain to an intelligent being that had no knowledge of this world. Although having a virtually infinite number of beliefs is implausible, it provides the premise for plausible and nontrivial explanations of behavior. Why, when your coat is not on the expected hook, do you look on other hooks and also on the floor but not on the ceiling? Folk theory of mind can not only explain this but could predict it, based on beliefs that it is reasonable to attribute to you.

Folk theory undoubtedly took shape to deal with just such mundane and commonsensical behaviors as those described. The same model has then been carried over into education and into the more studied and creative uses of knowledge. Education, on this model, is concerned with adding to and amending the contents of the storehouse of beliefs. Although this suggests a crude, 'stuffing the head' conception of education, the model underlies more refined conceptions as well, such as those of constructivism, values education, conceptual change, central conceptual structures, and whole language. Indeed, it is only the most hare-brained conceptions of education that disavow an

interest in content. Therefore, if mental content is to be eliminated from mentalistic descriptions, or if it is to be reduced to the status of an ungrounded metaphor, there is a serious question as to how education's unavoidable concern with individual knowledge acquisition is to be honored. The rest of this book addresses the challenge implied by that question, the challenge being to improve upon, to go beyond what folk theory has to offer education in its concerns with knowledge.

For the present, however, I avoid education and other weighty areas of application and instead deal with cases chosen for their illustrative value. The first case comes from well outside most people's serious concerns. It is dreaming. Dreaming is an unquestionably mental phenomenon, traditionally but, as it appears, wrongly thought of in terms of stored mental content. A contentless approach provides an account of dreaming more closely in accord with current scientific understanding of dreaming. The second case is a trivial example of misperception, which I analyze largely through introspection to show how an adequate account can be developed that restricts the mentalistic part of the account to what is actually experienced. The third case illustrates what is meant by having a relationship to an abstract knowledge object. Again, to avoid weighty issues I have chosen a minor kind of object, the joke, which nevertheless captures the essentials of cognition in relation to conceptual artifacts.

Making Sense of Dreams

Dreams are clearly a part of mental life, and dreams have content. But, with occasional exceptions, the content does not come from outside, via our senses, it comes from within. Dreaming therefore makes a nice first case for trying out the notion that we can account for mental phenomena without positing a storehouse of mental content. Aristotle (in Ackrill, 1987, pp. 214-217) speculated that dreams are a kind of aftereffect of waking experience, much like the visual aftereffects of staring at a bright object. That is an interesting conjecture, for it steers clear of implying that the mind is a container of past experiences. But it falls far short of accounting for the rich and novel content of dreams. The mental filing cabinet model does not do much better, however. Dreams are vividly visual, but even if one allows for some powerful mental

morphing, it is implausible to imagine them as composites of old film clips.

Modern explanations of dreaming start from electroencephalographic data, which reveal that, during REM (rapid eye-movement) sleep, sensory areas of the brain are more active than in the waking state. Evidently there are inhibitory centers in the brain that discourage sensory neurons from firing except in response to incoming stimuli; but during sleep the inhibitory centers shut down and allow the sensory areas—especially the visual—to act at random (Hobson, 1988). What about the meaningful, narrative quality of dreams, then? Why aren't dreams disconnected bits of visualization? The coherence of dreams, according to Hobson, arises from the thinking brain's efforts to make sense of perceptions. The brain, in effect, is making sense out of nonsense. It can't entirely succeed; hence the occasional bizarre jumps and transformations. But it does remarkably well, considering what it has to work with. Although this may seem a far-fetched theory, there are other evidences of the brain's relentless sense-making. In split-brain patients (people whose right and left cerebral hemispheres have been surgically disconnected) stimuli to the right hemisphere will cause the person to act for reasons that the left hemisphere has no knowledge of. Yet when questioned, such people will give confident although wholly fabricated explanations of their behavior (Gazzaniga, 1995; Gazzaniga & LeDoux, 1978). The kind of sense the brain makes of dream sensations will, of course, be influenced by past experiences, by fears and desires—by the same internal conditions that affect waking perception; hence the relevance of dream experience to our waking preoccupations and recent experiences.

Although this account of dreaming greatly simplifies what recent research is showing to be a very complex and also highly variable phenomenon, the main point I want to make from it is not called into question by any findings I am aware of. The idea that dream research drives home so dramatically is the human mind's relentless and largely automatic effort to make sense of things—an effort that commonly takes the form of weaving things together into a coherent narrative.

In folk theory, the meaning of events is either obvious (that is, conforming to expectations) or else it is problematic, to be

worked out through deliberate efforts to explain. Essentially the same view is embodied in information processing models, such as that of Schank and Abelson (1977), in which normal events are understood in terms of their conformity to stored scripts and explaining is driven by failed expectations. This seems quite valid as far as it goes, and it seems to accord well with the part of sense-making that we are conscious of. Dreaming, however, does not fit these categories. To the extent that the sensory events that occur during REM sleep are random, their meaning cannot be said to be obvious; yet the sense-making that occurs in dreaming is spontaneous, not a matter of effortful seeking of explanations. Just such automatic sense-making goes on throughout our waking hours. The conspicuous difference is that when we are awake sensory events having origins outside the brain predominate, but it is surely not a complete switch-over.

Reading provides the most striking evidence that automatic sense-making during the waking state is not altogether different from what goes on in dreaming. There is abundant evidence that our interpretations of what we read are driven by expectations (Anderson & Pearson, 1984; Rumelhart, 1980). A classic experiment used the following text:

> Rocky slowly got up from the mat, planning his escape. He hesitated a moment and thought. Things were not going well. What bothered him most was being held, especially since the charge against him had been weak. He considered his present situation. The lock that held him was strong but he thought he could break it. He knew, however, that his timing would have to be perfect. Rocky was aware that it was because of his early roughness that he had been penalized so severely—much too severely from his point of view. The situation was becoming frustrating; the pressure had been grinding on him for too long. He was being ridden unmercifully. Rocky was getting angry now. He felt he was ready to make his move. He knew that his success or failure would depend on what he did in the next few seconds. (Anderson, Reynolds, Schallert, & Goetz, 1977, p. 372)

Most people read this as being about a contemplated jailbreak. But when it was presented to members of a wrestling team, most of them read it as being about a wrestler thinking about how to

break a hold. Of equal interest, however, is the fact that, whichever interpretation people gave to it, the interpretation seemed obvious. It was not arrived at through a consideration of alternatives or a deliberate effort to resolve ambiguities.

When expectations drive reading beyond a certain point, we get demonstrable misreadings; beyond that, something verging on hallucination. Among students, misreading is so common that teachers at all levels regard it as a significant part of what they must deal with. The misreadings may readily be explained by what the students do not know or by the interpretive frameworks they bring to texts. What is most relevant to the present discussion, however, is the absence of feelings of confusion or incomprehension. Asking students to raise questions about parts of a text that were unclear to them usually produces an uneasy silence. The teacher who fancies a Socratic method must go to considerable lengths to get students to recognize that they have not in fact made good sense of what appears on the page.

The more hallucinatory misreadings become evident only when readers can be presumed to have ample background knowledge and skills. My experience in submitting scholarly articles to peer-reviewed journals is that the incidence of glaring misreadings among qualified and supposedly careful peer reviewers is not much different from that to be expected of students in an introductory course. It could be, of course, that the fault is in the manuscript, and that it is my sense-making that is hallucinatory. Evidence that what I describe is not an isolated phenomenon can be found, however, in those journals that publish peer commentary along with the 'target' article, usually followed by a response from the original author. I have been struck by the extent to which such responses consist of allegations of misreading—particularly complaints that the commentators criticized the article for things that were never said or that were at times the opposite of what was said.

Where reading most closely resembles dreaming, however, is in those cases where readers' anxieties are aroused. In recent years I have been involved in several different drawn-out exchanges of correspondence with people who were in high states of anxiety on matters of money, occupational security, or status. These people would respond with alarm or outrage to

proposals that were nowhere to be found in what was actually written. All that kept these responses from being utterly fantastic was that they centered on certain keywords that did appear in the texts.

It is not stretching a point terribly far to say that mental life consists of controlled halluncination. Our hard-wired perceptual systems provide the strongest sorts of controls, ensuring that we seldom see things that are not there or grossly distort or misrecognize things that are there. Reflective analysis provides another sort of control. If you ever find yourself uncertain as to whether something really happened or you only dreamed it, you are likely to resort to a consistency check to see whether it fits with other data in which you have more confidence. Finally, social intercourse provides strong and pervasive controls. The surest way to determine whether something really happened or was a dream is to check it with someone else who would know about it if it really happened. The need to coordinate our actions with those of others and to establish shared topics and common ground in conversation means that a great deal (sometimes too much) of what our minds have to make sense of is already highly processed material. But there is still a wild, internally generated flow of data that mixes indiscriminately with sensory inputs, leading at times to imaginative leaps, at other times to craziness, and probably most of the time to mere noise that gets factored out as we make sense of the more interpretable data coming from other sources.

I think I am in accord with a number of recent philosophers of mind when I say that this unremitting sense-making, carried on during our sleeping as well as our waking hours, is really what mind amounts to (Dennett & Kinsbourne, 1992). The mind is not some organ which does the sense-making. That role belongs to the brain. The mind is a product of our sense-making activity. It is what our sense-making postulates when it tries to make sense of itself.

Sense-making, in the view of most people who use that term, is largely a matter of making up stories (Bruner, 1986; Egan, 1989; Dennett, 1991). To say that sense-making is automatic is to say that our making up of stories is, most of the time, involuntary and effortless. (This obviously applies to dreaming, for dreams are almost always stories, however bizarre; and when they are

not—when the fevered brain fires away with results that cannot be fitted together into any kind of narrative—we awake tired and vaguely distraught.) A story requires actors and goals and usually a setting of some kind. In our waking as well as our dreaming states, we are usually central figures in our stories. So what kind of story do we create when the story is required to make sense of our sense-making itself? The central figure in the story is, of course, our self, but it has to be a self that can stand aside from the corporeal body and its sensations and rememberings so as to make sense of them. Hence the mind is this peculiar character that is at once indissociable from the physical being and yet able to view and speculate upon it. What about a setting? Descartes posited an internal theatre in which percepts and ideas do their turns while the mind observes and makes sense of them. In common usage, the mind is both the theatre and the observer. Cognitive events occur *in* the mind and are at the same time interpreted *by* the mind—a confusing notion which common sense tends simply to gloss over. Philosophers have puzzled over all this at great length. What I am suggesting (and I take this to be in accord with Dennett's [1991] more strongly reasoned proposal) is that we not press too hard for coherence in the consciousness story. The endless story that we construct to make sense of our lives must inevitably include the author as actor, object, observer, and setting, and there is only so much coherence that you can expect in a story like that. The important thing, from an educational standpoint, is just to recognize that sense-making happens regardless. It is the starting point of educational processes, not the ending point.

Knowledgeability without a Mental Filing Cabinet

The following example will serve to illustrate automatic sense-making and the role of knowledge in it. I have chosen a mundane example and will interpret it first in the conventional way, in terms of knowledge items in a mental storehouse. Then I will show how it can be interpreted without assuming stored mental content. My purpose at this point is not to show that the latter approach works better. For ordinary purposes the conventional approach works fine. My purpose is just to show the difference and to show that the second approach is at least convincing and not so outlandish as might be supposed.

> Looking out the window one wintery day, I saw in a nearby tree a small bird of a peculiar reddish brown color. So unusual was its coloring that I went immediately to an illustrated bird guide to see if I could identify it. Finding no immediate match, I went to an upstairs window from which I could get a better look. From this new vantage point it became immediately obvious that it was not a bird at all but an apple, which had lodged in a branch instead of falling to the ground and had rotted to this reddish brown color.

How could I have been so certain initially that it was a bird, and what made it so immediately obvious, viewing from a different angle, that it was not? After the fact, I can adduce reasons that make my cognitive behavior appear quite rational:

- Small solitary colorful objects in the branches of trees are usually birds.
- To protect themselves from the cold, small birds will fluff up their feathers and draw in their heads so that they are approximately spherical in shape.
- One does not expect to see an apple in a barren tree in the middle of winter.
- Viewed from above, a head should have been visible if it were a bird; there being no head, it could be inferred that it was not a bird.

Two quite different psychological accounts can be given of my mistaken recognition of a bird, one supporting the old and one the newer conception of mind. (To simplify, I am leaving out some intermediate kinds of accounts, such as one based on a bird schema). According to the older conception, items of knowledge such as the above were actually stored as content in my mind. What I experienced as the direct perception of a bird in a tree was actually the result of a lightning-fast and unconscious reasoning process, more or less similar to my after-the-fact justification.

The other kind of account starts with the premise that our brains are built in such a way that visual stimulation having certain properties of contrast and contour will result in perceiving an object. *Not* to have perceived an object in the tree, to have seen only patterns of light and color, would have required special training of a kind that artists sometimes undergo. There is nothing mental about seeing the patch of

reddish brown color as an object. It is not the result of a cognitive process, anymore than it is the result of a cognitive process when my brain sends signals to my heart to beat faster when I am climbing a flight of stairs. It is just how the brain works.

Now, what about seeing the object as a bird? Why not see it as a Christmas tree ornament or a wad of paper or simply as an unidentified reddish brown thing? Here, obviously, effects of past experience are operative and those effects *could* be described in the form of propositions like those listed. But that would lead us back to the first kind of psychological account. Instead, let us suppose that these effects take a form analogous to that of the weights in a connectionist network. That is, they take the form of a large system of relationships among neurons such that perception of the reddish brown object will be, as it were, attuned to the time of year, the features and habits of small birds, and so on. As a result, I do not first perceive an object and then figure out that it must be a bird, I do in fact directly see a bird. Much more than that, actually: I see a bird huddling from the cold on a barren branch. My sympathies are aroused. Also, I am struck by the bird's beauty. And my curiosity is aroused because I do not just see a bird, I see an unknown kind of bird. Mental activity *starts* with this affect-laden awareness of a strange, beautiful bird alone in a bleak northern winter. That is the first mental content to enter the picture. It does not start with a set of rules in the mind that fire off in rapid order to produce this mental state. Discovering my error does not involve undoing a chain of inference that led to a wrong conclusion. Instead, looking at the object from a different direction, I directly see an apple instead of directly seeing a bird. Different visible features cause my neural network to settle on a different perception. When I return to look out the original window, I no longer see a bird, I see an apple that looks somewhat like a bird. None of this is mental either. What is mental is my realization of having made a mistake and being amused by it and starting to think out how I could have made such a mistake.

Regardless of which way it is interpreted, a considerable amount of knowledge would seem to have been involved in the bird / apple episode, albeit knowledge of a commonplace sort: knowledge of what birds are and something about them, that a bird cannot turn into an apple, that apples may be found in trees,

that apples turn brown when they rot, and so on. But none of this knowledge was consciously recalled and employed in reasoning. What is this knowledge—and *where* is it—if it is not mental content? The challenge for a new theory of mind is not to answer this question but to find a graceful way of avoiding it. Unfortunately, the available ways are not very graceful, and so in order to carry on a fluent discussion we are driven to treating knowledge as stuff people have, which suggests that it must exist in some form and be somewhere. That is what we did in the preceding chapter. But when we are trying to achieve conceptual clarity, we are better off not talking about knowledge in that way at all and trying to speak instead of *knowledgeability*.

The bird / apple episode, to put it in this other way, reveals some knowledgeability on my part concerning birds and apples. Although propositions such as "Small solitary colorful objects in the branches of trees are usually birds" and "To protect themselves from the cold, small birds will fluff up their feathers and draw in their heads so that they are approximately spherical in shape" never actually entered my mind, I responded and behaved in ways that were consistent with them. Accordingly, a characterization (whether delivered by myself or an observer) of my knowledgeability with respect to birds might well include those two facts. Further justification for doing so might come from my exhibiting other behavior consistent with them and from my answers to questions such as "How could you recognize a bird in a tree?" and "What do small birds do to protect themselves from the cold?" Yet I may never have thought of those facts until asked questions that called for articulating them.

Of course, those facts are just a sample of the bird-related facts with which my behavior is consistent. My behavior is consistent with an infinitude of such facts and principles, this despite my level of ornithological sophistication being, I suspect, below average. Ornithologists and experienced bird watchers will act in ways consistent with strata of facts and principles far beyond those pertinent to my behavior. They can not only identify many more birds, but can identify their nests and eggs. They will notice characteristics of birds and their behavior that I overlook and they will be able to offer explanations beyond anything I would think of.

The common way of wrapping up these contrasts is to say that ornithologists and bird watchers have much more bird knowledge in their heads than I do and that this knowledge enables them to do a variety of things that I cannot. Thus there are two elements in the account, knowledge and abilities, with the first being the cause of the second. If we eliminate the idea of mental content, however, these elements are reduced to one—knowledgeability. A person's knowledgeability can be described in various ways. One way is to describe it in terms of content, as if it were a book. Another way is to describe it in terms of abilities. But these are different descriptions of the same thing, and even in combination these descriptions do not exhaust what might be said in characterizing a person's knowledgeability. The itemization of content will capture the more textbook-like features of a person's knowledgeability but will trail off into metaphor and then inarticulateness just when is starting to get interesting—that is, when it is starting to touch upon the more profound aspects of a person's understanding. A description of abilities will suffer the same fate. There will be some obvious abilities, easy to describe, but when it comes to the subtleties that distinguish one expert from another, description will begin to fail. There is nothing occult about this. Similar difficulties would arise in trying to characterize a painter's or a poet's or a philosopher's *oeuvre* or to describe teaching as it is carried out by a masterful teacher. In all of these there is more than can be captured in a description or in several descriptions done from different vantage points. What we need to be clear about is that the boundaries of the mental include knowledgeability. Items of specific knowledge may enter into our descriptions of knowledgeability, but they are not components of it in anything like the way that pages are components of a book or premises are part of a syllogism. You can to a certain extent describe a pain, and there are situations (such as seeking medical assistance) in which it is desirable to do so; but the description is not a component of the pain. So it is with descriptions of knowledgeability.

Understanding a Joke

Jokes are among the clearest examples of abstract objects. We recognize the same joke even though it is told differently at various times by various people. Jokes have histories. A joke that

may have first been recorded during the Crimean War will turn up with suitable updating in subsequent wars. An ethnic joke will be adapted so as to ridicule a different group. When no longer socially acceptable, it will be altered to remove ethnic allusions and dialect, but will still be recognizable as the same joke. There are of course boundary problems: Are two examples variations on the same joke or different jokes? But these problems are not serious enough to discourage us from such worthy efforts as *telling the same joke that we found so hilarious at last night's party, but modifying it to make it suitable for children*.

Although jokes may be written down and preserved in books, the joke is not any particular literal version. It is the abstract thing that may be variously told, written, sometimes dramatized, sometimes merely referred to as "the one about...." (Jokes, it should be noted, are not *conceptual artifacts* as we have been using the term. An explanation of a joke would be a conceptual artifact—it could be debated, applied, and so on; but, as is well known, the explanation and the joke are two quite different animals.)

Without getting into the fascinating psychology and rhetoric of jokes, we can make do for present purposes with a few commonplace observations:

- To tell a joke well you must understand it.
- Understanding is not purely intellectual. If you do not 'see the humor' in a joke you cannot be said to understand it.
- Although liking a joke without understanding it is improbable, you can dislike a joke somewhat independently of understanding it. Distaste for the joke's theme or content may override appreciation of its humor.
- 'Seeing the humor' or 'getting the point' of a joke is not necessarily accompanied by ability to explain it. People demonstrate their understanding of a joke by laughing in the right way at the right time.
- Explanation is seldom helpful in getting people to see the humor in a joke. Thus, understanding a joke is rather like perception. Either you see it or you don't.
- You can continue to appreciate the humor of a joke even though you have heard it enough times that it no longer makes you laugh; you can delight in telling it or in hearing it well told, and suffer in hearing it badly told.

- Sharing jokes is a common form of social bonding. ‘Inside’ jokes serve to separate insiders from outsiders.

Notice that in all these commonsense observations, the joke is treated as something existing in its own right, to which people have various responses and attitudes. Nothing is implied about content in people’s minds. Sharing a joke is like sharing a bottle of wine or enjoying a concert together. It does not depend on two minds having the same content. It depends on two minds (and bodies) being attuned to respond in concordant ways to something in the external world. And yet with jokes—as with wine and music—there is an important intellectual component. To respond in a concordant way, people have to have a fair bit of background knowledge in common. But we could not hope to pin down the essential elements of that knowledge. Let us say, rather, that the shared knowledge is *in* the attunement.

What I have been arguing for, especially in Chapter 3, is treating conceptual artifacts of all sorts in the same way that we treat jokes, legends, and other kinds of abstract cultural objects—as existing out there in the cultural environment. Our minds do not contain these objects. Our minds have various abilities, attitudes, dispositions, habits in relation to these objects.

Memory is probably the hardest aspect of cognition to bring under this idea of the contentless mind. Jokes, again, provide a nice test case. Most people, it seems, have trouble remembering jokes, despite a desire to do so. You hear a good joke and try to remember it. You may even think of yourself deliberately putting it into a mental filing cabinet, perhaps even with notes attached as to where you might use the joke or who you would like to tell it to. But when you go to retrieve it the joke is missing; in fact the whole joke drawer of your mental filing cabinet seems nearly empty. But there may be a few stray items in it, so that if required to tell a joke you can bring forth something, although it may be some childish joke of decades past and quite unsuitable to the occasion.

What does it mean to remember a joke? It does not mean remembering the exact words. But it does not mean just remembering the gist or the punchline, either. You may think you remember a joke and will start telling it, only to find that you have said something that spoils it. Remembering a joke heard

only once is more like remembering a route that you have traveled only once before or remembering how to carry out some complicated procedure that you have practiced only once or twice. It is a task of reconstruction and also of execution. At your leisure you may be able to mentally reconstruct the joke or the route, but when it comes to actually telling the joke or driving the route, you make mistakes. If you *really* remember a joke, which typically means that you have told it on a number of occasions, then, as with a familiar route or procedure, you can not only reconstruct it easily but you can vary and improvise without losing hold. The ability to reconstruct and execute, therefore, is what your mind has or what your knowledge consists of. This ability depends, of course, on alterations in your brain brought about by previous experience, but we are in no position to say precisely what those alterations are—nor, as educators, should we have to care. As long as the learner's brain is intact, we need hardly concern ourselves at all with what goes on in it. Instead, we can turn our attention to the outer world, to the material and immaterial objects it contains, and to the abilities and dispositions of the learner with respect to those objects.

Putting It Together

The preceding examples—dreaming and overly imaginative reading, mistaken perception, and the telling and appreciation of jokes—are all intended to show how everyday cognition makes more sense if we abandon the idea of a mind operating on stored mental content and replace it with the idea of a mind continually and automatically responding to the world and making sense of whatever befalls it. I call this the 'connectionist view of mind' because connectionist AI currently provides us the only way of imagining how such a mind might work. By itself, the connectionist way of talking about the mind offers advantages that are offset and perhaps exceeded by its disadvantages. One advantage, as I discussed in Chapter 2, is that it enables us to talk about feelings and cognition in the same breath. When I saw what I took to be a bird in the apple tree, I had an immediate esthetic feeling—I saw the bird as peculiarly beautiful—and also feelings of sympathy for its solitariness and misery. But thinking it was a bird was not one kind of event, which then gave rise to the feelings. They were all part of one global response, triggered by what I saw, conditioned by all my past experience, and

regulated by my whole constitution, which is itself a complex of interacting systems. For decades, educators have been struggling to establish a more wholistic view of learning, signaled by expressions such as “teaching the whole child” and “whole language.” Folk psychology offers no support for such efforts, and so in this regard the connectionist view of mind should be welcomed.

The other big advantage is that the connectionist view of mind enables us to talk constructively about kinds of learning that elude efforts to treat them as rules or propositions or pictures in the mind: number sense, moral sense, esthetic sense, a way with words, the effects of immersion in great literature—in short, just about everything that distinguishes a well-educated person from a product of rote learning. The failure of folk psychology to encompass these kinds of learning must take much of the blame for the fact that, against practically everyone’s better judgment, we are still stuck with tests that almost exclusively emphasize mastery of facts and rules, which in turn drive curricula with the same emphasis. How much can actually be gained from adopting a connectionist view of learning remains, however, to be demonstrated in later chapters.

These advantages are offset by a lack of explicitness about what a person knows. Folk psychology encourages us to believe that when we say Jack does or doesn’t know *x* we are making much the same kind of assertion as when we say that Jack does or doesn’t have a Harris tweed coat.Of course, it is much easier to inspect the contents of Jack’s closet than to inspect the contents of his mind, but the principle is the same: If we had access, we could supposedly itemize the contents of Jack’s knowledge repository just as we could itemize the contents of his wardrobe. Far from being ridiculous, this is an extremely useful notion. I don’t see how teachers could ever get along without it. When Jack is in trouble scholastically, a smart thing to do is to inventory his relevant knowledge and when you find something important to be missing, teach it to him. This is the cornerstone of direct instruction (Bereiter, 1968), which I will discuss in Chapter 8 and which, by the way, all teachers resort to whenever they are prepared to suspend ideology in favor of getting something across. The connectionist way of talking about

knowledge and mind does not lend itself to any such direct procedure of identifying what needs to be taught and teaching it.

For that reason, along with other equally down-to-earth reasons, I would never suggest doing away with folk psychology. Like folk physics and folk biology, it serves our daily needs by providing rough-and-ready connections between observation and action. It is just that we should not be limited by it. However, I am not suggesting, either, that the connectionist view of mind should only be activated for special purposes. Quite the contrary. I think it should constitute the fundamental way that we think of the mind, even while for practical purposes we treat the mind as a container of rules and beliefs. It is the same as with cosmography. For many practical purposes, such as planning a garden or positioning a patio umbrella, we must think of our patch of earth as stationary and of the sun as passing overhead. Problems would be too difficult for us otherwise. But even while doing so—assuming we have really absorbed the Copernican view of the world—our basic attunement is to a rotating earth and a stationary sun.[4] We do not have to remind ourselves of it when we read about space travel or giant ice balls entering the atmosphere. The Copernican view has become an integral part of *the way we think*. It was not always thus. It is probably not true for most of the world's population, and I suspect that even for educated Westerners the Copernican view was not really incorporated into their world views until the advent of intercontinental air travel.

The connectionist mind is much more directly in contact with the world than is the symbol-processing mind conferred upon us by folk psychology. Being 'more directly in contact with the world' does not, of course, mean bypassing sense organs and neural pathways. Remember, it is the mind, not the brain we are talking about, and so the directness of contact is psychological. I like to think of psychological directness along the lines that Joseph Church (1961), a phenomenologist, talked about the experience of writing with a pencil on rough paper. The

---

[4] The sun also moves, but except for astronomers this fact is of such limited significance that it need not become part of out cosmological attunement, that is our functional image of the cosmos. It is enough that be able to recall the fact on the occasions when it is relevant, whereas we should not have to recall that the earth moves around the sun. That should always be part of what Wittgenstein called the 'scaffolding of our thoughts.'

sensation comes to us through vibrations of the pencil moving over the paper. But we do not experience the vibrations of the pencil. Instead, we feel the texture of the paper directly. Similarly, according to the connectionist view, we do not receive a pattern of visual stimuli which we interpret or recognize as a bear. Our visual system sees things. Given certain inputs, we will see a bear. There is similar directness with respect to action. In the symbol-processing mind, the belief that a bear is approaching activates fear and a variety of escape schemata, while the connectionist mind has us clambering up the nearest tree.

Putting it more generally, the connectionist view of mind is at its best in making sense of people interacting in real time with things in the external world. The symbol-processing view interposes too many unlikely mental events. The symbol-processing view is best suited to rationalizing such behavior after the fact—to explaining, for instance, how one might reason one's way into climbing a tree to escape from a bear. In this case, the rationalized version probably bears little resemblance to the actual mental events, which are likely to have been a jumble of thoughts and impressions accompanying rather than causing the physical actions. How do we explain someone's climbing a tree to escape a bear when they know that bears can climb trees? For those who adopt the symbol-processing view, such behavior raises questions about human rationality. But from a connectionist standpoint, there is nothing paradoxical about it. Given the real-world situation, there are only a few alternatives on which the cognitive system would have much likelihood of settling: confront the bear; cower; turn tail and run (with the prospect of being seized from behind); step aside to let the bear go past; and climb the tree.[5] Assuming there was no strong bias from prior experience, which option prevailed would likely depend on particularities of the situation. The sight of a reachable branch could well tip the balance in favor of tree climbing, whereas an only slightly higher branch or a slightly less menacing bear might tip it in the direction of one of the other

[5] These options do not include what I have heard is the smart thing to do on meeting a bear, squat down so that the bear does not perceive you as challenging it. The behaviors I listed are all ones we are likely to have acquired in our encounters with menacing human beings. They may be thought of as templates for the rapid construction of a response in the bear situation. Our squatting template, however, is attuned to entirely different kinds of situations and so is unlikely to come into play unless we have been instructed.

possibilities. Similarly, having previously seen a bear climb a tree (as opposed to having only heard that bears can do so) would no doubt greatly lower the attractiveness of the tree-climbing alternative.

Most of the concern in education and knowledge work, however, is not with split-second decisions. It is with situations in which we have time to reflect on what to do, to recall relevant knowledge, to weigh alternatives. In such cases, a symbol-processing approach often works well. That is because, when thinking carefully, we try to rationalize or justify each idea as we go along. How the ideas come about is often not of much concern. They may already be laid out for us as choices to be made, and so we are wise to focus on the rationalizing—on the checking for consistency, the scanning for alternatives, the building up of a convincing story about why what we think is right. But to build an educational program or a knowledge management system on the basis of such a limited view of cognition is to ignore a very large part of what makes human beings different from logic machines. It is not just that human beings have feelings. If you could program a machine to get angry when it is thwarted and sad when its plans fail, that would not make it into a machine that can create knowledge, fall in love with ideas, develop number sense and historical sense and a literary style that goes with its personality, and become dissatisfied with the shallowness of it understanding. Connectionist simulations cannot do these things, either—at least not yet—but they occupy a design space in which such things are reasonable to contemplate as programming goals. That should be enough to make us question folk theories according to which all the distinctly human aspects of intelligence are mysterious and elusive and the only part that seems comprehensible is the part we share with machines.

Self-Organization and Emergence

As I noted in Chapter 2, connectionism is one part of a much larger research movement aimed at providing scientific explanations of emergent phenomena. *Self-organization* is an idea that ties together many different research programs carried out at various levels of description from the atomic to the cultural. Adopting the way of thinking about knowledge and mind that I have been trying to put across in this chapter

amounts essentially to acquiring a mindset that sees learning, thinking, knowing, and the creation of new knowledge as forms of self-organization.

Common sense and whatever biological heritage underlies it have not prepared us for the idea of self-organization. When common sense sees a puzzling phenomenon it looks for a causal agent. When it sees organization it looks for an organizer. This works amazingly well for purposes ranging from the diagnosis of diseases to the creation of governments. But it cannot account for emergence, which E. O. Wilson (1998, p. 86) defines as "the appearance of complex phenomena not predictable from the basic elements and processes alone." Pricing in a free market is perhaps the one example of emergence that is widely recognized. No causal agent or higher authority sets the prices. Prices emerge from the behaviors of buyers and sellers, but you could not predict the price of cucumbers through an analysis of those individual behaviors. A monkey colony exhibits a clearly defined and functional hierarchical structure, but the monkeys did not get together and draw up an organization chart. The complex social organization emerges from the interactions of monkeys, each of which is executing much simpler behaviors of dominance and submission. Human reason tends toward five-year plans and organization charts. Self-organization tends toward farmers' markets and the informal social networks that exist within and often in spite of bureaucratic structures.

Not surprisingly, then, the idea of emergence has been the darling of romantics in their opposition to what they often seen as a campaign by scientists to reduce life's mysteries to chemical reactions. What must be understood about complexity theory, chaos theory, self-organizing systems, and the like, is that this is not the romantics starting to win, it is the scientists assaulting yet another bastion. It is not the reinstatement of mystery, it is an important new phase in science's unending program of demystification.

Self-organization and emergence are the rule in nature, and it can hardly be otherwise, unless you introduce the guiding hand of a deity. New structures have to arise from what already exists. The complexity has to result from structures and processes that do not themselves embody that complexity. Understanding self-organization always requires that we

consider two levels and try to understand how the lower level phenomena can produce the higher level phenomena. To understand knowledge and mind in this way, we need to consider four different kinds or levels of self-organization:

1. *From neurons to mind.* This is the self-organization that takes place, as Hofstadter (1985, p. 649) put it, "when thought emerges from billions of in-themselves-meaningless neural firings." It is what connectionists are trying to model. Common sense has to posit a mind that sits like an executive on top of the brain, but then leaves the executive as something that cannot be explained. Connectionist models try to show how intelligent, knowledgeable behavior could arise from a brain that does not itself contain stored propositions and rules.

2. *From individual behavior to social organization.* How can ants have such complex social behavior, waging military campaigns and the like, and in some species building elaborate dwellings, when no individual ants have anything like the intelligence or knowledge that seems to be exhibited at the group level? Mitchel Resnick (1994) has provided striking computer demonstrations of how self-organization can work at the social level. One of Resnick's programs produces virtual ants, each of which embodies a simple set of rules that cause the ant to scurry about the computer screen, nibbling food that it encounters and laying down a trail of scent. When there are only a few ants, their behavior appears random and disorganized. But once the ant population gets into the hundreds, you start to see systematic antlike behavior. They follow one another around in a line. They all nibble away at one food object until it is devoured and then spread out until another is discovered. This is a particularly elegant demonstration of behavior that emerges at the group level. Somehow, what accounts for the complex social behavior of ants or people must be embodied in the behavior systems of individuals. There is no collective nervous system. But the individual does not have to embody the complexity and the purposefulness exhibited by the society. It only has to embody capacities and dispositions that, in interaction with others of its kind, produce the systematic social behavior as an emergent. Thus, each ant has the disposition to move toward the scent left by other ants. When there are few ants this has little effect on behavior. But when there are many ants each moving toward the

strongest scent and at the same time laying down scent of its own, trails of scent begin to emerge. And, since each ant also has a tendency to move toward food, the trail will tend to lead toward food. Virtual cows would create cattle paths in much the same way, behaving according to different but equally simple rules. In human societies, self-organization is obscured by the involvement of deliberate planning. To see it we have to look in the interstices—at the black markets that emerge, for instance, when governments try to regulate economic activity too closely. Yet we must assume that self-organization is at work within deliberate planning and management processes as well. Social systems acquire complexity of their own, beyond that embodied either in the components of the system or in the plans and policies meant to control it.

An important characteristic of self-organization of social behavior is that the low-level components entering into self-organization include not only the individual organisms but also things in their physical environment: the size and location of food objects for the ants, the shape of the terrain for the meandering cows, the hideability and transportability of goods for the black marketeers.[6] From the standpoint of self-organization, therefore, what we have is not organisms behaving in ways that are adapted to the environment. Rather, behavior is an emergent phenomenon of organisms interacting with each other and with the physical environment. Where individual behavior is concerned, this amounts to saying the same thing in different ways, but when we are talking about the emergence of social behavior—of teamwork, of economic transactions, of feuds and warfare, of customs—the self-organization view allows us to make sense of things in ways that the behavioral adaptation view does not.

---

[6] Along Spadina Avenue in Toronto's Chinatown, there are open-air stands facing out on the street in front of nearly every shop. On the street side there are much smaller stands facing inward, mostly selling produce. One day as I was walking up Spadina Avenue, a sudden flurry of activity swept through the street-side vendors. Goods were stuffed into shopping carts, tables were folded, and within a minute the vendors had disappeared into the crowd. I had previously wondered at the small quantitites of produce that individual vendors had on display—not enough to make a living from, it seemed. Now the answer was obvious. Unlicensed street vending, which undoubtedly has a complex social structure invisible to the passer-by, must accommodate to the physical constraint of involving no more goods than can be removed in a minute.

3. *From adaptation to niche construction*. Organisms change the environment they function in, and such changes influence the behavior that emerges, whether through evolution or through learning or both. As termite behavior evolves in such a way that it produces a massive sort of castle for the termites to live and breed in, this changes the conditions of survival, resulting in the evolution of termite behavior adapted to that kind of termite-made environment. We end up with urban termites. Construction of one's own niche reaches dramatic proportions in human societies, having a profound effect on cultural evolution and ultimately, perhaps, on the co-evolution of culture and organism (Laland, Odling-Smee, & Feldman, 1999). Human behavior of nearly every kind is mediated by tools that are themselves human constructions. And these tools are themselves emergents of individual and social behavior. But 'niche construction,' the term used by Laland et al. is somewhat misleading. Neither we nor any other organisms construct our ecological niche. We construct things that have an impact on our environment, but the results ramify in unpredictable ways, as we are gradually coming to appreciate. What settles out as our environment is not the pre-existing environment plus what we have constructed in it. It is a new environment that is an emergent result of unimaginably complex interactions that involve our human constructions in nontrivial ways. Unimaginably complex as the interactions may be, environmental scientists are making some headway in making predictions by modeling self-organization in simplified models.

4. *From physical construction to knowledge construction—the emergence of World 3.* Conceptual artifacts—the theories, defined concepts, histories, designs, problem formulations, and so on that make up Popper's World 3—are themselves emergents. All three of the preceding kinds of self-organization are normally involved. Newton's laws arose from self-organizing processes in his own mind, but they were also a product of a recently emerged international community of scientists whose communication with one another accelerated scientific advancement. Universities provided niches in which theory construction could go on with minimal interference, and towns provided niches within which universities could evolve safely. The emergence of Newton's laws depended on existing

conceptual artifacts—the theories of Galileo, for instance, whose work in turn depended on tools, such as the telescope, which emerged from a quite different set of circumstances. Conceptual artifacts in turn become part of the environment. They can serve as tools for the construction of new conceptual artifacts, and so we get the emergence of progressive scientific disciplines. A whole new plane of activity opens up in which self-organization among ideas produces new ideas. The prodigious impact of this fourth level of self-organization can only be appreciated by comparing societies where it has occurred to those where it either has not occurred (the vanishing stone-age societies) or those where religious strictures or the conditions of survival have held it in check. The other three kinds of self-organization are universal among human beings. The fourth does not seem to be a given of human nature. It has to be recreated in each generation. That—largely unrecognized—is the task of schools in a modern society.

The shift from understanding phenomena in terms of causal chains to understanding them in terms of self-organization is a bigger shift than that from a Ptolmeic to a Copernican view of the solar system but probably no greater than that from a creationist to a Darwinian way of thinking about life forms. There are emotional resistances to all three, because they all involve demoting ourselves. Furthermore, none of them offers a payoff in daily life. Most of our everyday problems are best handled by looking for causal agents, by assuming a stationary earth and a moving sun, and by regarding species as fixed and human beings as a special class. Why, then, are most educated people Copernicans and Darwinians? Because the world makes more sense that way; there are fewer gaps and inconsistencies; it is easier to relate one fact to another. Moreover, because that is the way things are heading. If you have not assimilated the Copernican and Darwinian models, then each new discovery in astronomy and evolutionary biology leaves you farther behind. For specialists, of course, there are more pointed reasons. As for self-organization in human knowledge and cognition, all these reasons apply. Things make more sense. Self-organization is the way the relevant sciences are heading. And for specialists in educational design or knowledge management, the old models are too limited and crude.

Related and Opposing Views

Situated Cognition

Researchers and theorists in this tradition have done a great deal to increase our understanding of the second and third kinds of self-organization described in the preceding section—the emergence of productive social behavior in what are called "communities of practice" and the essential role of tools in such behavior (Engestrom, 1987; Wenger, 1995). The first and fourth kinds of self-organization have been largely ignored, however. Individual cognition has been treated mainly as the internalization of socially constituted knowledge. This has had implications for the social organization of instruction but otherwise has had little to offer for the solution of educational problems. Situativity theorists have also generally been opposed to treating knowledge as detachable from social behavior and from material tools, and this has meant that they have little to offer regarding the creation of knowledge. Although it started out with very wide-ranging ambitions, situated cognition seems to have found its applied niche in management consulting, primarily in helping organizations to foster and take advantage of the implicit knowledge that grows up among workers (Wenger, 1995; Brown & Duguid, 1991).

Situativity theory is very much in the modern mold of scientific treatments of emergence, and thus is potentially fully compatible with the connectionist view of mind and with the idea of cognitive artifacts. But I think proponents of situated cognition hobble themselves by not making a clear distinction between the situated knowledge inherent in the practices of any productive group and the *non*situated knowledge which for some groups is the exportable product of their work and for other groups is the stuff they work with (Bereiter, 1997; compare Engestrom & Cole, 1997). Making this distinction amounts to recognizing World 3. Recognizing World 3 leads naturally to recognizing World 2 as well, the world of personal abilities and dispositions, for it is here that we take account of the abilities and dispositions that people have with regard to objects in World 3. Such an enlarged conceptual framework has room in it for the important insights of situativity theory, but it does not force the concept of situated cogntion to do more work than it can handle.

Brain Science

With new ways of observing the brain at work, cognitive neuroscience is throwing new light on the nature of learning. Research is not only showing what happens during normal learning, as in becoming a skilled reader, but also what happens when learning goes awry, as in dyslexia (Shaywitz, et al., 1998). A sizeable literature has developed out of efforts to derive educational prescriptions from brain research (one educational publisher alone offers 36 titles of "materials to apply new knowledge about the brain and learning to the classroom"). I have not actually read any of this material, but from attending conferences where it has been touted and from Internet browsing I have not noticed any ideas that are less than 30 years old. The people I have talked to who have some sophistication in both neuroscience and education are cautiously optimistic that findings of educational significance will start to emerge in a decade or so, but they are not sure what form they will take (see Bruer, 1997). The first contributions are likely to be to the treatment of pathologies of learning. Beyond that—this is happening now—brain research findings will figure in practical matters mainly through lending weight to ideas that are drawn from other sources and that enjoy other kinds of support as well.

The advances in brain science are, nonetheless, impressive; moreover, there are signs of convergence between the behavioral and brain sciences (as shown, for instance, by examining the journal of that name, comparing what was being published in the 1980s with what appeared in the 1990s). Suppose that there is a revolution in understanding of mental phenomena, comparable to the 19th-century revolution in the understanding of infectious disease. Will it affect education the way the discovery of bacteria affected medicine? And what will it do to the theoretical premises I have been advancing here? On the practical side, there is reason to keep optimism in restraint, as Bruer (1997) suggests. Medicine is essentially unrestricted in the range of treatments it may undertake (which, incidentally, is why it keeps running into ethical problems). Discoveries in biochemistry may be translated into chemical therapies, advances in cellular genetics may be translated into gene therapies, and so on up through anatomy, psychology, and anthropology. Education, along with other knowledge arts, is much more restricted in what its practioners can do. Basically,

teaching is limited to various forms of communication plus a modest ability to modify the environment within which communication takes place. The separation in levels between brain processes and educational practice may be so great that advances in one can have no direct effect on the other. To assume that they should is as unrealistic as expecting that advances in molecular biology should have an effect on the doctor's bedside manner.

But surely scientific advances in understanding learning, memory, thinking, emotion, and the like ought to influence how we *think* about educational issues. No doubt. But, again, the separation in levels argues against high expectations. The question of immediate concern is, what is the likelihood that scientific advances at the level of brain processes will overturn the ideas I have been presenting here? The connectionist view of mind, even in the very general form that I have presented, is certainly vulnerable. But what is the likelihood that further research will show that folk psychology is right after all? I would say it is about as likely as that further advances in paleontology will vindicate creationism. As for the other main idea, conceptual artifacts, that has nothing to do with natural science at all. It will rise or fall on the basis of its usefulness.

Postmodernism

In academic circles, most of the controversy surrounding knowledge does not have to do with its relation to mind but rather with its social status. I refer especially to what is being called the 'science wars,' a controversy that pits defenders of mainstream science against various antifoundationalists, postmodernists, sociologists of knowledge, and feminist theorists who are trying to lower scientific knowledge from its privileged status and put it on a par with other belief systems as found in religions, myths, and political ideologies. Although there are very important issues here, some of which I have addressed elsewhere (Bereiter, 1994; Bereiter, Scardamalia, Cassells, & Hewitt, 1997), they are tangential to theory of mind. Take the issue of how something gets established as a fact. Much of the lighter-weight critique of science is devoted to attacking the notion that facts are established on the basis of correspondence with the way things really are. According to what might be called the liberal scientific view (Rauch, 1993),

facts are facts by virtue of withstanding empirical tests and cohering with other putative facts. According to what seems to be the modal left-wing view, establishing something as a fact is more in the nature of a political victory, having little or nothing to do with an external nature of things (Collins, 1981). But across the whole spectrum of opinions—which range in both directions beyond those I have indicated—it is accepted that there is a something, the so-called fact, whose origins and status are in question. There must be some sense in which adversaries are talking about the same thing, otherwise there could be no argument. At the least, 'fact' has to serve as a 'peg' in the sense discussed in Chapter 3, something to hand assertions on; and that amounts to granting abstract ideas a sort of World 3 status. When that much is granted, opinions may differ wildly and still be consistent with the view of knowledge and mind advanced here.

There is one respect, however, in which the science wars might seem to impinge on the present discussion. I have maintained that one of the inadequacies of folk theory of mind is its inability to deal with advancement of the state of knowledge—with knowledge creation and knowledge improvement, as in the progress of a science. One of the implications of the stronger antifoundationalist positions is that there is no progress in knowledge, only change. Such a position, if you take it seriously, calls into question why anyone would want to pursue science, why governments should be so foolish as to fund research, and so on. But if you are going to go ahead anyway, if only to advance your career, you are going to need tools. One of those tools, I argue, should be a theory of mind suitable to the task of creating new conceptual artifacts. You do not need to be a postmodernist to recognize that some conceptual artifacts are no improvement over their predecessors.

In most segments of society postmodernism is having its main effect on personnel policies, through its campaigns for minority rights, gender equity, and protection from harrassment. But in education it is a movement of more far-reaching consequence. At its worst it means censorship and propaganda; at its best it brings a heightened level of sensitivity to all of the human consequences of power differences and prejudice. In either case,

however, the result is to seal off certain ideas from critical examination—specifically, any ideas closely associated, positively or negatively, with the identity of a disadvantaged group. Objectionable as it may be on several grounds, this is a constraint that teachers from kindergarten to university are learning to work within. And so the contemporary problem is how to carry on effective education within this constraint. Because folk theory conceives of both the personal development of understanding and the general progress of knowledge as the same kind of event taking place within individual minds, it cannot make much sense out of one proceeding while the other is blocked. With the kind of theory I have been advocating, developing understanding of a knowledge object and improving the knowledge object are different though related activities. In any fundamentalist doctrine—be it religious, political, or of the postmodern breed whose names typically start with 'anti'—certain knowledge objects are treated as absolute and unimprovable. If, as it appears, we are in a period in which fundamentalisms of many different kinds are flourishing, we need a theory of mind and knowledge that does not confuse understanding with true belief and that can distinguish advancement of knowledge from individual mental states.

Constructivism

Constructivism has at least three meanings, all of which have some relevance to theories of knowledge and mind. Broadly speaking, there is a World 1 meaning, a World 2 meaning, and a World 3 meaning. Starting with the World 2 meaning, constructivism as it has come to be known mainly through the works of Jean Piaget asserts that knowledge is acquired by a process of mental construction. Although this was a bold proposition during the reign of behaviorism, it has since lost most of its force. As long as you acknowledge that constructive activity can go on unconsciously and automatically as well as deliberately, it is compatible with every cognitive theory going, including both folk psychology and connectionist models. That is why I have not made an issue of it. In educational practice, Piaget's World 2 meaning has degraded into a World 1 idea. Constructivism has become a synonym for 'learning by doing'—in other words, all kinds of hands-on activities and projects. The term has been taken up by reformers who either

want more or less of that kind of thing, and so you will see constructivism vigorously advocated and denounced in current phases of the endless debate about pedagogical methods. As I will elaborate in later chapters, one of the reasons education needs a new theory of knowledge and mind is in order to rise above such fruitless controversy. For that to happen, 'constructivism' in its degraded World 1 sense is one of the terms that will have to go. That leaves the World 3 meaning, which has two parts. One part is what I have tried to capture in the term 'conceptual artifacts.' It is the idea that theories and the like are human constructions much like material artifacts. The other and much more controversial part is that the *truth* of propositions is a social construction. The important thing to realize is that the first part does not imply the second part. You can buy the idea of conceptual artifacts and still believe anything you want to about truth. I think this permissiveness is essential if progress in the knowledge arts is not to get hung up on a philosophical controversy that shows no signs of settlement.

Conclusion

To common sense, mental content is real. It is something we experience. I am arguing that we should reduce it to a metaphor. Conversely, ideas, as common sense regards them, are not quite real once they get outside people's heads; they are not things that exist out in the world along with parakeets and pencil sharpeners; at most, they are public manifestations of things people have in their minds. I am arguing that ideas should be treated as real things out in the world—that there are conceptual artifacts, just as there are material artifacts. They deserve the fullest possible recognition because producing, improving, and applying them is one of the most important things human beings do. Especially in a knowledge society. In fact, work with conceptual artifacts is what makes a knowledge society different from an industrial society.

If you happen to be thinking about parakeets and want to call parakeets mental content, there is no problem with that. The problem is with the millions of items that you are not thinking about right now but that are supposedly stored away in your mind. Such a notion not only becomes implausible under examination, it leads to warped and constricted conceptions of learning, knowing, and thinking. The examples I worked

through in this chapter—of dreaming and of misrecognition—combined with the discussions of everyday understanding in Chapter 4, will I hope go some way toward showing that we do not really need a notion of stored mental content in order to deal sensibly with human knowing.

The alternative to the folk conception of mind as a container of beliefs and other mental objects is the connectionist view of mind as a self-organizing system—a system that does not actually contain mental objects as data but that produces knowledgeable behavior as an emergent. I am pretty sure that this conception of mind is going to win out, that it will continue to develop long after the research program known as connectionism gets assimilated into some more advanced program. It makes better sense of the mind's relation to brain on one hand and to the physical world on the other and of the inseparability of thought and feeling. Two of the most serious limitations of folk theory are its inability to make sense of the pursuit of understanding and of knowledge creation. The connectionist view of mind can make sense of these, but only if it is complemented by a conception of ideas as things to which the mind can become attuned and with which it can interact. That is where the notion of conceptual artifacts comes in. There is self-organization at the neural level that produces thought as an emergent. There is self-organization at the level of intellectual artifacts that produces things like sciences and disciplines as emergents. Obviously the two are related, but how? That is a question that is going to require the collaborative efforts of brain scientists, cognitive scientists, and social scientists to answer. In the meantime, however, those of us who are concerned with practical matters like education and knowledge management—with the knowledge arts, to put it more broadly—do not have to soldier on with only a centuries-old folk theory to guide us. The main components of a new theory of knowledge and mind are available, even though the theory proper is not. These components are the connectionist view of mind and the Popperian view of conceptual artifacts. In succeeding chapters I hope to show that these are serviceable even in their pretheoretical state, that they can help us toward a more enlightened approach to education.