Chapter 5

Aspects of Knowledgability

> When I say what I know, how is what I say *what* I know?
>
> —Wittgenstein, 1980, par. 88

How many kinds of knowledge are there? When it refers to knowledge objects in World 3, the question obviously has no definite answer. It is like asking how many kinds of books there are in a catalogued library. The answer depends on the cataloguing system in use. When the question refers to World 2, however, to knowledge as a property of human minds, the possibility arises that nature may have imposed its own cataloguing system, that there may be natural categories based on ways that the brain holds information. In cognitive psychology a number of distinctions have been made on evidence, for instance, that brain injuries may affect one kind of memory while leaving another intact. Such evidence has been used to support a distinction between 'declarative' and 'procedural' knowledge (Anderson, 1983) and between 'explicit' and 'implicit' memory (Schacter, 1989) and it played a part in Howard Gardner's (1983) identification of seven different 'intelligences.'

For a practical theory of mind, however, brain-related considerations are only indirectly relevant. The issue is what kinds of knowledge it is useful to distinguish. Useful distinctions would be ones that affect how we teach or how we work with knowledge. Tests to qualify for a driver's license typically contain two parts, a 'knowing that' part and a 'knowing how' part. The first might include knowing when it is legal to make a U-turn; the second might include being able to execute a U-turn. Very likely different brain systems are involved, but you should study differently whether this is the case or not. The reason has more to do with how the Ministry of Transport works than with how your brain works. To pass the first test, you must be prepared to deal with the unpredictable wording and focus of test questions. If you simply memorize the words of the driver's manual without understanding them, you are liable to be tripped up. Therefore, you will do well to work at *understanding* the U-turn regulations well enough that you can think through whatever question is posed. Not everyone knows how to do this, with the result that many people fail the written part of their driver's test,

despite having worked hard at memorizing the rules. To pass the second test, you must be prepared to deal with unpredictable features of the situation in which the examiner asks you to execute a U-turn. Accordingly, you should not do all your practicing in one spot, but should practice U-turns in various settings. In the course of such practice, however, you should acquire knowledge that goes beyond mere skill in getting the car to reverse directions. You should acquire knowledge that links up with the first kind. The rule book likely says that U-turns are permissible in certain situations 'provided on-coming traffic permits,' or words to that effect. There is only so far you can get by analyzing such statements. Only through practice on real streets will you come to know whether an oncoming car is distant enough and proceeding slowly enough that you can execute a U-turn and be on your way before it gets to you. Eventually, rule-book knowledge and skill should fuse in such a way that the instant the thought of making a U-turn enters your head you see the situation as one in which it is or is not okay to do so. In educational terms, therefore, we need to be able to distinguish two kinds of knowledge at one stage of learning but also to hold as an objective that these two kinds of knowledge will come together and form a third.

Practically speaking, then, how many kinds of knowledge are worth distinguishing? Under the influence of cognitive science, the currently favored number is two. These are our old friends, knowing-that and knowing-how, better known in cognitive science circles as declarative and procedural knowledge. These two types can be stretched to cover pretty much everything. In the U-turn example, what I have referred to as a third kind of knowledge can instead be treated as a more advanced level of procedural knowledge in which previous declarative knowledge (about when it is permissible to make a U-turn) becomes 'proceduralized' and incorporated into the skill (Anderson, 1987). The possibility of covering all knowledge by two types does not, however, mean that we should do so. We can have as many kinds of knowledge as we like. By an argument analogous to the one that explains why Eskimos need to distinguish many kinds of snow, it can be maintained that educators and others who work extensively with

knowledge need to distinguish many kinds of knowledge.[1] We are still talking about knowledge in the World 2 sense, however—as a property of people's minds. The desirable set of categories will not necessarily resemble the set we would use for cataloguing books or for designing an information system. Rather, the categories ought to divide things up in ways that help us deal with problems of learning and intelligent action.

The Declarative-Procedural Dichotomy

The declarative-procedural dichotomy enjoys a natural fit with conventional educational thought, as embodied in curriculum guidelines, scope-and-sequence charts and the like. In these, the dichotomy is more likely to be expressed as one of content and process, and process is likely to be described in terms of activities to be carried out rather than things to be learned, but underlying it all is still the Bloom's *Taxonomy* view criticized in the preceding chapter. 'Content' is the itemizable content of memory, roughly equivalent to textbook content reproduced in the mind of the student. 'Process' is all the cognitive skills that the curriculum activities are intended to develop and that are supposed to enable the student to do something with the content. I have already criticized this view for its inability to make sense of problems of understanding. In a later chapter I will discuss its limitations when it comes to identifying subject matter worth learning. In this chapter, however, my concern is with the dichotomy itself and what it leaves out.

There is such a thing as declarative knowledge, and 'declarative' is a good name for it. It is knowledge that you can state or declare, thus exposing it to correction by verbal means. According to the view of knowledge and mind that I have been advancing, declarative knowledge of some subject just is the ability to state its content in your own words, diagrams, or whatever. It is not some set of propositions in your mind that is the source of your utterances—but that issue has been belabored enough already. There is also procedural knowledge, which conforms to the ordinary sense of the word 'procedure'; that is, a routine that can be represented by a set of rules, steps, or clearly demonstrable actions. Procedural

---

[1] I am aware of reports that Inuktatuk does not in fact contain a multiplicity of snow words; but the reason the myth will not die is that there are such compelling reasons why the Inuit *should* have a rich snow-related lexicon, and it is a *should* type argument that I am making here about the need for us knowledge workers to have a richer lexicon for the stuff we work with.

knowledge is ability to carry out such a routine, with or without an accompanying ability to 'declare' the procedures that make up the routine. In actual use these terms have gotten stretched over much wider territory than I have just indicated, but the stretching has been at a price. The categories become too broad to be useful.

No matter how far we stretch them, the categories of declarative and procedural knowledge still form a dichotomy. There is no gradual blending of one into the other. Consequently, they leave out—or at least make it easy for us to ignore—a vast range of knowledge that does not clearly belong in either category. The following is a suggestive, by no means exhaustive list of kinds of knowledge that are given short shrift in the declarative-procedural dichotomy. Some items we have discussed before, others are thrown in to extend the range of possibilities:

Number sense

Place knowledge—knowing your way about in a city, a harbor, a forest

Connoisseurship—refined tastes in food, wine, literature, music, romance

Tact

Moral sense

Acquired instincts - "a nose for news," "an eye for the main chance," ability to recognize promising leads, problems, or ideas in one's area of expertise (Bereiter & Scardamalia, 1993, Ch. 5)

Human relations skills, very broadly conceived—ability to read body language and other clues to people's unstated thoughts and dispositions, ability to judge character, to make friends, to gain trust, to persuade or to influence feelings

Style—in speech, dress, art, management, etc.

Background knowledge—what you need, for instance, in order to understand a Jane Austen novel in the context of the time and place and social world of its characters and of Austen herself as a part of that world

It may be objected that some of these items are not knowledge, that they are abilities, dispositions, perhaps personality traits. To object that way, however, is to backslide into folk theory of mind, which requires that knowledge be something statable, or at least amenable to demonstration of the "here's how you do it" kind. The items listed all involve learning (however much they might also

depend on native traits). They are things that people at least try to convey to others through explanations, demonstration, and guidance. That it seems a bit strange to think of them as varieties of knowledge should be taken as a signal that our conception of knowledge needs expanding.

Functional Classifications

Distinguishing types of knowledge on practical grounds is not new to education. Indeed, knowledge typologies abound. In a literature review, Alexander, Schallert, and Hare (1991) found 26 different kinds of knowledge being addressed. Of course there is redundancy in such a list, as the reviewers recognized, but since researchers tend to generate new typologies without considering existing ones, there is not much to go on in reducing the redundancy.

Perhaps the most empirically grounded and functional of functional classifications in the educational literature is one worked out by Robert M. Gagné (1977). Gagné's classification is based on evidence that different conditions apply to different kinds of learning. There are different kinds of prerequisites, different appropriate methods of instruction, and different ways that learning is manifested. Repetition, for instance, is important for learning some things, not for others.

Whereas Gagné's criterion for distinguishing a knowledge type was *needs to be treated differently*, the one I shall apply is a bit looser and accordingly more inclusive: *deserves special attention*. The reason a type of knowledge might deserve special attention could be that, like Gagné's types, it needs to be treated differently for optimal learning to occur; or it could be simply that it needs special attention lest it be neglected, lest it get lost among more conspicuous and high-status kinds of knowledge. Accordingly, the types I identify here include ones not found in Gagné's typology.

On 'Having' Knowledge

In the folk conception of mind, *having knowledge* means having the knowledge in your head. In the conception of mind that I have been arguing for, a mind that is knowledgeable but that does not contain knowledge, things are not so simple. Knowledgeability is a property of the cognitive system. To talk about what people know is to talk about and try to describe that aspect of their minds. But what terms are we to use? 'Knowledgeability' itself is an ugly term and the language cannot be expected to get any better if we start trying to tease apart different varieties or aspects of knowledgeability.

Rather than introduce a cadre of awkward neologisms, I have decided to accept the risk that goes with treating knowledge as something that people can have. This means using the same vocabulary to talk about World 2 knowledgeability as is used to talk about World 3 knowledge. The risk, of course, is that this will contribute to conflating the two—the very thing I have been inveighing against and blaming on folk theory of mind. If someone wants to start a movement to establish different vocabularies for talking about World 2 and World 3, I shall be happy to join; but my personal opinion is that we need a few more believers first.

Michael Polanyi (1964) used the term 'personal knowledge' to distinguish what characterizes individual minds from knowledge in the objective or public sense. Still, the term 'personal knowledge' has unfortunate correlations with terms like 'personal computer' (the now largely forgotten words that 'PC' stands for) and 'personal pizza' (a little empty box so labeled happens to clutter my desk at the moment of drafting this paragraph)—in other words, a diminutive version of the larger thing after which it is modeled. Of course, Polanyi meant no such thing by it. 'Personal knowledge' remains the best term I can find to use when it is necessary to draw distinctions and one does not want to drag in Popper's three worlds for the purpose.

There is no harm in talking about 'having' knowledge so long as we do not slide from there into assuming that this means having particular symbolic objects, such as propositions or rules, in our heads. We should think of 'having' knowledge in the same way that we think of 'having' abilities, attitudes, tastes, and traits of character. That is, we 'have' them as attributes or properties. There is, of course, something in our neural organization that corresponds to 'having' these attributes, but we would not expect to probe the brain and find them. We do not expect them to represent or correspond to anything in the outside world, although they commonly relate to the outside world—in fact, characterize the way we relate to the outside world. We do not expect to be able to enumerate them, although we may speak of someone as having an unusual diversity of abilities or tastes. We may note similar traits in two people, but this does not imply that there is some *thing* that their minds hold in common. So let it be with knowledge, when we are treating it as what individual people have. The complication to keep in mind is that, as people's knowledge relates to the outside world,

some of what is out there for it to relate to is also known as knowledge, in a different sense.

With this extended caveat, we are now ready to examine the kinds of knowledge people have that deserve special attention from an educational point of view.

Six Kinds of Personal Knowledge

What follows is a further development of a typology introduced in *Surpassing Ourselves* (Bereiter & Scardamalia, 1993). There we identified five kinds of expert knowledge, two of which were the familiar 'declarative' and 'procedural.' The remaining three were kinds of knowledge that tend to be neglected by that dichotomy. The same general intent persists here, but I have dropped the two standard terms. They have become stretched to cover too much, thus losing the neat meanings that a narrower sense of the two terms conveys. The narrower sense of 'declarative' is now captured in what I call 'statable' knowledge. 'Skills' now take the place of procedural knowledge, so that 'procedural' can be recognized as a particular class of skills characterized by explicit steps. The newcomer to the list is 'episodic knowledge.' The neglect of this important kind of knowledge in the earlier work probably indicates that our conception of knowledge was still too constrained by the prevailing declarative-procedural dichotomy. Perhaps it still is, but progress is being made.

Statable Knowledge

This is knowledge that the knower can actually put into some explicit form—usually sentences, but possibly diagrams, formulas, stories, or enactments—such that it can be conveyed, argued about, compared to alternatives, and evaluated by others. It is part of what cognitive scientists refer to as 'declarative' knowledge. It is the explicit part. It does not include the vast stretches of unarticulated, largely unconscious knowledge that cognitive theorists such as John Anderson (1983) have also put under the 'declarative' umbrella. I will erect some other umbrellas to take care of that sort of thing. Admittedly, it is a very fuzzy line that separates what is statable from what is not, and the line may shift depending on the time of day or on how much effort one puts into formulating an idea. But it is still well worth trying to distinguish statable knowledge because of its unique cultural significance.

Statable knowledge has been the over-riding and often the only concern of epistemologists through the ages. It is also, of course,

what formal education is most obviously about, and certainly what the testing of subject-matter knowledge is about. All overt actions and products can be discussed, evaluated, reflected upon, but statable knowledge is distinctive in that it can be discussed, evaluated, and reflected upon *as knowledge*. It can be contradicted. It can be argued with. By the same token, it can be defended as true, supported by other assertions, organized into larger wholes. If I pick up a violin and play it, my performance can be discussed as music or as demonstration of skill. But if I say, "This is how the violin was played when it was first invented," and then proceed to play, I have made the playing part of an assertion. It may now be treated as knowledge. People can agree or disagree with it. You cannot agree or disagree with a portrait, regarded as a work of art, but you can agree or disagree with it when it is being considered as a representation of how someone really looked.

Statable knowledge is thus the World 2 counterpart of World 3, the world of abstract knowledge objects. It is the personal knowledge that we can objectify and thus bring into the social processes of knowledge building (Nonaka, 1991). However, as we shall see later, its role is broader than that, influencing all the other aspects of knowledge that we shall consider.

Implicit Understanding

Work on expert systems, knowledge engineering, and expertise has led to a heightened appreciation of the role of knowledge that people apparently have and use but cannot state. Unstated, tacit, or implicit knowledge covers a very wide range, however. In fact, it covers most of the aspects of knowledge I am discussing here. Accordingly, the qualifier 'understanding' is meant to separate off one kind of implicit knowledge from several others. One could say that migratory birds have an implicit knowledge of celestial navigation, but few would say that they have an implicit understanding. 'Understanding,' as discussed in the preceding chapter, carries the implication of an intelligent relationship to what is known. Thus, implicit understanding refers to those aspects of our knowledge that characterize intelligent relationships to things or situations in the world.

Implicit understanding need not have a very exalted character, however. An example I like to use is that of predicting whether crockery of various kinds, dropped from various heights on to various surfaces, will shatter, chip, or remain whole. Most people

would be able to make intelligent, though imperfect, predictions. But except in the rare case where they could recall a very similar instance, they would be quite unable to articulate the knowledge on which their predictions were based. It is knowledge *about*, crockery, gravity, and so on. It is not skill. But it is not knowledge you would find in a physics textbook or a book about ceramics. It is knowledge gained from experience and it probably owes little or nothing to formal education.

Daily life is full of occasions when intelligent action depends on predicting what will occur. Indeed, that may be about all that intelligent action amounts to in familiar circumstances: If you can predict results, the appropriate action is obvious. But in almost all cases, whether it is predicting the behavior of another person or predicting the effect of adding a touch of vinegar to a sauce, the knowledge on which the prediction is based is largely unspecifiable. It is just a sort of residue of past experience.

Cognitive theories based on the mind-as-container metaphor nevertheless require that this implicit knowledge exist as propositions or some other kinds of symbolic representations in the mind. Much of the argument of preceding chapters was devoted to making a case against this implausible notion and replacing it with a connectionist view of mind. In saying that we make predictions based on knowledge, I am slipping back into the folk idiom, and it is well to be reminded that this is just a manner of speaking. When the teacup tumbles from the saucer, we do not experience any calling up of our knowledge of gravitation or ceramics or any forming of a prediction, on the basis of which we select an action. We just respond in a manner appropriate to the perceived situation. We make a grab for the cup or put out a foot to break its fall, steel ourselves for the inevitable crash or find ourselves anticipating the stain on the plush carpet rather than the unlikely breaking of the cup—all in a flash, as if by reflex. And yet the reflex-like response is conditioned by our past experience in ways that make it reasonable to think of there being a residue of knowledge that makes the response an ‘intelligent’ one.

Implicit understanding is more like perception than like having propositions in the head (Clancey, 1991). In Wittgenstein’s terms, implicit understanding is not knowing *that* the world is round but seeing the world *as* round—which may not be how we see it much of the time but is how we are likely to visualize it when thinking about a flight from New York to Moscow or about a lunar eclipse. If we do

not see it that way, if we have to recall and interpret information about the earth's shape, we are liable to find the fact that the New York-Moscow flight passes over the North Pole very strange. Experimental evidence suggests that most people, even though they know about acceleration due to gravity, do not see falling bodies as accelerating. From a literally perceptual standpoint this is to be expected. Acceleration is not easy to gauge visually, especially over short distances. But by not seeing falling bodies *as* accelerating, people tend to ignore this factor in their predictions and explanations. For instance, many adults I have questioned can offer no sensible explanation of why one suffers worse injury falling out of a second floor window than out of a first floor window—this despite the fact that on further questioning they demonstrate the knowledge that acceleration due to gravity means that bodies gain speed as they fall.

Episodic Knowledge

Research on human memory has yielded considerable evidence that memory for events is basically different from memory for meaningful content or 'semantic memory,' as it is called (Schacter, 1989). Different kinds of brain damage can affect one and not the other. Thus, it is one thing to remember that moth larvae eat woolens and another to remember the occasion when moths chewed up a whole corner of an oriental rug. Remembering one can help you recall the other, but they are distinct.

From a functional standpoint, which is what I am trying to sustain here, episodic and semantic memory have different roles as knowledge. Semantic memory covers everything discussed under the previous labels of statable knowledge and implicit understanding. Episodic knowledge much more closely fits the filing cabinet model. Remembered episodes can be retrieved and considered in new contexts. The fact that moths destroy woolens is a useful but limited bit of knowledge. The remembered oriental rug episode, however, may yield many suggestions and questions. The ravaged part of the carpet was under a sofa. Do moths prefer dark enclosed places? Did several generations of moth larvae do the work or would they all have come from one generation? Other relevant episodes might come to mind, contributing to a more elaborate understanding of the clothes moth way of life. At another time the episode might be recalled and contribute to practical considerations about the purchase or arrangement of furnishings.

The functions of episodic knowledge have been extensively pursued by Roger Schank and his coworkers (e.g., Schank, Collins, & Hunter, 1986), who have developed AI programs based on stored and indexed episodes. 'Case-based reasoning' is the name for this lively and expanding area of AI work. Schank notes that much of conversation involves recounting, comparing, and drawing inferences from episodes, and that informal teaching does as well (Schank & Cleary, 1995). Much of the value of business consultants resides in their repertoires of cases relevant to their clients' problems. Reasoning based on cases is distinct from reasoning based on principles. It is reasoning by analogy rather than by deductive inference. But effective thinking and problem solving will make flexible use of both.

Schank's work particularly highlights the 'reminding' problem. We cannot search our episodic knowledge systematically. One thing reminds us of another, and most of the time the connections are superficial. Not very often does the falling of an apple remind someone of the planets in their orbits. Episodic knowledge would seem to represent a great intellectual resource that is largely wasted.

It could be questioned whether memory for episodes in itself constitutes knowledge. Many memories seem to have no function. I remember all kinds of scenes from childhood—a bluebird on a fencepost; the crawlspace under the front porch, frequented by cats; a scary ride in the rumble seat of a car. Psychoanalysis might make something of these, but in their unanalyzed state they do not seem to constitute knowledge. They have no role. Is not episodic memory, therefore, just raw material out of which knowledge may at times be constructed? We may also, however, remember many facts to no evident purpose—names of casual acquaintances long gone, the combination of a high school locker, and so on. Should such trivia count as knowledge while miscellaneous episodes do not? The definitional issue evaporates if we remind ourselves that we are not really talking about episodes and facts stored in the mind but about a mind with the ability to recall past experiences and previously encountered facts, coupled with a disposition to do so spontaneously as well as under conscious direction. What is recalled may amount to significant knowledge at some times and not at others, but there can be little doubt that the recall of past experiences is an important part of knowledgeability.

Impressionistic Knowledge

Beyond statable knowledge and beyond our more confidently held implicit understandings lies a realm of feelings and impressions that also influence our actions, that function like knowledge even though we tend not to think of them as such. Something about a situation or a decision doesn't feel right, or some prospect inspires confidence although we cannot say why. "I woke up that morning with a good feeling about American Motors," says the young stockbroker in Walker Percy's *The Movie Goer*. We trust or mistrust others, find actions morally okay or reprehensible, often with nothing specific we can point to as a reason.

All personal knowledge has an emotional aspect (although it may be quite attenuated, as in the case of knowledge of number facts for those free of mathematics phobias). What distinguishes impressionistic knowledge is that the feelings *are* the knowledge. This becomes most evident when our feelings appear to run contrary to reason or the weight of evidence. We have to decide which to trust. If we decide to go with our feelings, it means that in that circumstance we believe them to constitute more reliable knowledge. But feelings and impressions also constitute important knowledge in circumstances where reason and evidence offer no guidance. This is generally the case with creative efforts (Bereiter & Scardamalia, 1993, Ch. 5), almost by definition. If you can reason your way to a result or get there by systematic empirical methods, the result does not count as creative (although, as in the case of mathematical inventions, you may be able to construct a logical path to the result after the fact). To achieve a creative goal you have to make decisions of uncertain result. The reason creativity isn't mere chance is that creative people become very adept, within their particular fields, at making risky choices that turn out to be good ones. They go by feeling, impression, or what in this context is often called intuition. Creativity remains clouded in mystery, however, unless we accept impressionistic knowledge as knowledge that grows and improves with experience like any other.

Impressionistic knowledge might be regarded as just extremely vague implicit understanding. If I were trying to create a rigorous taxonomy of knowledge I would have to treat it that way. But I am distinguishing types of knowledge on the basis of their deserving special attention, and on this basis impressionistic knowledge clearly warrants a category of its own. Because it is practically unmeasurable, it is almost totally ignored in education. Although

growth in impressionistic knowledge may not be very obvious over the span of a school term, it becomes very obvious over long time spans, however—time spans long enough to be the stuff of biography. It may be reckoned as growth in wisdom, acumen, or moral sensibility, or as becoming a good judge of character or a connoisseur. In all cases it manifests itself in judgment, but in judgment based on knowledge acquired slowly, distilled from long experience. Without it, one grows older but not wiser.

Impressionistic knowledge is what we are left with after we have forgotten all the explicit content of a great literary or artistic work. When we speak of connoisseurship or of a person's having "standards" in language, literature, art, musical performance, cinema, professionalism, or personal integrity, we are talking about impressionistic knowledge. Efforts to formulate such standards as rules generally fail, or provide only crude approximations to the underlying impressionistic knowledge without which the rules are useless. Impressionistic knowledge is also the stuff of prejudices, phobias, and crazes, however. Pointing this out is only to recognize that any kind of personal knowledge can be dysfunctional, can lead us to act in ways that may seem intelligent to us at the time but that may be judged quite differently by others or by ourselves from a different vantage point.

Skill

Skill learning is ubiguitous. No matter what you do, if you do it repeatedly you will become more skillful at it. You may, however, become skillful at doing something wrong, such as grinding your teeth or slicing a golf ball or reading a textbook without grasping its content. Skills have both a cognitive and a subcognitive component, and it is worth distinguishing them, even though they are closely intertwined. The cognitive part is the knowing-how. When we say that someone *knows how* to read, to solve quadratic equations, or to swim we mean that they can voluntarily call forth actions sufficient to achieve a certain result. No matter how learnedly I may expound on the art of swimming, I will not be credited with knowing how unless I can provide evidence of having got into the water and actually done something that would count as swimming. I might now be physically incapacitated and no longer able to swim, but if I could marshal witnesses who had seen me swim in years past, I would probably still be credited with having the knowledge. That is the sense in which the knowing-how is cognitive.

The subcognitive part is the inevitable change in any skill that takes place with practice. The performance becomes smoother, more automatic, and more economical of effort. This is learning, and so the result could perhaps be called knowledge, but I cannot think of any purpose in doing so. When you do push-ups daily, the increased endurance that enables you to do more push-ups is not mainly a result of increased muscle mass. It is a result of your nervous system's learning to distribute the work more efficiently to the muscle fibres you already have, so that, in relay team fashion, they can keep going longer. It is, thus, skill but of a definitely subcognitive kind. If, in addition, you learned a technique for controlling your breathing so as to increase your endurance, that would be a different matter. That would be a gain in knowledge.

In so-called cognitive skills, the cognitive and subcognitive components are harder to distinguish, but still evident. If you have to do a lot of adding columns of figures, part of your gain in performance will come from discovering shortcuts and memorizing combinations. These will constitute gains in knowledge. But your performance will also get faster, easier, and more automatic just through repetition. The difference between the two kinds of learning becomes evident in cases where the cognitive part, the knowing-how, is faulty. If your way of adding figures is wrong in the first place, and no knowledge advance takes place, you will just get increasingly proficient at doing it wrong. This can often be observed with school children, if they are in the hands of a teacher who believes that the way to remedy deficiencies in arithmetic is to assign more pages of exercises.

The cognitive and the subcognitive parts of skill learning can cooperate or they can get in each other's way, which is what makes skill learning an interesting challenge. Cooperation occurs when the automaticity gained through practice frees up mental resources that enable you to think about what you are doing while you are doing it and thereby, perhaps, discover ways to improve your procedures (Karmiloff-Smith, 1992). Cognition can interfere with subcognitive performance, however, as suggested by the old centipede aphorism. Interference can take several forms. "Inner tennis," "inner skiing," and the like are based on the premise that knowledge and reasoning have got to somehow be held in abeyance so that the body can do its subcognitive learning (e.g.,Galway, 1974). On the other hand, automaticity renders your actions less accessible to examination and

deliberate change. Thus a well-practiced bad golf swing becomes hard to correct, not only because of ingrained habit but because you literally don't know what you're doing. The same goes for your accent in speaking a foreign language. You can't hear it and the articulatory actions that are responsible for it are controlled at levels of your neuromusculature too far down to be cognitively accessible.

Folk psychology rather arbitrarily separates knowledge, skill, and emotion. There is some sense to this, even at the level of brain anatomy, but functionally there are no sharp boundaries. Statable knowledge trails off into implicit understanding which then trails off into the feeling states I call impressionistic knowledge which in turn trails off into feelings that have no knowledge function ("I'm feeling low"). On another side, statable knowledge and implicit understanding trail off into skill, which trails off into subcognitive bodily adaptations. It is therefore important, when dealing with a psychomotor skill, to be aware that your body may be learning in ways over which you have little control and that may or may not cooperate with the cognitive part of your skill learning.

A striking example of this appeared in a study by Pam'la Ghent (1989) on expertise in piano performance. A concert pianist was working on a novel score, a transcription for piano of Indonesian wayan music, normally played on a variety of drums. The pianist was working through conceptually how he would have to adopt a more percussive style of playing in order produce the desired effect. Then he stopped himself and said that he had better start practicing, explaining that he did not want his mind to get too far ahead of his body, because his body would then start taking shortcuts that would defeat the intention. Here was someone who seems to have understood clearly that skill is a form of knowledge, but that it depends on a body that also learns in its own unknowing way. That understanding of himself as a pianist and as a learner, however, exemplifies the sixth kind of personal knowledge, to which we now turn.

Regulative Knowledge

Originally this section was going to be titled "Self-Regulatory Knowledge," which is what the corresponding section was called in an earlier pass at this topic (Bereiter & Scardamalia, 1993). As such, it would cover rather familiar and noncontroversial ground, much of

which has been studied under the rubric of 'metacognition.'[2] The general idea is that in any realm of activity there is knowledge that pertains to yourself as a factor in that activity. If it is history, it is knowledge of yourself as a student of history, not only of how you function in that role but of how to get yourself to function in it: how to stay awake while sifting through official documents, to take a lowly example. A higher-level example would be knowledge of your own biases and shortcomings and how to take proper account of and deal with these as you go about the study of history.

But there is also regulative knowledge that pertains to collective activity, and here we get into a controversial zone, in which prejudices run rampant. There are what some philosophers refer to as *regulative ideas*. Prominent among these are *truth*, *objectivity*, and *perfection*. Let us throw in *equity* for good measure. In ordinary language these are referred to as ideals. They are goals which no one can reasonably expect to attain, yet they are thought to have value in directing and sustaining desirable kinds of action. Or at least they used to be so judged. Artists and artisans who pursue the perfect tone, the perfect risotto, or the perfect welded joint continue to be honored, but those who pursue truth or objectivity are nowadays frequently condemned as deluded, arrogant, and self-serving. Those who do the condemning are frequently pursuers of the ideal of equity, a goal that is no more attainable or explicable than truth and objectivity, but that does not come under the same kinds of attack.

To pursue in any depth the question of why *truth* and *objectivity* have become so tainted while *perfection* and *equity* remain pure would take us far afield. Instead, I will only deal with one factor, which has to do with distinguishing among types of knowledge. Ordinarily, truth and objectivity are thought of as attributes of statable knowledge. If I am defending a scientific proposition, I am likely to be seen as implying that it is both true and objective, whether I say so or not. If, as is widely believed these days, no empirically based proposition has or ever could have those two attributes, then I must be deluded, arrogant, etc. What I propose instead, however, is that truth and objectivity are not attributes of scientific knowledge, they are *components of the knowledge that regulates the conduct of inquiry*. In short, they are part of regulative knowledge.

---

[2] "Metacognition," according to Swanson (1990, p. 306), "is defined as the knowledge and control one has over one's thinking and learning activities."

As components of regulative knowledge, truth and objectivity are still open to criticism, but of a quite different kind from the popular postmodernist sort. We can examine them within the context of the whole body of regulative knowledge in a discipline and consider whether they are necessary, what their positive and negative effects are, and whether some other regulative ideas might be better. I have argued elsewhere (Bereiter, 1994) that truth and objectivity are not necessary regulative ideas and that scientific enterprises (including, especially the teaching of science) would be on firmer ground if the pursuit of theory improvement were substituted for the pursuit of truth and if progressive discourse were substituted for objectivity. It remains, however, that most of science has been carried out and still is being carried out by people who conceive of science as the pursuit of truth and objectivity, and it would be arrogant in the extreme to deny that this has proved to be a very productive and generative conception.

The issue before us here, however, is not judging one kind of regulative knowledge against another but rather appreciating the role of regulative knowledge in general. I see regulative knowledge as covering a very wide range, from explicit principles that may be debated and codified as codes of ethics or by-laws down to idiosyncratic personal knowledge, such as how you as a social worker deal with your aversion to assertive people. It may include knowledge of all the kinds discussed so far, but it is always knowledge pertaining more directly to the actors than to what is acted upon.

Writing provides a nice illustration of this distinction, because it is an activity in which self-regulative knowledge plays a crucial part, over and above the part played by other kinds of knowledge. Even skilled writers often find writing difficult and stressful. Procrastination is the rule (Wason, 1980). Competence in writing of course involves many kinds of knowledge besides self-regulative. There is the kind of knowledge represented in rhetoric texts, plus a vast amount of implicit knowledge of conventions, principles, and strategies. Impressionistic knowledge plays an enormous role, especially in stylistic decisions. Epidsodic knowledge, in the form of knowledge of relevant examples, undoubtedly also plays a very large role. And then there is all the knowledge that goes to form the content of what is being written. Statable, implicit, episodic, and impressionistic knowledge will figure in varying proportions,

depending on the genre. But the more of these kinds of knowledge you have and the more you try to bring them all to bear in a coordinated way, the greater the mental burden becomes (Bereiter & Scardamalia, 1987). That, we believe, is why experts generally find writing to be harder work than do nonexperts (Scardamalia & Bereiter, 1991). Their greater knowledge enables them to formulate more complex problems. Solving these problems advances their knowledge still further, thus leading to more complex problem formulations, and so on in a feedforward loop that generates better but increasingly effortful writing.

Managing themselves, shepherding their mental, emotional, and physical resources, thus becomes a major concern of writers. It is one of the few aspects of their craft that professional writers seem to show much inclination to talk about (cf. Plimpton, 1992). Most serious writers, it seems, have carefully nurtured habits as to where they write, when they write, with what kinds of instruments and materials, and in what sequence of stages. Deviating from these rituals tends to disrupt the process for them. Of course, one writer's successful self-regulatory strategies may be the opposite of another's. Nevertheless, writers often make bold to recommend their strategies to others. For instance, Peter Wason (1980) argued for the universal desirability of his own novel approach to writing, which involves among other things scrapping everything at a certain point and starting over.

In the last decade, there has been a trend toward assigning regulative knowledge a larger share of responsibility for what were previously counted as abilities or skills. Critical thinking is still being treated in schools as something to be taught and exercised, but as McPeck (1984) has pointed out, people are generally pretty good at criticizing propositions they are opposed to and need no further training in that. It is thinking critically about one's favored ideas that is difficult, and that is largely a matter of getting oneself to do it. I have made a similar argument about transfer of learned intelligent behavior more generally (Bereiter, 1995). It tends to be bound to the situations in which it was learned, and breaking free of such bondage is on one hand a matter of understanding what has been learned at an abstract enough level that it can transfer and on the other hand a matter of assimilating the learning into the way one responds to new situations, which is therefore a matter of regulative knowledge.

Well-Rounded Knowledgeability

Of the six aspects of knowledgeability that I have been discussing, only two—statable knowledge and skill—offer much to the external view, and so it is not surprising that they should have received almost all the attention of educators and others who trade in knowledge. Episodic knowledge plays a lively role in people's private mental lives—people commonly devote considerable time to mulling over past experiences—and it is the stuff of small talk, but it is too idiosyncractic to be judged by a standard. It would be difficult to compare one person's episodic knowledge with another's or to judge whether a person's episodic knowledge in some area was adequate. Consequently, it has tended to be neglected and undervalued. Intuitive understanding, being mostly invisible even to the person who has it, continues to receive little more than a nod, except when it is judged to be erroneous, as in the case of scientific misconceptions. Impressionistic and regulative knowledge are usually not recognized as knowledge at all, or at least not as part of a person's knowledge of a domain.

Competence in any domain will likely involve all six kinds of knowledge. The prominence of the different kinds will of course vary from field to field, but it is hard to imagine a field in which any of them would be absent. When we look at a highly competent person, however, we cannot help but note that (a) the whole of the person's competence appears to be greater than the sum of identifiable parts and (b) the parts themselves are not easily distinguished. Everything seems to blend together. Statable knowledge trails off gracefully into intuitive understanding, significant parts of which could be rendered explicit if the need arose. Intuitive understanding in turn blends imperceptibly into impressionistic knowledge. Similarly, episodic knowledge blends into both implicit understanding and impressionistic knowledge. On occasion a specific episode may be recalled and explicitly applied to a current situation, but more often past experiences influence actions and choices without any recall and analysis. Regulative knowledge becomes integrated into habit and character.

This merging of aspects of knowledge takes place even at levels of competence well below those of the expert, however. If, for instance, you are painting a room, and this is something you do infrequently enough that you do not have well-practiced routines for the task, you may do some recalling of principles or past instances if a problem arises, but for the most part the residue of past wall

painting episodes is experienced simply as confidence that you will know what to do as the need arises and that things are going all right. It would be impossible to mark off a place where this kind of tacit knowledge leaves off and skill in painting begins.

Put somewhat differently, the more fully developed and well-rounded a person's knowledge is, the more artificial seem the distinctions among its components or aspects. This is a universal characteristic of competence and has nothing to do with the particular set of distinctions I have drawn. AI theorists have had to deal with this phenomenon. 'Chunking' (Rosenbloom &Newell, 1987) is one way. Items of information or procedure that are repeatedly used together are consolidated into larger units that are retrieved as a whole. In 'proceduralization' (Anderson, 1987), statable knowledge that is used in solving problems is coverted to procedures that are simply executed. Opposite to this is the process whereby the results of frequently used procedures are stored as factual knowledge and simply recalled. If you frequently have to convert between miles and kilometres, for instance, as in driving between the United States and Canada, you will soon learn equivalents corresponding to typical speed limits (e.g., 30 m.p.h. is approximately 50 km / hr) and no longer have to compute them. In all of these hypothesized processes, according to the mind-as-container metaphor, one kind of object in the mind is converted into another. This still leaves cognition with a kind of 'chunkiness,' which does not accord with the impression of everything blending together. In connectionist approaches to modeling learning, however, the chunkiness disappears. Instead there is a whole system becoming increasingly differentiated in its response to environmental inputs. Parts of the connectionist network may come to be distinguishable. For instance, in a network that is learning to read, certain nodes may come to be consistently associated with certain letters of the alphabet. But it is still the system as a whole that responds to each new event and is modified by it.

This part-whole issue, which at one level seems academic and fussy, has nevertheless been close to the heart of the main educational debate that has raged through most of the past century and will no doubt continue well into the 21st. On one side are advocates of teaching the 'whole' child and 'whole' language, strenuously opposed to any approach that decomposes learning into separate skills or objectives. On the other side are people who

advocate systematic teaching of subject matter and skills, and this typically involves identifying and sequencing elements to be learned. As so often happens in education, what could be treated as a fairly straightforward and practical pedagogical issue comes, without much justification, to be treated as a theoretical and finally as a moral and political issue.

At a theoretical level, the part-whole issue is a false issue. No one, I believe, denies that mature acquired abilities have the seamless, wholistic character that we have already discussed. At the same time, no thoughtful person will deny what Paul Attewell (1990) has aptly stressed, that

> *all* human activity, even the most mundane, is quite complex. Things that everyone does—such as walking, crossing the road, and carrying on a conversation—are amazing accomplishments requiring a complex coordination of perception, movement, and decision, a myriad of choices, and a multitude of skills. (pp. 429-430)

The reason for educators to pay attention to the parts is so as not to neglect something that needs work. This is well understood in sports, where the part-whole issue tends to be handled on a purely pragmatic basis. Some trainers put more emphasis on the whole—on playing whole games, running whole races—while others put more emphasis on special exercises to strengthen particular skills; but they all recognize the importance of both. Only in academic areas do we find the kind of extremism that insists on one to the virtual exclusion of the other.

The teaching of reading provides the clearest example of the need to pay attention to the whole as well as to different kinds of knowledge that must come together for successful action to occur. On one hand, reading dramatically exhibits the wholistic character of mature competence. For the skilled reader engaged in light reading, the meaning seems to come right off the page, with no intervening steps or sub-processes. Partly because of this, however—because introspection does not yield much sense of what actually goes into the reading process—reading education over the years has tended to neglect or undervalue every one of the six kinds of knowledge previously discussed, although usually not all of them at the same time. Thus, one of the effects of research and scholarship in reading has been to call attention to various facets of reading competence, which seem obvious once they are pointed out, but

which nevertheless become the objects of reform or resistance. The following are noteworthy ones of recent decades, classified here according to the type of knowledge brought into attention:

• *Implicit understanding in reading.*. Cognitive research of the 1970's, particularly that carried out by Richard Anderson and others at the Center for the Study of Reading, brought to light the vast amount of informal knowledge required to understand the simplest texts. The same was being discovered by AI researchers trying to build text understanding systems. *Becoming a Nation of Readers* (R. C. Anderson, et al., 1985), a monograph that featured the role of background knowledge, became the basis for a wave of reform in curriculum and textbook adoption guidelines.

Another kind of implicit understanding more specific to reading has been brought out by more recent research. This is 'phonemic awareness.' Although variously treated as theory or as skill, it most clearly fits the category of implicit understanding as I have sketched it: It is understanding that manifests itself in a perceptual way—not merely knowing *that* spoken words are composed of identifiable sounds but hearing them *as* composites of such sounds (known as phonemes). As with other kinds of perception-like understanding, this does not mean hearing words that way all the time, but hearing them that way when appropriate. It is appropriate, for instance, in producing rhyme and alliteration. It is also appropriate, and by some accounts essential, in learning to read an alphabetic language.

Infants' babbling and much of the spontaneous vocal play of toddlers involves phonemes and their recombination. So do many of the songs and stories that appeal to young children, the Dr. Seuss books being outstanding examples. As a result, many children start school well equipped with phonemic awareness, and as they begin to attend to the spellings of words, the connections between letters and sounds become obvious. But there are other children who show hardly any such awareness, and the prognosis for their success in reading is poor. However, experiments indicate that phonemic awareness can be taught and with beneficial effects on beginning reading (Murray, 1998).

• *Statable knowledge in reading.* Statable knowledge as an outcome of reading is, of course, well recognized. It is what 'book learning' is all about. But its role in the process of reading has been much less appreciated, so much so that a book arguing for its importance was one of the more maligned books of the 1980s in educational circles. In *Cultural Literacy,* E. D. Hirsch, Jr. (1987), pointed out the extent to which even popular media presuppose extensive, though low-level knowledge of historical, scientific, and literary facts. The reader of an American news magazine, for instance, will need some familiarity with names like Eisenhower, Castro, and Edison, and with concepts such as the speed of light, natural selection, and computer virus. Such terms will not be explained and yet are essential to understanding texts. His argument for the importance of such knowledge, although set out in plain language and with vivid examples, proved sufficiently novel that to this day most educational commentators seem to miss his point.

• *Impressionistic knowledge in reading.* In "What Lesson Does This Poem Teach You?" Louise Rosenblatt (1980) criticized school people for treating literature as a cryptic form of expository writing—that is, as if its purpose was to convey statable knowledge. Rosenblatt argued for focusing literature teaching instead on what matters most to people who actually know and understand literature, the personal experience that comes when one is absorbed in a good poem or story. Rosenblatt's writings have made "reader response" a new slogan in school literature curricula. Now even the most pedestrian basal reading program will include among the questions accompanying the reading selections a few that ask children to report how the selection made them feel or how they liked it or what personal experience it reminded them of. Such trivialization of literary experience, which in the worst case eliminates concern with actually understanding the text, comes from placing reader response in contrast to knowledge. What ought to be the concern is treating literature in the classroom in ways that advance impressionistic knowledge of the kinds implied by the broad meaning of the term 'literate.'

• *Episodic knowledge in reading.* Reading allows the accumulation of episodes, both factual and fictional, that can

serve as episodic knowledge. A classically educated politician, contemplating a return to private life, might recall the return of Odysseus as well as the experiences of recently retired colleagues. On a less literate plane, discussions of sexual harassment in the workplace at the time of this writing are more likely to draw on examples from a series of sensational movies on this theme than on real-life cases. They are widely known, the facts are clear, and they are safer to talk about. Roger Schank and colleagues at the Institute for the Learning Sciences have built instructional programs around real and fictitious narratives. Their value at the school level has yet to be exploited, however. A strong 'don't look back' policy guides the reading curriculum like all others, so that, except for testing memory, there is seldom any recall of old stories for new purposes.

The importance of narrative as a form of knowledge has been emphasized in a number of critical examinations of school practice (Bettelheim & Zelan, 1982; Bruner, 1986; Egan, 1989). Remembered narratives are episodic knowledge. Yet, even though the basal readers used in elementary schools consist largely of stories, there is little regard for their potential as knowledge, except for the moral lessons that may occasionally be extracted from them.[3] Although watchdog groups pay close attention to the impressionistic knowledge that a story might convey, their influence has been deadening as far as episodic knowledge is concerned. In order not to give offense, stories are rendered trite and unmemorable, mythic lovability being perhaps the most common theme. Narratives are also much used as attention-grabbers in expository text material, but Hidi, Baird, and Hildyard (1982) found them to deal with trivial or irrelevant points, and thus to draw attention away from main ideas.

• *Skills in reading.* In some educational circles it is impolitic to refer to reading as a skill and even worse to refer to it as a composite of skills. To these whole-language advocates, 'skill' is inseparably linked to 'drill' and therefore is antithetical to

---

[3] One exception, I make bold to point out, is the reading series I have helped author, *SRA/Open Court Reading*, which groups together selections that relate to a common 'explorable concept,' and engages students in work aimed at advancing their understanding of the concept, using story content along with other knowledge sources.

their view of reading as a naturally developing sense-making process. Many of them reject phonemic awareness, mistaking it for a skill. Yet there is mounting evidence that word recognition is a skill quite separate from the kinds of knowledge involved in comprehension (Stanovich & Cunningham, 1991). An indication of this separability comes from the fact that for less than a hundred dollars you can buy text-to-speech software that does a better job of oral reading than the bottom quartile of American high school students. Of course, it does not understand the text, but poor readers listening to it can do so, which dramatizes just how separable the decoding part of reading ability is. It is probably more separable than the skill of catching fly balls is from the other skills of a baseball outfielder. This is because in a baseball game the outfielder has to pay attention to base runners and several other things at the same time as trying to catch the ball, whereas for the skilled reader word recognition is so swift that it is completed before other processes arise to compete with it (Perfetti & Roth, 1981). Yet baseball training normally includes practice in chasing fly balls, with little worry that this may destroy the wholistic character of fielding ability.

A reason that skills have aroused so much opposition among reading specialists is, however, that they have often been dissociated from reading altogether. It is as if baseball players, instead of practicing catching flyballs, sat doing workbook exercises on ballistics. The most abused skill, and also the most controversial, has been phonics. Instead of being taught as a way to identify unrecognized words, it has been taught as a subject in its own right (largely through workbook exercises) and never applied to its intended task. Skill-oriented reading programs have also been full of exercises of so-called skills that are merely the inventions of workbook and test authors: ‘classification’ skills, ‘seriation’ skills (supposed to help children to get story events in the right order, as if that were a problem), and ‘inference’ skills. As is well recognized in more successful areas of skill development, such as sports and music, there is a place for specific skill training, but it needs to be carefully designed so as to strengthen what actually needs strengthening and so as to insure that what is learned is

incorporated into the global competence that is the ultimate objective of the training.

• *Regulative knowledge in reading.* Perhaps the strongest impact of cognitive research on reading instruction has been in the identifying and teaching of strategies for self-regulation of reading processes. Under the rubric of 'study skills,' there had already developed a body of practices aimed at improving learning from texts—summarizing, raising questions, reviewing, etc. However, these were generally things to do after or between cycles of actual reading. Through use of thinking-aloud procedures, cognitive researchers were able to gain insight into what skilled readers do 'on-line,' during the actual reading of texts (Bereiter & Bird, 1985; Scardamalia & Bereiter, 1984).

Changing the Unschooled Mind

In *The Unschooled Mind,* Howard Gardner (1991) has described the vast body of implicit understanding of the world that children have already developed by the time they first encounter formal education. Much of what we know or infer about children's implicit understanding comes from research on misconceptions, research that reveals the many ways in which naive understanding is inconsistent with expert knowledge. Here we encounter a nice irony, however. Students' naive understandings of physics, biology, economics, and so on are interpreted by educators according to their own naive understanding of knowledge and mind.

To folk epistemology, ideas in the head and ideas in the textbook are the same kind of thing, except that ideas in the head may be unconscious. What are controversially called misconceptions, then, are simply cases of a mismatch between the ideas in students' heads and the ideas endorsed by experts. Resolving the mismatch ought to be a straightforward instructional problem. When it turns out not to be, when misconceptions are found to persist even after university-level instruction, this is taken as evidence that naive beliefs are very deeply held.

There are instances in which a new idea encounters a deeply held belief and the result is stubborn and sometimes violent opposition. This can happen, for instance, when science teaching contradicts religious beliefs. It can also happen when science teaching flies in the face of what appear to be indisputable facts. A recent book, *The Nurturance Assumption* (Harris, 1998), marshals evidence against

the belief that parenting makes much difference in the intellectual and personality development of children. The common response of friends of mine, including psychologists who might be expected to show an interest in the evidence, is instant denial, followed by declarations about how important they have been in making their own children the exemplary human beings that they are. This is a classic reaction to heresy. But students' reactions to new knowledge are not usually like that at all. More commonly, students fail to notice any inconsistency between their prior knowledge and what is being taught.

Treating the problem as one of belief change makes it unsolvable. Making it solvable requires abandoning the folk way of thinking about it and seeing it in terms of different kinds of knowledgeability related to conceptual artifacts. The problem is that you can acquire considerable knowledge of a conceptual artifact without undergoing any change in your implicit understanding of the world. This is so obviously true that it almost goes without saying. And yet if you don't say it, you are liable to fall into misconstruing the problem as one of new beliefs colliding or failing to collide with old ones. There is no collision because the phenomena in question occupy different worlds. Here is an object in World 3: Newton's first law of motion, which states that a body in motion will continue in motion unless acted upon by a force. Here is our World 2 implicit understanding, which makes sense of a world in which everything is in the process of running out of steam. Where is the conflict? Conflict arises if we start to work out the practical implications of Newton's law and test them against what our intuitive understanding leads us to expect or if we try to create a theory of our own—to produce a competing conceptual artifact that is consistent with our intuitive understanding. Modern science educators try to induce such conflicts (cf. Hunt & Minstrell, 1994), and that is laudable. The trouble is that those conflicts can be worked through and resolved and yet implicit understanding is *still* largely untouched. We have resolved some theoretical issues. Our statable knowledge has changed. But we still see the world through preNewtonian eyes.

Let me amplify this point by using an example that may apply to you, the reader. The World 3 object in question is the idea of 'knowing your way around' in a knowledge domain. David Perkins (1995) has elaborated this idea into what he calls 'realm theory.' It is one of the central ideas in his revisionist treatment of intelligence.

James Greeno (1991) has used this idea to define the elusive concept of number sense. I made extensive use of it in the preceding chapter in developing a conception of understanding. It is not a difficult idea. Neither is it controversial. It does not provoke opposition or rejection. So if you have read one or another presentation of the idea I will assume that you understand it, according to conventional critera of understanding. Let us assume, furthermore, that you accept the idea as reasonable and worthwhile. I would nevertheless conjecture that this idea has had no discernible effect on how you normally think about intelligence, understanding, or number sense. I base this speculation on the scarcity of mention of the idea in World Wide Web documents, despite the fact that Perkins and Greeno are widely cited for other ideas and the book in which Perkins presents the idea is itself widely cited.

Why does the 'knowing your way around' idea gain so little attention? I suggest that it is precisely because it does *not* contradict an existing belief. It offers nothing to get upset about, nothing to marshal arguments for or against. Instead it offers a different way of thinking about knowledge and mind. Although that is vastly more important than merely offering a different opinion about something, people do not have ready-made ways to deal with it. People know how to argue, criticize, and persuade. That is to say, people know how to deal with statable knowledge. If schooling does nothing else it sharpens those skills. Bringing about changes in statable knowledge may in turn bring about changes in implicit understanding, but these changes are liable to be partial and unsatisfactory (Vosinadou & Brewer, 1987). Instead of regarding this as proof of the strength of naive beliefs, however, we ought instead to regard it as evidence of the limited overlap between statable and implicit knowledge.

One of the puzzles about misconceptions is why they remained undiscovered for so long. In an essay that might be credited with launching educational inquiry into scientific misconceptions, David Hawkins (1978) remarked, "The partial recognition of these problems is very old, probably as old as formal instruction, but somehow they have not been brought into sharp focus." Evidently the kinds of examinations, recitations, and laboratory exercises used in science education enable students to acquire and display statable knowledge without ever revealing the substrate of implicit understanding where misconceptions reside. This does not mean

students only learn by rote—Hawkins recognized that. Students may be able to display some competence at all levels of Bloom's taxonomy, being able to paraphrase what has been taught, apply it to textbook problems, and in some degree evaluate, analyze, and synthesize the material. So what is lacking? Folk epistemology offers no further possibilities, and thus the educational problem not only cannot be solved, it cannot even be formulated.

The problem is how to bring about intended changes in implicit understanding. We know reasonably well how to share or promote all the other kinds of personal knowledge. We share statable knowledge through explanation and argument, episodic knowledge by telling stories, skills by demonstration and coaching, regulatory knowledge by precept, example, and various subtle and not-so-subtle means of inducing conformity to norms. Impressionistic knowledge may not fare well in schools, but in daily life we know that if we want others to feel the way we do about something (a work of art, for instance) it will not do much good to explain or argue; instead we rely on expressive language and gesture and on directing others' attention to significant features. But how do we change the lens through which people view some aspect of the world? For that is what changing intuitive understanding amounts to, and that is why remedying basic misconceptions is no easy matter. It is getting students to look at motion through a Newtonian rather than an Aristotelian lens, to look at the varieties of life forms through a Darwinian rather than a Lamarkian or a Book of Genesis lens.

One thing that is obvious about intuitive understanding is that it evolves gradually through experience. If only that much were taken into account, the battle with misconceptions would take a turn in education's favor. David Hawkins reminds us that conceptual changes that took centuries for scientists to work their way through cannot be expected to happen with students in a couple of lessons. But it is not simply a matter of repetition, of convincing students and getting them adjusted to a matter of fact.

> The round Earth-body was not simply a new fact to be stored along with other facts; it was a fact which required a radical reorganization of the whole category structure of geographical and cosmological thinking. If it were taught merely as a fact, without appreciation of the need to help it penetrate into the subsoil of understanding and to rebuild the mind's category

> structures in the process, it would remain something merely bookish and abstract, to be entertained nervously and then forgotten. Perhaps children of today can grow up without this particular conflict of understanding, one which many of us can remember from our own childhoods. The educational time scale here, that of the transition from opaque fact to intuitive widespread grasp, has been at least a couple of millenia. We ought to do better. (Hawkins, 1978)

But how can an idea, a conceptual artifact, "penetrate into the subsoil of understanding"? As I discussed in Chapter 3, conceptual artifacts such as Newton's laws may on one hand be regarded as truth claims—statements about how the world really is—and on the other hand as tools—tools whose principal use is in solving problems of understanding. Both are important, but educators tend to put too exclusive an emphasis on the first. The result of treating conceptual artifacts as truth claims is change in students' statable knowledge. This does not mean merely that they learn to parrot textbook statements. It may mean quite significant changes in belief, but these changes are likely to have only limited effect on implicit understanding. Students in an introductory physics class may do plenty of problem solving using Newton's laws, but these are problems of *how fast, how far, what direction*, and so on. As John Sweller (1988) has shown, solving such formulaic problems has little effect on understanding. Sweller suggests replacing them with problems of *figure out whatever you can.* As I will argue in Chapter 9, it is even more important to replace them with problems of *why*. And these should not be esoteric problems about juggling tennis balls in a plummeting elevator. They should be problems about explaining the phenomena of everyday experience. That is where intuitive understanding holds sway and where new learning must have its effect.

Knowledge at the Sociocultural Level

The work of sociocultural theorists has highlighted two kinds of knowledge that belong neither to Popper's World 3 nor to World 2. That is, they are kinds of knowledge that cannot be treated as objects in their own right but that are not attributes of individual minds either. As discussed in Chapter 3, these are knowledge constituted in the practices of groups and knowledge embodied in tools. The finely coordinated artistry of an improvisational comedy troupe illustrates the former. The knowledge of photography and physics embodied in

an automatic camera illustrates the latter. Both the comedy troupe and the camera are, of course, treatable as objects, but there is knowledge involved that cannot be treated separately from them. 'Situated' is perhaps the most widely used term for referring to such knowledge.

A useful way of separating personal knowledge from knowledge embedded in the practices of a group is to ask what happens when the group dissolves. Consider that paradigm of collective knowledge, the tango, as in "It takes two to tango." The fact is that people take lessons to learn ballroom dances such as the tango for the very reason that it is knowledge that they can take away with them and that will serve them in new situations. But suppose a couple takes ballroom dancing seriously and starts entering competitions. They will begin to develop skill and artistry as a pair. If they break up, each will carry away some of that skill and artistry, but a distinctive part of it will have been lost and will need to be redeveloped and will not be quite the same with different partners. Thus, we can distinguish a large, basic part of tango knowledge, which is an attribute of individuals even though it is learned in combination with partners, and a smaller but very significant part that is an attribute of a couple and that is unavailable to them as individuals.

Skills are the most obvious kind of knowledge that may be situated at a group rather than an individual level. What of the remaining five kinds of personal knowledge? Do they have group-level counterparts as well? This is an interesting question which I will not pretend to explore in depth, offering only a few observations which suggest that answers may vary:

Statable knowledge is fundamentally social, for to call knowledge statable presumes that it is communicable to others. Most of it originates through discourse, and even when we fabricate it by ourselves, we do so in the imaginary presence of peers (Harré, 1984). Still, statable knowledge is an attribute of individuals. After participating in the same seminar, different people will walk away with different statable knowledge, as may be ascertained by questionnaire or examination. It is, as I said before, the World 2 counterpart of World 3 knowledge. This suggests that World 3 knowledge is the group-level counterpart of personal statable knowledge, and this seems to be the common opinion among people who take a sociocultural view. World 3, under this construction,

becomes the social world, the world of public discourse. That is quite different from how Popper conceived of it, and I have argued at length for the pragmatic value of Popper's conception of World 3 as a world of abstract knowledge objects that can be dealt with in their own right, apart from the individual or collective minds that relate to them. Supposing we accept Popper's conception, then is there anything comparable to statable knowledge that groups can be said to know but that individuals do not? The closest thing to it is that vague but often mentioned something called the 'state of knowledge' in a field.

In any knowledge-rich field, the 'state of knowledge' exceeds what any individual knows. Nevertheless, there are periodic efforts, through annual reviews, handbooks, and the like, to encapsulate the state of knowledge in various fields. Usually these are compilations of efforts by different specialists to summarize the state of knowledge in different sectors; but authors of the separate contributions will frequently confess that their little subfields are still too vast for them. Thus the state of knowledge remains elusive. But it is not reducible to the World 3 objects pertinent to the field. It has to do with the status of those objects—what World 3 objects are recognized, which ones are favored, what are perceived to be outstanding problems, where it is perceived that progress has been made, and so on. A good state-of-knowledge review will address such issues. The result will be new World 3 objects, debatable propositions concerning the status of the primary objects, the theories and putative facts comprising the field. What those secondary propositions are about is this collective, transpersonal knowledge—statable knowledge that is a property of the group rather than of particular individuals, although individuals may know about it. When, for instance, feminist theorists criticize 'androcentric' science , they are criticizing this collective knowledge, not the knowledge of particular individuals and usually not any particular World 3 knowledge objects comprising the field, either. Their criticisms imply, rightly I think, that statable knowledge at the group level can have its biases, gaps, and unexamined assumptions, just like statable knowledge at the individual level.

Whether it is useful to think of implicit understanding and impressionistic knowledge at a group level I am not sure. Talk about public opinion or a public mood or attitude often seems to imply something that transcends individual beliefs and attitudes, but this

may be nothing more than convenient journalistic language, like saying that the stockmarket is optimistic or that it has adopted a wait-and-see attitude. Naturally, if a certain expectation or attitude is widely held and expressed it will tend to spread, but that does not mean something exists at the group level beyond what exists among individuals. de Kleer and Brown (1985, p. 13) say of physics that although it is expressed formally as what we would call World 3 knowledge, "the laws are all based on the presupposition of a shared unstated commonsense prephysics knowledge." Everyday discourse as well presupposes vast shared implicit understanding, but does this imply some kind of understanding beyond the individual level? I see no compelling reason to suppose so.

Where the idea of suprapersonal understanding and impressionistic knowledge makes most sense is in the realm of ceremony and ritual. A traditional funeral service embodies a constellation of beliefs and feelings that may not correspond to the personal beliefs and feelings of anyone present, yet people may participate wholeheartedly, without hypocrisy, and would take offense at a service that reflected their actual sentiments. There is a related literary genre called 'occasional' verse. This is verse written for public occasions such as coronations, graduation ceremonies, and commemorations. It is not intended to express the poet's own sentiments, and the poet is of course not expected to be privy to the actual sentiments of the listeners. Rather, the verse is intended to express the beliefs and sentiments most appropriate to the occasion. If you consider how often in daily life you say things or hear things that are not so much personal expressions as expressions 'appropriate to the occasion,' you begin to get a sense of a large expanse of implicit and impressionistic knowledge that is part of the culture but that is not attributable to the individuals within it.

Episodic knowledge at the group level will consist primarily of episodes that members of the group have experienced in common. These episodes take on a value well beyond that of episodes known only to individuals. An episode can be jointly examined, and individuals can offer their distinctive perspectives on it, thus greatly enriching its potential as a source of ideas, cautions, abstract models, and shared metaphors. First-hand episodic knowledge blends gradually into history and eventually into legend and myth, as old-timers leave and people enter the group who were not around when the episode took place. Myths embody episodic knowledge important

to a culture. Knowledge of these objects is personal knowledge, some of it statable, some of it episodic. If there is anything beyond this, it would seem to exist in the relation of myths to the global entities referred to by terms such as ‘ethos’ and ‘national character.’ The value of such notions is controversial, and I suspect that most action or situativity theorists would reject them.

With a bit of a stretch, it seems, we can make out kinds of statable, implicit, impressionistic, and episodic knowledge situated at the group level. The case for doing so is not very compelling, however. Most of what needs to be considered beyond the personal level can be handled by positing World 3 knowledge objects and not positing knowledge attributable to a suprapersonal mind or suprapersonal cognitive processes. The same is not true for the remaining two kinds of knowledge, however. There are group level skills. Recognizing and understanding them is vitally important for the design and training of expert teams, and expert teams are becoming the heroes and geniuses of the knowledge society. There is also regulative knowledge at the group level, however, and this too takes on social importance when what is being regulated is the production of knowledge itself.

An expert team not only develops skills that inhere in the team rather than in the individual members, it also develops ways of managing and regulating itself that are something more than the ways individual members manage and regulate their actions. This is true whether or not there is a designated manager or leader. There are regulative skills that enable the group to adapt to new situations, to maintain performance under stress, and to keep from breaking up when dissensions arise. There are also norms that regulate the collective activity as well as the individual activity of members. When is a job finished? What constitutes a ‘good enough’ result? There are ethics that apply to group as well as individual behavior. Beyond that, there are delicate problems of trade-offs that involve collective knowledge. For instance, how much efficiency or quality may be sacrificed in order to preserve harmony within the team?

These kinds of regulative knowledge are easiest to discern in small, closely interacting groups, but they may be found as well in larger and looser collectives, such as those of professions and disciplines. The official standards of a profession are World 3 objects and they are meant to regulate individual behavior. But there are unwritten standards that are only to be found implicit in the

workings of the discipline or profession as a collective. For instance, in a learned discipline, what constitutes a discovery, a contribution to knowledge? This is something that award committies and reviewers of manuscripts and proprosals are called on continually to judge, on the assumption that their judgments will not be idiosyncratic but will somehow reflect the unwritten (and evolving) standards of the discipline.

Looked at from the outside, the operative standards of a discipline may seem strange and questionable. To make sense of them it must be understood that they do not descend from a disciplinary god but evolve to meet a complex of social requirements. Until very recently my only contact with the world of intelligent tutoring systems had been through reading about some of the more notable ones and being acquainted with the work of some of the leading thinkers and designers. Then I was roped into reviewing proposals for an international conference in this area and for the first time began to get a glimpse of AI in education as a subculture. Thousands of people make their careers in this area; many are academics, and so publishing and presenting papers at conferences are important. But here problems arise. Intelligent tutoring systems take a long time to build, there is usually little research involved, and the result, even if excellent for its purpose, is not very newsworthy. So how does one fill up a conference program in ways that will demonstrate that AI in education is a discipline and not just a trade? How is a team of developers to produce enough different papers that they can all get on the program and thereby enhance their resumés and perhaps get their travel expenses paid? I doubt if these questions were ever addressed directly by any body, and yet solutions seem to have evolved. From what I have been able to make out (and this is speculation introduced here to illustrate a general point; I would not want to have to defend its validity), a whole substrate of academic writing has evolved in which the game is to translate commonplace ideas into elegant flow diagrams or logical or mathematical formalisms. Thus, an art form has developed that has a superficial relationship to the actual work of the discipline, but that is regulated according to esthetic norms rather than norms of usefulness and originality of content. The attendant regulative knowledge exists at the group level—it is constituted in the practices of program

committees and review panels—and only makes sense when the activity of the discipline is considered in its fullest human sense.[4]

The sociology of science has recently taken an iconoclastic turn in which the regulative norms of scientific communities are seen as not much different from those of most businesses. Truth and objectivity, rather than being regulative ideas that guide research, are a sort of merchandize that scientists package and sell (Latour, 1987). The regulative ideas are those of highly competitive enterprises, armies, sports teams, and the like. Winning and losing are central. This is not the place to evaluate these claims, but it is the place to remark that regulative knowledge at both the group and the individual level is complex and does not necessarily have a great deal of internal consistency. Central to the practice of all sciences and disciplines are what Dreyfus and Rabinow(1983) call "serious speech acts." These are assertions meant to have significance and validity beyond the immediate situation in which they are made. Failing that, you may have a social club but you do not have a discipline. Accordingly, any discipline develops knowledge concerning what constitutes a serious speech act and how such speech acts are to be judged and responded to. Truth and objectivity are likely to figure prominently as regulative ideas, but if not, there will be other regulative ideas such as moral or political rightness, phenomenological richness, or novelty of insight.

Group level regulative knowledge is an attribute of long-term social organization. Engestrom (1987) refers to this aspect of social activity as 'rules,' and it can often be described in that way. But we should exercise the same caution in attributing unwritten rules to groups as we do to individual minds: Rules are creations of the observer; the actual participants may or may not have rules in mind. In an elementary school classroom, one often sees explicit rules of conduct posted on the walls. In a well-managed classroom observable conduct will show a fair degree of correspondence to these rules, usually as a result of conscious effort by teacher and

---

[4] A somewhat similar story has been told in the domain of cattle breeding. Selecting breeding stock to maximize milk production is straightforward on the female side but problematic on the male side. Nowadays geneological records make it possible to select bulls from family lines in which the males have good records of siring milk producers. But before genetics became advanced enough to support such strategies, breeders looked for promising physical characteristics in the bulls themselves. Hence developed the esthetics of cattle judging and the breeding of show animals, which became an art form increasingly remote from its original purpose.

students. At a deeper level, however, there are norms not represented in posters that are probably not even considered by the teacher and that could not be translated into rules if they were.[5] Yet these are norms at the heart of the educational enterprise. For instance, what constitutes an acceptable level of text comprehension? The norm is obviously different in an elementary school classroom from what it is in a law office, but in each case it is unspecified, though real. There is no limit to how hard one might work to tease out every implication of a text. What constitutes good enough is, accordingly, what will result in favorable response within the particular group.

A philosopher friend of mine used to teach in such a way that it took a whole term to get through a page and a half. 'Good enough' comprehension in his class was well beyond the level required in ordinary life. By contrast, I have been in elementary school classrooms where 'good enough' comprehension for ten-year-olds meant being able to identify the topics in a text, without necessarily making out anything that was said about them. That level of comprehension would be sufficient to remind a child of personal experiences related to a story and thereby participate in the kinds of discussions that constituted book talk in those classes.

Paul Cobb and his associates (1997) have studied the development of what they call "sociomathematical" norms in primary grade classrooms. The classrooms were ones in which much of the work in mathematics was focused on producing, comparing, and judging alternative ways of solving a problem. In this context, sociomathematical norms determined what would be judged as a *different* way of solving a problem and what would be accepted as justification for a method or solution. Neither of these could be formulated as explicit rules, at least not rules children could understand, and yet over the course of a year children's judgments in these matters become stricter and more in line with those of mathematicians. Much of what develops will, we trust, be personal knowledge that the children will carry away with them. That, after all, is the point of teaching mathematics. But an important part of

---

[5] An interesting perspective on norms of conduct comes from the study of antisocial children, who seem bent on violating them. In *The Aggressive Child* (1957), Redl and Wineman observed that, although these children go out of their way to violate explicit rules, the vast majority of their behaviors conform to implicit social norms. Evidently you have to know a rule in order to disobey it.

what develops will be regulative knowledge at the group level, that will need to be recreated in the next class the children go to.

Ways of Knowing

In feminist scholarship there is talk about "women's ways of knowing," which are thought to be more subjective and contextual and less analytical than men's ways (Belenky, Clinchy, Goldberger, & Tarule, 1986). Much the same distinction has been made between nonEuropean and European ways of knowing. There is also developing a body of scholarship on "narrative knowing," which may be distinguished from "paradigmatic" or propositional knowledge (Bruner, 1986). Kieran Egan (1997) has elaborated a theory of different forms that understanding takes at different stages of cognitive development. In young children it takes a mythic form, with story characters representing stark contrasts of good and evil, strong and weak. Later it takes a romantic form; tales of adventure are the model within which information is most readily understood. Only in young adulthood, and then not for everyone, does it become possible to understand things in terms of autonomous theories, the principal occupants of Popper's World 3. But throughout life, according to Egan, narrative remains the most natural and accessible form of understanding.

A number of issues get entangled here, ranging from 'learning styles,' which is a kind of quack pedagogical medicine, to the nature of truth. The central, but at the same time least controversial issue is qualitative differences in personal knowledge—the claim that the mental states we call knowing are qualitatively or structurally different for different people, varying with gender, ethnicity, and age. From a connectionist point of view, this almost goes without saying. If personal knowledge is a sort of attunement of the brain, then each change in that attunement—each new learning, in other words—is conditioned to some extent by the person's entire past history, which is surely going to reflect age and cultural experience and quite possibly hormonal differences as well. This notion only becomes problematic if you hold to a folk mind-as-container view. For if you imagine knowledge to consist of objects in the mind, then saying that there are basic cultural differences implies that there are radically different *kinds* of objects in different minds. What could these possibly be? Get rid of mental objects and you get rid of that problem.

The real problem—the one that can lead to departmental coups and political donneybrooks—is not with claims that there are different ways of knowing, it is with claims that out of those different ways of knowing come different kinds of *truth*. Logically, the second claim does not follow from the first; in Popperian terms, the first is an assertion about World 2, the second is an assertion about World 3, and you cannot infer from one world to the other. But what if, acccording to your kind of *truth*, the second claim does follow? Where can we go from that impasse? It is an impasse that has given rise to all sorts of epistemological vagaries; in an effort to avoid them, let us consider a real case.

One place where the idea of different ways of knowing has received more than academic attention is in environmental policy. There is a movement whose identifying labels are “indigenous knowledge” and “traditional ecological knowledge” (often refered to by acronyms) and whose political program is “co-management” of environments by governments and people indigenous to the environment. The current state of the movement may be tracked by doing a Web search on those terms. The idea is that when plans are being made for protection of an environment or for sustainable use of its resources, people indigenous to that environment ought not only to have a place at the table but their ecological knowledge ought to weigh in with the knowledge of outside experts. This movement has taken hold strongly in the far north. There have been conferences of native peoples from the entire circumpolar region devoted to what might be called native knowledge rights.[6] Traditional environmental knowledge (TEK) became a hot issue throughout Canada, however, when a government-appointed Environmental Assessment Panel declared that TEK should receive "equal consideration with scientific research in assessing the environmental and socio-economic impacts" of a proposed mining development (Howard & Widdowson, 1996, p. 34). But what would “equal consideration”mean?

Howard and Widdowson, whose article in a policy journal touched off the controversy, evidently saw it as meaning that TEK and scientific knowledge were to have equal claims to truth. They objected that TEK’s truth claims are untestable:

---

[6] In 1996 a seminar was held in Inuvik, NWT, on “Documentation and Application of Indigenous Knowledge.” The 58 participants included people from Alaska, northern Canada, Greenland, and northern Russia.

> there is no mechanism, or will, by which spiritually based knowledge claims can be challenged or verified. In fact, pressure from aboriginal groups and their consultants has made TK a sacred cow for which only uncritical support is appropriate. Traditional knowledge is thus granted a sanctity which could lead to the acceptance of incorrect conclusions.(Howard & Widdowson, 1996, p. 34)

The academic response to Howard and Widdowson was to splatter them with pop epistemology—criticizing "the simplistic view that there is such a thing as objective, value-free science" (Berkes & Henley, 1997, p. 30). Although, in response to several criticisms along these lines, Howard and Widdowson (1997) denied that they held such a view, they dug the epistemological hole deeper for themselves by saying, “"There are not different ways of knowing. There are different beliefs about the same phenomena." (p. 46)

Clearly, the argument among academicians was destined to go nowhere. But what about the native people whose knowledge was being contested? As I browsed the websites devoted to Inuit knowledge exchange I found a refreshing lack of postmodernist and antiracist posturing. Native writers emphasized the compatibility between traditional knowledge, handed down by the elders, and scientific knowledge brought in by outside experts. “The elders are our experts,” one said. Both kinds of knowledge are derived from experience and are modified by facts. Put the two together and you should get knowledge that is superior to both. A distinction between World 2 and World 3 is implied here, and it would be helpful to make it explicit. There are differences at the World 2 level—differences in personal knowledge—that go beyond “different beliefs about the same phenomena.” There are different ways of seeing the world. But when these different ways of seeing the world give rise to explicit truth claims—to statements of fact or principle, to conceptual artifacts, in other words—it should be possible to compare, combine, or choose between them through a reasonable process of negotiation.

One of the marvels of human history has been the success of trade, despite all the barriers of language, custom, and knowledge that should have made it impossible. Part of that success is no doubt due to the fact that we are all members of the same species and are therefore, in ways we seldom appreciate, much more alike than different. But the other part has to do with the nature of trade. I

have yams, you have a chicken. You want yams, I want chicken. There has got to be a way to make a deal. To do that, we will find a way to communicate about yams and chickens. These real objects provide a basis for communication. If, instead, we are meeting to hammer out an environmental policy that matters to both of us, we still need common objects of reference, only now some of these will be abstract rather than concrete. We need conceptual artifacts—concepts, ideas, facts. We may perceive them differently, have different feelings about them, disagree as to their status, but unless we have some reason to believe we are talking about the same things, our discussion has no chance. The academic arguments about TEK have no chance, because the only objects in common among the disputants are words. With the Inuit there is at least the suggestion that conceptual artifacts of substance can be brought forth as objects of discussion and inquiry.

But traditional knowledge, as described by its native advocates, does not consist of theories or knowledge objects of that sort. The examples offered have to do with skill or know-how. For instance, the elders teach the young how to read many signals, in the look of the sky, the movement of the ice, the behavior of wildlife, that give them information vital to their survival. This is surely knowledge that the visiting scientist does not have. But how is it to be put on the table so as to enter into the formulation of environmental policy? That is not explained, and so it leaves room for what I take to be Howard and Widdowson's real concern: That native participants will use their undeniably superior personal knowledge of the environment in order to claim authority for statements that do not in fact derive from that knowledge.

This is what seems to be at issue in all the controversies where alternative ways of knowing enter the picture: Some statement is made that would normally be subject to criticisms based on reason and evidence, except that its validity is claimed to rest on the distinctive experience or world view of a certain group and to be immune to criticism by outsiders. This is not much different from claiming revealed truth. I have no intention of tackling the vast societal problems arising from this kind of situation—problems that have often been settled by bloodshed. It should be enough for present purposes simply to recognize that one can appreciate and respect kinds of personal knowledge profoundly different from one's own while at the same time insisting that any truth claims, theories, or the

like arising from such knowledge are subject to negotiation. Of course, the terms of such negotiation should also be subject to negotiation, and so on in what one hopes will not be an infinite regress.

But what is the point of recognizing different ways of knowing if these are not credited with validity in upholding truth claims? The point is much the same as recognizing the value of biological diversity: it is a great resource and it is what makes adaptation possible. Having a lot of often subtle intellectual diversity present in a group that is jointly tackling a problem or a project probably does more than anything else that can be done at a managerial level to ensure a creative outcome. (There is more about this in Chapter 7.) To that extent the Canadian Environmental Assessment Panel's controversial recommendation was on the right track and it is what the Inuit writers seem to be advocating. But putting it as giving equal weight to indigenous knowledge and scientific research is just the wrong way to think of it. That only sets the stage for polarization around non-negotiable truth claims. Evaluating truth claims (or however you want to put it so as to avoid the term 'truth') is always a problem in argument and problem solving. Much has to be taken on authority, yet authority has to be questioned and kept in check. Not to do so is to suppress diversity, not encourage it. When there are power differences, as between imported experts and ordinary village folk, negotiation is notoriously difficult. I am in no way trying to minimize the difficulties. All I am saying is that assigning prior value to World 3 objects on the basis of their source is not a way to overcome the difficulty. It is likely to breed a contempt that then carries over into contempt for the people associated with those objects and for the resources of personal knowledge they embody.

Narrative versus paradigmatic knowledge presents a somewhat different problem, which in practice, however, often merges with the kind of problem I have been discussing. Narrative gets identified with women's ways of knowing, and paradigmatic with men's ways. In education, there has been a rapidly growing movement in which narrative is identified as the teacher's natural way of knowing and more generally as the way to advance understanding of teaching, in contrast to the conventional paradigmatic way of knowing represented in educational research journals (Carter, 1993; Clandinin & Connelly, 1996). Activity goes on at different levels, from teachers sharing and discussing their personal narratives to

researchers using them as data. Sometimes the discourse surrounding these narratives looks like journalism, sometimes like psychotherapy, sometimes like literary criticism, and sometimes like philosophizing. It also, of course, varies greatly in quality, along almost any dimension you might care to assess. There is also a split, similar to that arising with TEK, between those who see narrative as yielding a fundamentally different kind of 'truth' and those who see narrative and paradigmatic knowledge as comensurate and combinable (Bullough & Baughman, 1998).

To me what is most significant about this narrativity movement in education, however, is its divisiveness. It defines a community of practitioners and closely allied education professors who are creating their own domain of inquiry in isolation from the established educational research community with its roots in the behavioral sciences. It is easy to see why teachers find more of value to themselves in the sharing of narratives than they do in reading scientific research reports, but this is symptomatic of something deeply wrong with the way the practitioner and research cultures have evolved in education. Institutionalizing what is taken to be a different kind of knowledge widens the gap between the two cultures—a gap which, I argue in Chapter 11, must be closed if education is ever to become a progressive profession.

So it is important to decide whether, in a context like teaching, narrative knowledge represents a distinctive kind of 'truth.' Narratives are cultural artifacts. They are things out in the world, not things in people's minds. But they are not all the same kind of cultural artifact. Some are tantamount to theories. They are *conceptual* artifacts and can be treated as such. A story about how the dinosaurs became extinct, for instance, can be weighed against an explicit theory (set out in paradigmatic form) and can be judged to be compatible or incompatible and to do a better or a poorer job of accounting for various facts. But then there are fictional narratives, which are also cultural artifacts but not conceptual ones; they do not assert anything to be the case and are not anything like theories. They can be judged as to how realistic or authentic they are, but such judgment is made within the framework of literary criticism, not of theory.

The narratives that figure in the teacher story movement are closer to the literary than the theoretical end of the spectrum. They are not knowledge objects but they may contribute to people's

personal knowledge in various ways. Mostly autobiographical accounts of teaching, they can provide other teachers with episodic knowledge that expands the store derived from their own experience and that may be useful in making decisions or understanding events. As literature, they can have an effect on readers' impressionistic knowledge. That is, they can influence how readers feel about things, how they judge them morally or esthetically. As reflective writing, they can affect readers' intuitive understanding of their students, their situations, themselves. Also, and not insignificantly, as narrative inquiry becomes institutionalized stories inevitably influence stories. In accordance with the principle that art imitates art, the kinds of stories teachers write are bound to be powerfully influenced by the stories they have read. One of my concerns about the teacher story movement is that a literary genre is emerging that is derived from poor literary models, that leads to stereotypic descriptions and sentiments, and hence to teachers unwittingly falsifying their own experience. The long and the short of it, however, is that the personal knowledge that results from writing, reading, and discussing autobiographical narratives is sure to be highly individual and, by definition, entirely subjective. How its effects combine with other effects, such as those of first-hand experience and the reading of scholarly literature, is unpredictable and probably impossible to evaluate.

A reasonable conclusion, nevertheless, is that teachers writing and reflecting upon personal narratives is a good thing. It may be more valuable to teachers than to most other people because of their isolation from one another behind classroom doors. If it were treated just as a way of enhancing personal knowledge, then attention could focus on how to make it serve this purpose better. A problem such as the distorting effects of genre constraints on personal experience would receive due consideration. But none of this implies that narrative is a different kind of knowledge, that it is somehow in competition with knowledge of other kinds and to be judged by a different standard.

Where narrative competes is at the level of representation. If you can represent your theory in multiple ways—as a narrative, as a set of propositions forming an argument, as a diagram, as a computer simulation—there are advantages in using all the ways. They may reveal different implications, different deficiencies. For purposes of communication, you are likely to find that the story reaches more

people, that the diagram helps some people and confuses others, and that the argument is best for getting critical response. The differences become more profound when we consider knowledge creation. Stories have the great virtue in constructing explanations of human affairs that they allow motivation and causality to be brought together in a coherent way. Argument forms tend to keep these separate and diagrams tend to end up as boxes connected by arrows every which way—the refuge of the conceptually handicapped—with motives in one box, causal factors in another. Because they favor a one-thing-after-another chain of events, however, stories are not a good way to formulate or understand constraint-based explanations, which are the norm in the physical sciences.

In this way narrative surely does represent a *way* of knowing, or perhaps a family of ways. It is a way in which the human mind, with all its evolved dispositions, capacities, and limitations, can come to grips with problems of understanding and create conceptual artifacts to serve as objects and tools of understanding. If some people are better disposed than others to narrative representation, then it is reasonable to say that they have a distinctive way of knowing, different from those who favor some other form of representation. Out of these different ways of knowing are likely to come differences in all six of the kinds of personal knowledge discussed in this chapter. People will see the world differently, will recall different events, will have different feelings and intuitions, will develop different skills and norms of conduct, and are likely to end up with different statable beliefs. All of these differences can work to the benefit of collective problem solving and knowledge advancement, unless—and this is a very big 'unless'—they become the basis for non-negotiable truth claims.

Conclusion

This chapter could have been titled "Ways of Knowing." That is certainly a more graceful expression than "Aspects of Knowledgeability."But as we have seen, 'ways of knowing' is a phrase that has collected about it a large bundle of epistemological confusions and loose strings. It illustrates how wrongly things can go when people fail to distinguish World 2 from World 3. As a result, a perfectly noncontroversial statement—that people's personal knowledge is richly varied in ways that reflect the differences in their lives and backgrounds—gets entangled with the notion that there

are different kinds of truth. The latter, although controversial, is also a defensible statement; but it draws no backing from the first statement. ‘Different kinds of truth’ is a casual way of saying that there are incommensurable theories. “Incommensurability,” according to Ronen (1998), “occurs in science where two theories lack a common measure, a standard reference, or an external criterion that could have served as grounds for comparison.” (Ronen, 1998). It is, thus, a characteristic of conceptual artifacts. It has as much to do with the way theories affect how we perceive the world as it does with how we perceive the world affects our theories. But in no way does incommensurability imply that the justification for a theory, belief, or policy is unique to the personal knowledge and experience of its authors. Even the stoutest advocates of incommensurability allow that incommensurable theories can be separately described and criticized in a common forum (Feyerabend, 1988). There is, in short, no non-negotiable World 3 knowledge, however inaccessible one person’s World 2 might be to others.

Folk psychology and epistemology, which treat knowledge as objects in individual minds and public knowledge as merely the expression of such knowledge, lead to undervaluing both personal and public knowledge. Public knowledge, according to this view, is ultimately just people’s opinions. That being the case, why shouldn’t my opinion count for as much as yours? It is difficult to argue that in general it should, but equally difficult to argue for a way of deciding when one opinion should count for more. But seeing public knowledge as the externalization of personal knowledge diminishes personal knowledge as well. What is inside the head is assumed to resemble in form if not content what is outside the head: in short, a collection of sentences. That could be taken as a rough characterization of what I have called ‘statable’ knowledge. It leaves the other five kinds of knowledge, except for skill, in limbo. They tend either to be ignored or to be romanticized. Either way, they fail to receive the systematic attention that would be required to make them into realistic educational or knowledge management goals.

A workable approach to personal and public knowledge requires recognizing first of all that they are very different. One is not a copy or a manifestation of the other. Popper’s distinction between World 2 and World 3 may be too stark a distinction for some purposes, but I have not found anything that works better for dealing with issues of education and knowledge creation. In order for that epistemology to

work well, however, it has to be coupled with a psychology that gets past the folk notion of 'things in the mind.' A connectionist view of the mind sees learning as the continual tuning of an almost unimaginably vast neural network to the world with which it interacts. It is an embodied mind; in fact, the only reason for talking about minds rather than bodies is in order to address such phenomena as thoughts, beliefs, plans, intuitions, certainty and uncertainty—the phenomena of subjective mental life. Out of this interaction with the world come those kinds of attunements that for practical purposes we may distinguish as intuitive understanding, episodic knowledge (memory for events), impressionistic knowledge (feelings that serve as knowledge), self-regulative knowledge, and skills. Something like these kinds of personal knowledge might also be attributed to other intelligent animals. But the human mind has a capability that is evidently unique. It is the ability to create artifacts—objects of deliberate design, which then become part of the world to which personal knowledge becomes attuned. Some of those artifacts—some of the most important in the modern world—are immaterial. They are conceptual artifacts. People can develop intuitive understanding of an idea or theory, which is not the same as intuitive understanding of the things that the idea or theory is about. They develop feelings about an idea, which can develop into impressionistic knowledge. They develop skills in use of the idea; they recall episodes involving the idea; they develop norms of conduct related to the idea. And they acquire ability to state some of this knowledge, which therefore enables them to teach, criticize, argue about the idea. The result is a dynamic process which, viewed from the outside, we call by terms such as 'scientific progress' and 'the advancement of knowledge.' As for what goes on inside, we lack suitable terms. 'Intellectual growth,' 'education,' 'self-actualization'—these are terms that do not quite make it. In Chapter 7, I will discuss further what is supposed to happen to the individual in the process of interacting with conceptual artifacts, because it is central to defining the purpose of formal education. As a place-holder, we might call what happens 'becoming a fully functioning member of a knowledge society.'