Chapter 4

The Knowing Mind

> [N]atural selection built the brain to survive in the world and only incidentally to undertand it at a depth greater than is necessary to survive.
> —E. O. Wilson, 1998, p. 61

What does it mean to know something, if it does not mean having something like a sentence or a picture in the mind that constitutes that knowledge? It seems to me that the prospects for a new theory of mind turn on this question. If it cannot be answered, or if the answer is some kind of conceptual monstrosity, then we are probably better off sticking with the theory acquired at our mother's knee.

As I remarked in Chapter 2, it is misleading to use the trivial examples of knowledge that so often appear in philosophical treatments—knowing the capital of Malaysia, the chemical composition of water, and so on. Any account of such knowledge that does not posit sentences in the mind will surely appear to be awkward and a lot of work for little return. But education does not rise or fall on its ability to deal with such knowledge, and neither do most other human enterprises. Education does rise or fall on its ability to deal with problems of understanding, so let us turn to those as a fitting challenge for a modern theory of mind.

Understanding: The Ghost in the Taxonomy

> Specifically, what does a student do who 'really understands' which he does not do when he does not understand?

That is the question as put in the introduction to the mightily influential work known among educationists as Bloom's *Taxonomy* (1956).[1] One will look in vain through that work for an answer, however. There was hardly any way the authors could have produced an answer out of the mixture of behaviorist and folk theory that prevailed at that time.

Instead, the *Taxonomy* set out a series of levels of cognitive attainment, with Knowledge (which means, roughly, literal knowledge) as the lowest level, followed by Comprehension,

Application, Analysis, Synthesis, and Evaluation. Putting Comprehension next to the bottom of the hierarchy gives the unfortunate impression that the taxonomists considered understanding to be a rather low level of attainment. But a close reading of the monograph suggests that the authors viewed the whole taxonomy as addressing the issue of understanding.[2] The ordering of the levels was based on logical dependency: In order to comprehend a text, you must be able to process it at a verbatim level; to apply information you must comprehend it; and so on.
Most of the monograph consists of sample items for testing the various cognitive levels. Although they undoubtedly get at worthwhile educational attainments, they leave understanding as the ghost in the taxonomy. Success on each item may require understanding of some kind, but the emphasis on skills obscures any picture that might emerge of the examinee's understanding. For example, one item intended to test application of knowledge in physics asks what would happen to a ball if it were thrown into the air by a passenger in an elevator that was in free fall. This obviously demands a substantial understanding of Newtonian mechanics, but a student might understand the relevant principles perfectly well and yet be unable to put them together to solve such a complicated problem. One could give an array of such test items, at all of the levels, all bearing on the physics of falling bodies, and yet not end up with any clear idea of students' understanding of the topic. That possibility is not just hypothetical. Years of testing, using the commonsense approach, failed to reveal the deep and pervasive misconceptions[3] that students are now being found to hold in practically every area of science (Wandersee, Mintzes,& Novak, 1994).

The strong point of this commonsense approach is the elaboration of a whole range of performances that are relevant to judging understanding. Its weakness is in its reliance on the mind-as-container metaphor of folk theory of mind. Knowledge really figures in the *Taxonomy* only at the Knowledge level, the lowest level in the hierarchy. The higher levels are taken to represent "intellectual abilities and skills" (p. 204). The authors are quite explicit in their adoption of the container metaphor:

> It may be helpful in this case to think of knowledge as something filed or stored in the mind. The task for the individual in each knowledge test situation is to find the

> appropriate signals and cues in the problem which will most effectively bring out whatever knowledge is filed or stored. (p. 29)

Comprehension, Application, Analysis, Synthesis, and Evaluation—the higher levels in the taxonomy—accordingly represent intellectual skills that operate on the contents of the mental filing cabinet (pp. 38-43). This view dominates much of the discourse of educational practice to this day. Such a view not only fails to account for understanding, it does not even make a place for it. Knowledge is reduced to stored reference material; intelligence and understanding are all in the skills that make use of this material.

A conspicuous failing of the *Taxonomy* is in providing no grounding for the notion of *depth* of understanding. Without that, we really have no handle on understanding at all. For knowledge is almost always characterized by understanding in some degree, however slight. Another item from Bloom's *Taxonomy*, this one intended to test skills of 'Analysis,' illustrates this point:

> Galileo investigated the problem of the acceleration of falling bodies by rolling balls down very smooth planes inclined at increasing angles, since he had no means of determining very short intervals of time. From the data obtained he extrapolated for the case of free fall. (Bloom, 1956, p. 151)

Examinees are asked to identify an assumption implicit in Galileo's extrapolation. To do so requires understanding the concepts of acceleration due to gravity, rolling friction, and extrapolation. High school students who have taken physics and algebra may have encountered all three of these concepts, but I would expect most of them to fail the item nevertheless. Indeed, I expect that most high school students would not even be able to understand what was being asked of them. Is this because they lack analytic skill? It is not clear what that means or how it could be inferred from performance on such test items. What is obvious, however, is that this problem requires a *deeper* understanding of acceleration due to gravity, rolling friction, and extrapolation than high school students are likely to have acquired. If you only understand acceleration due to gravity as it applies to free fall, if you do not understand how rolling friction relates to incline, and if you think of extrapolation only as the extension of a straight line, then you have no hope of understanding the logic of Galileo's experiment. The *Taxonomy* , however, offers no help in determining how deeply a concept or principle is understood. An unfortunate consequence is that it has served to encourage spending time on hidden assumption problems and the like (practicing what are called 'higher-order thinking skills') while neglecting deeper understanding of subject matter. This glorification of high-sounding skills and contempt for knowledge is a large part of the reason why I say that the Knowledge Age has not yet come to the schools. I am not sure whether Bloom's *Taxonomy* deserves blame for this or whether its continued prominence merely reflects education's backwardness. At any rate, it is clearly part of the problem rather than part of the solution.

Understanding as an Attribute of Cognitive Structures

With the so-called 'cognitive revolution,' knowledge assumed a much larger and more complex role in human cognition (Bereiter & Scardamalia, 1992). A science educator of the 1950s, encountering a student who knew the textbook facts about gravity and acceleration but who failed problems like the two discussed above, would conclude that the student lacked necessary cognitive skills (cf. Bloom, 1956, p. 126). A cognitively sophisticated science educator of today would take, as a first hypothesis, that the student did not really understand gravity and acceleration well enough to draw the appropriate implications. This hypothesis would be grounded in research indicating that such failures of understanding are common even among college students of physics (Kaiser, Proffitt, & McCloskey, 1985). But what does it mean to understand something well enough to draw the appropriate implications? How is this mental state different from having the facts stored away in memory and relying on general cognitive skills to apply them?

In what I take to be the prevalent cognitive view, knowledge is still thought to consist of things in the mind, but the things are considered to be large and to have complex internal structures. A student's knowledge of gravitation may be in the form of a schema, which may be visualized as a hierarchical outline with some lines filled in and others containing blanks waiting to be filled. My relativity theory schema, for instance, would be mostly empty and would lack even the major headings and the right blanks to constititute what a physicist would credit as understanding. Alternatively, a student's knowledge of a topic may be in the form of a network of concepts and meaningful connections among them. According to either model, the answer to the preceding questions is much the same: The student with a good understanding of gravitation has the relevant concepts embedded in a complex cognitive structure with many connections and implications already filled in.

Understanding, thus, is a characteristic of the schema or network as a whole. To evaluate a student's understanding, we may compare the student's schema or concept net to that of an expert or to some ideal or minimally adequate schema or network. As Nickerson (1985, p. 222) put it,

> One understands a concept (principle, process, or whatever) to the degree that what is in one's head regarding that concept corresponds to what is in the head of an expert in the relevant field.

We may call this the 'correspondence' view of understanding.[4] Such a view, of course, presumes that mind-as-container is not merely a handy metaphor but is how things really are. If the mind is not actually a container of mental objects, if it is only a system of capacities and tendencies as the connectionist view of mind suggests, then Nickerson's idea evaporates. There is nothing on which to base a correspondence. You can, of course, get people to draw network diagrams representing what they understand and you can compare one diagram with another (Novak & Gowin, 1984), but only the most naive believe these are pictures of what is in the head. They are documents, in a class with essays and movies, not 'mind maps' as some enthusiasts say. Like other documents, they may provide useful information, but comparing one with another is difficult and a far cry from comparing one mind with another. Presented with a document—whether it is an essay, a concept map, or something even more exotic—we cannot say that we are looking at its author's understanding. We are looking at words or lines or icons. To make inferences from them about the author's understanding, we need a concept of understanding that will guide us in what to look for. The correspondence view fails to do that.

A Relational View of Understanding

For a start on a more illuminating view of understanding as it relates to educational purposes, we need to go back a decade before Bloom's *Taxonomy*, to a time before understanding had been decomposed into cognitive behaviors and later reconstituted as mental structures. In 1946, the National Society for the Study of Education devoted a yearbook to the measurement of understanding. In a chapter titled "The Nature of Understanding," Brownell and Sims (1946, p. 27) noted that "A technically exact definition of ‘understand’or ‘understanding’ is not easily found or formulated.” Observing, however, that most people seemed to be able to use the terms with “considerable agreement as to meaning,” they proposed to take a practical and nontechnical approach to the topic, beginning with “a broad statement (not a definition) about the nature of ‘understanding.’” The broad statement was the following:

> *As a start, we may say that a pupil understands when he is able to act, feel, or think intelligently with respect to a situation.* (Brownell & Sims, 1946, p. 28; italics in the original)

Thus, understanding was seen, not as some entity in the head but as an aspect of how we deal with the world. That is a valuable intuition, very much in tune with situated cognitive theory as it is developing today. What the statement gains on this account, however, it loses on another. It directs attention not toward understanding itself, as some distinctive form of human competence, but toward the ways that people give evidence of understanding. Carried to the extreme, this can lead to the sort of empiricist reductionism represented in statements like “Intelligence is what intelligence tests measure.” In less extreme forms it leads toward what we have seen in the *Taxonomy* or in present-day performance assessment.

A recent and sophisticated version of this reduction of understanding to performance comes from David Perkins. He calls it a “performance perspective”:

> Our "performance perspective," in brief, says that understanding is a matter of being able to do a variety of thought-demanding things with a topic—like explaining, finding evidence and examples, generalizing, applying, analogizing, and representing the topic in a new way. (Perkins & Blythe, 1994)

When itemized, these “understanding performances” (Perkins, 1992, p. 77) cover the same ground as Bloom’s taxonomy—explaining, applying, etc. Also in keeping with Bloom, these performances are thought to represent not only ways of demonstrating or testing understanding but ways of teaching it as well.

> This performance perspective on understanding illuminates what a pedagogy of understanding should attempt: to enable students to display a variety of relevant understanding performances surrounding the content they are learning.
>
> ....

> If we want students to understand, we should make the choice of teaching them understanding performances about Newton's first law or anything we want them to understand. (Perkins, 1992, p. 78)

Is Perkins saying that understanding just is these performance capabilities or is he saying that understanding is a mental state that makes these performances possible? Perkins is not entirely clear on this point, and it may seem to be one of little consequence. But if understanding is nothing more than a set of desirable performances, why do we bother talking about understanding at all? Why not just focus on the performances? Consider driver education. We could identify a set of performances and habits that constitute good driving and teach them, without giving any thought to understanding—*unless we conceive of understanding as an objective that deserves attention in its own right, either as something beyond the performance objectives or as a prerequisite to mastering them.* To get serious about teaching for understanding, as I see it, means treating understanding as something that warrants attention in its own right, over and above the attention we give to performance. It means giving attention to *what* it is that needs understanding, to the problems of understanding it, and only further down the line to *how* we will assess it. But this implies that understanding is a state that enables performance, that it is not the performance itself. How can we characterize this state, without falling back on the folk / cognitivist notion that it is a characteristic of objects in the mind? The answer, I believe, lies in considering understanding to be neither something in the mind nor a set of overt performances but rather a relation between the knower and an object of understanding.

This is not an easy conception to grasp, because it depends on both of the departures from folk theory previously discussed; it depends on a connectionist view of mind and a Popperian view of conceptual artifacts. Let us start with an analogy. Gravity is like understanding, in that we only know it through its effects. The childish belief that gravity must be a substance may be likened to the folk belief that understanding is something in the mind. With considerable difficulty, science educators try to get across the idea that gravity is a relationship between objects. Suppose we apply the same idea to understanding—*that understanding is a relation between the knower and an object of knowledge.* You are not going to dig into the brain and find it, anymore than you are going to find gravity by digging into the earth. But we can learn a lot about the relationship. And, as is not the case with gravity, we can influence the relationship through our actions.

Adopting this relational perspective, we may restate the question that introduced Bloom's *Taxonomy*:

> What distinguishes the relationship that we call 'understanding' from the many other ways that a person may relate to an object of knowledge?

What, for instance, distinguishes understanding genetic engineering from being good at, having an opinion about, being interested in, knowing some things about, being worried about, maintaining an open mind about, or being in the business of genetic engineering? To relate to something is to have a certain disposition toward it, supported by the necessary abilities. The relation of 'being good at' implies ability to carry out the work of genetic engineering along with a disposition to do so. The relation of 'being worried about' implies a disposition to think dark thoughts about genetic engineering, which in turn implies some ability to infer consequences of it. 'Understanding,' by comparison, is broader and less definite in its implications. *Understanding implies abilities and dispositions with respect to an object of knowledge sufficient to support intelligent behavior.*

The difference between what I am proposing and what Perkins has proposed may seem too subtle a semantic distinction to be worth the extra trouble that my version entails. The difference takes on clearer significance, however, when we think about teaching for understanding. Perkins's version points inevitably toward teaching for the test—defining a set of performances that will be accepted as evidence of understanding and then schooling learners in those performances. According to the view that I am advancing, teaching for understanding is a matter of cultivating the learner's relationship to objects of knowledge, developing it into a relationship capable of supporting intelligent action. Later I will expand on what this entails, but first I need to show how this relational view fits with what we already understand about understanding in everyday life.

Everyday Understanding

Almost all discussion of understanding deals with understanding abstract entities such as concepts, principles, or theories. Because folk theory locates these things in the individual mind, it is natural to locate understanding there as well. But in our day-to-day lives we are often much more concerned with understanding objects and events in the external environment than we are with understanding ideational entities.

High on the list of things we try to understand in our daily lives are other people—not people in general, but the particular people who matter to us, each one of whom presents a distinctive problem of understanding. Although understanding other people is widely recognized as hard to do, we seem to have little difficulty understanding what it means to understand other people. No one is waiting for the experts to provide us a rigorous definition. Let us, therefore, see whether we can draw any generalizations from the understanding of people that may be usefully extended to other kinds of understanding. The following are commonplace observations that would apply to understanding any particular person:

1. What constitutes understanding depends on your relationship to the person. Your understanding ought to be quite different depending on whether you are the person's child, spouse, assistant, client, or psychotherapist. Thus there are many different 'right' ways to understand a person.
2. Understanding is intimately bound up with ability to act intelligently in relation to the person. Each advances the other. Although, conceptually, understanding and intelligent action are not the same, it would be virtually impossible in any actual case to disentangle one from the other.
3. Understanding is also intimately bound up with interest. It is not just that interest motivates you to try to understand the person. Indifference toward another person would imply lack of understanding.[5]
4. Understanding individual people may depend on understanding their relationships and the people, institutions, occupations, and so on, with which they have these relationships. In this sense, understanding does involve a network—but it is a network out in the world, not inside your head.
5. Understanding of a person is not necessarily accompanied by ability to explain. Explanation is an indication of understanding, but, more importantly, it is a *means* by which understanding is developed and shared.

6. Although there is no single correct, complete, or ideal understanding of a person, there are recognizably wrong understandings, and these are potentially correctable.
7. Many absorbing discussions in everyday life consist of two or more people trying to achieve a mutual understanding of somebody else. Although there may be considerable speculation about things in the mind of that third person, the discussants make little reference to objects in their own minds. Instead they talk about the other person or about the conjectures and explanations that have been put forth for discussion.
8. A major way in which understanding of others is manifested, developed, and shared is through narratives in which the other person figures importantly. If the stories you tell about Smith strike others as implausible or incoherent or as failing to ring true, this will count as evidence that you do not understand Smith. Superior understanding will be demonstrated by telling stories that other people who know Smith find revealing and enlightening.
9. Having a *deep* understanding of another person means understanding deep things about the person: underlying motives; basic values, dispositions, strengths, weaknesses; recurrent patterns in thought and behavior; and so on.
10. Deep understanding of a person is most readily demonstrated by insightful solution of problems involving the person and by telling stories that have depth of characterization.
11. Depth of understanding normally implies deep and complex involvement with the other person. “No understanding at a distance” is a sound enough rule that those who claim exemption from it are claiming mystical powers.

Without much adjustment, these same observations may be applied to understanding other things in our environments: devices, such as a computer or an automobile; institutions, such as a university or a bank; games, occupations, and rituals—in short, pretty much the full range of things that we have some practical reason to understand. In all such cases, understanding inheres in the relation between us and the thing to be understood. The nature of our understanding will depend on the nature of that relationship and will be intimitely tied to our goals, interests, and skills. Everything that needs to be said about our understanding can be said without referring to objects in our minds. It will in all cases be sufficient to refer to the object out there or to various conjectures, explanations, and narratives that enter into discussions of the object.[6]

In all these practical cases, deep understanding means understanding deep things about the object in question, which in turn implies deep and extensive involvement with the object. A complex device like a computer can be understood at a number of different levels—ranging roughly from outside to farther and farther into the components of the system. Different levels relate to different kinds of action requiring understanding at that level. Understanding that is limited to the front view of the computer may be sufficient for normal operation, but understanding what is on the back of the computer becomes necessary for installing peripherals; understanding something about the computer's memory and about memory allocations becomes useful for dealing with system crashes; and so on down to levels of understanding required to write system software. Practically everything dealt with in the practical world admits deeper levels of understanding relevant to certain purposes. An eggbeater is much simpler than a computer, but still there is a level of understanding required for designing an eggbeater that is notably deeper than the understanding required to use one.

The term 'deep' is not meant to have a precise meaning in these contexts. It is often useful, especially for purposes of education and training, to think of a progression of levels of understanding that a learner may go through, with one level perhaps being necessary before the next can be achieved. But this does not mean all kinds of understanding can be ordered as to depth. There are things to be understood about a stone that are relevant to the work of a stonecutter and others that are relevant to mineral collecting. They are different enough that little would be gained by disputing which is deeper, although both are obviously deeper than the things one needs to understand about a stone in order to use it to hold down a picnic blanket.

The important point for our discussion is that these various kinds and levels of understanding can all be fully described and analyzed without ever talking about a mental object—some internal representation of the computer, the eggbeater, or the stone—which embodies the understanding. Understanding is an attribute of the person's relation to the thing that is understood. To behave with understanding is to act in ways that are attuned to relevant properties of the thing. To have deep understanding means that one is attuned to nonobvious structural or causal properties of the thing and to that thing's relations to other things. Of course, there are problems in characterizing understanding in this way, but it does not seem that they are alleviated by introducing the idea of mental content.

As far as understanding in the practical affairs of life goes, I have been preaching to the faithful. That is, what I have been proposing as a way of treating understanding is the way people already treat it in most walks of life. It is only when we talk about understanding for its own sake, the kind of understanding that is propagated in the schools and the learned disciplines, that we start positing schemas, mental models, or concept nets, and treat understanding as a property of those mental objects. I shall argue that that shift is unnecessary, that the way we treat practical understanding is a good way to treat theoretical understanding as well. As a stepping stone to that radical notion, however, I want next to consider what it means to understand a tool.

Understanding a Tool

Concepts, theories, explanatory principles, and the like may be thought of as tools. And so understanding the Newtonian concept of force, for instance, ought to be something like understanding a pipe wrench. Of course, there is a level of understanding a pipe wrench that requires understanding the Newtonian concept of force, but among people who use pipe wrenches, few would have reason to understand them in that scientific a way. Yet clearly there are differences between untutored people whose understanding of pipe wrenches is superficial and inadequate and those who can be credited with deep understanding.

There is a truly wonderful study of depth of understanding as it appears among those who live in the world of pipe wrenches and the like. It is *Working Knowledge* by Douglas Harper (1987). Harper devoted the study to Willie, who ran a repair shop in a rural area and was a genius at fixing everything that came his way. In our book on expertise (Bereiter & Scardamalia, 1993), we used Willie as the prototype of the expert, partly to overcome the stereotype of experts as people who wear white coats and have doctor's degrees. The following is our summary of Harper's findings about the nature of Willie's genius:

> It becomes immediately clear that Willie's skill is primarily mental. But it is not merely a matter of raw mental power, although something like that may also be involved. Willie knows a tremendous amount. He has a deep knowledge of materials, what can be done with them, how they will respond to various treatments. He understands his tools in the same way. His ability to fix things depends on understanding how they were made and why they were designed the way they were, and this often includes understanding shortcomings in manufacture or design that were responsible for the eventual breakdown. With this understanding, Willie is often able to modify the object in the process of repairing it, so that it is better than it was originally. His repair of a broken car door handle, for instance, does not involve simply fixing or reproducing the broken lever. It involves inventing a new lever that fits into the same handle but is less vulnerable to stress. Even Willie's manual skills have a basis in understanding. He understands why filing the teeth on a chain saw often results in rounding the edges, and this enables him to handle the file in such a way that it cuts straight.[7] This is not to say that Willie's knowledge is all conceptual, explainable. He knows how things sound or look or feel under different circumstances, and often there is no way to convey this knowledge except by demonstration. But it is knowledge nonetheless. A brilliant general problem solver, if there is such a thing (it is what Sherlock Holmes purported to be), would be helpless trying to solve Willie's problems without Willie's knowledge.

Let us try mapping this kind of understanding on to the set of characteristics developed earlier for understanding a person:

1. What constitutes understanding depends on your relationship to pipe wrenches—whether you are an ordinary user, a plumber, a hardware dealer, a tool designer, a physics teacher, etc.
2. Understanding is intimately bound up with—but not the same as—ability to use and treat the tool intelligently.

3. Understanding is also intimately bound up with interest. Failure to appreciate the ingenious design of a pipe wrench would suggest lack of understanding.
4. Understanding a pipe wrench may depend on understanding its relation to other things—for instance, how it is different from a monkey wrench and from a vice grip, what it is about pipes that makes a pipe wrench suited to gripping them.
5. Understanding of a pipe wrench is not necessarily accompanied by ability to explain. Explanation is an indication of understanding, but, more importantly, it is a *means* by which understanding is developed and shared.
6. Although there is no single correct, complete, or ideal understanding of a pipe wrench, there are recognizably wrong understandings, and these are potentially correctable.
7. Discussions aimed at advancing understanding of pipe wrenches will seldom make reference to mental states or mental contents of the discussants. Instead, discussion will focus on pipe wrenches themselves, their uses, and their relations to other things in the world.
8. A major way in which understanding of pipe wrenches is developed, manifested, and shared is through deliberating on and sharing insights into *problems* involving them.
9. Having a *deep* understanding of pipe wrenches means understanding deep things about them: why they are designed the way they are, the functions of various parts and design features, causes of faulty performance, underlying physical principles (this is where understanding force would come in), etc..
10. Deep understanding of pipe wrenches is most clearly demonstrated by insightful solution of problems concerning them.
11. Deep understanding normally implies deep and complex involvement.

Although people are vastly more complex than pipe wrenches, it seems that we can apply all the same principles in elaborating what it means to understanding them. Will the same principles, then, suffice to characterize the kinds of understanding that school curricula are supposed to promote?

Situational Understanding

Often the thing to be understood is not a single object or kind of object or event but is a whole situation or complex. We talk about people “knowing the ropes” or “knowing their way around.” To say that you understand baseball or banking or human relations or, for that matter, that you understand numbers or Latin American history is to say not only that you understand a lot of individual things but that you understand the whole complex of which they are a part. What we may call ‘situational understanding’ fits all eleven of the characteristics previously discussed with reference to persons and tools. The connection to intelligent action is particularly strong, however. Knowing the ropes or knowing ones way around implies an ability to act intelligently and resourcefully in the situation or domain but it also clearly implies that this ability is rooted in thorough knowledge. It is not at all the same as native wit. A none-too-bright person who knows the ropes will often outsmart the ignorant genius.

What does it mean, however, to understand a whole situation or domain, other than understanding its parts? The geographical metaphor is helpful here. Knowing your way around in a city or a wilderness means, at the minimum, knowing how to get from one place to another by efficient routes. It is something more than knowing a lot of individual landmarks and routes, but is it understanding? In a modest way it fits the description. It is having a relationship to the place that is sufficient to support intelligent action of a certain kind. There are, moreover, deeper levels of understanding that also go by the name of ‘knowing ones way around,’ and that support more wide-ranging intelligent action. They amount to knowing deeper things about the place. We once vacationed on a small island in the Caribbean, staying in a house owned by a Toronto friend. The friend advised us to put ourselves into the hands of one of the taxi cab drivers who would meet the boat and let him “look after us” for the week. His looking after us turned out to include not only taking us places but finding us tuna and lobster on an island where these were thought to be unavailable, taking us to the right beach for the weather, fetching the police detective when our house was robbed, and generally solving any problems that arose. He knew his island, obviously, knew it better than most of the others who lived there, because he got around more and solved more various problems.

You learn your way around geographically by solving navigation problems. You learn your way around the shops and markets by solving supply problems. In general, you learn your way around in a domain by solving problems in it. Yet the textbook problems that are the mainstay of school mathematics and science seem not to have that result. It is not that the problems are not challenging enough. The problems you solve in finding your way around in a city may all be relatively easy (especially if you are not averse to asking directions), but they do the job. They engage you in putting information together in different ways to achieve different destinations or itineraries. Eventually everything fits together into what is often called a mental map, although that is not an apt metaphor. When you really know your way around you know where everything is in relation to yourself, wherever you happen to be, not in relation to points on a map. Textbook problems ought to have you putting information together in many different ways to achieve many different purposes. They ought to engage you in the equivalent of criss-crossing terrain in many different directions. In the end, you ought to have something that might be called a mental map or a mental model except that, again, you are not aware of consulting a model; you simply feel confident that you know how to deal with problems in that domain.

What we know about textbook physics problems, which have been the most studied, is that they can often be solved without reference to the physics principles they are supposed to be teaching (Sweller, 1988) and that students tend to classify problems according to the concrete situations they are about (inclined planes, pulleys and weights, etc.) rather than according to the principles involved (Chi, Feltovich, & Glaser, 1981). This sounds like students are learning their way around in a domain, all right, but it is the domain of textbook physics problems, not the domain of physics, of which those problems occupy a small and isolated part. They may acquire understanding, which supports intelligent behavior in the context of doing textbook problem assignments, but it is not an understanding of the physical world, and much less is it an understanding of the theoretical world from which textbook problems are drawn.

Theoretical Understanding

Perhaps the most thoroughly researched problems in theoretical understanding are those associated with Newtonian mechanics. Although such principles as universal gravitation, inertia, and the relationship between force and acceleration are common parts of the school curriculum and are usually taught systematically in beginning physics, there is evidence that most students, including many who have studied physics, understand motion in a preNewtonian way. They believe continual force must be applied to keep a body in motion and that bodies propelled along a curved path will continue along the curved path (McCloskey, Caramazza, & Green, 1980).[8] Although they may be able to solve problems that explicitly deal with acceleration, they seldom make spontaneous use of this most fundamental concept of Newtonian mechanics in solving other problems or in explaining phenomena.

Let us, accordingly, use Newtonian mechanics as an example for investigating the extent to which theoretical understanding may be treated in the same way as practical understanding. The eleven observations made earlier about understanding people are here applied to what we shall for simplicity call 'Newton's theory.'

1. What constitutes understanding Newton's theory depends on your relationship to it. Different kinds of understanding (not just different levels of understanding) are appropriate to an elementary science teacher, a historian of science, a mechanical engineer, a theoretical physicist, and a croquet player.
2. Understanding is intimately bound up with ability to act intelligently in relation to the theory. Acting intelligently with regard to Newton's theory might mean using it intelligently as a tool or it might mean using it to make sense of physical phenomena. A person who has thoroughly "internalized" Newton's theory will perceive and act in ways consistent with the theory, without having to recall or deliberately apply its principles.[9]
3. Understanding is also intimately bound up with interest. Although exceptions are conceivable (as they are in the case of understanding persons), it seems a reasonable supposition that the person who has no interest in Newton's theory does not understand it.
4. Understanding Newton's theory depends on understanding its relationships to cosmology, vectorial representations of forces, etc. (Reif & Heller, 1982)
5. Understanding of Newton's theory is not necessarily accompanied by ability to explain. Explanation may,

however, play a crucial role in acquiring and furthering its understanding.

6. Although there is no single correct, complete, or ideal understanding of Newton's theory, there are recognizably wrong understandings, and these are potentially correctable.
7. Discussions of Newton's theory will seldom make reference to mental states or mental contents of the discussants. Instead, discussion will focus on the theory itself, its implications, applications, limitations, and so on.
8. A major way in which understanding of Newton's theory will be manifested is through narratives in which ideas such as gravity, acceleration, and inertia figure. An account of placing a communications satellite in orbit, for instance, could be expected to reveal quite a bit about the extent and accuracy of the speaker's understanding of Newtonian physics. Inadequate understanding would be manifested in narratives that are incomplete, implausible, and incoherent.
9 Having a *deep* understanding of Newton's theory means understanding the deeper things about it—derivations, proofs, nonobvious implications and applications.
10. Deep understanding of Newton's theory is most clearly demonstrated by insightful solution of problems involving it.
11. A deep understanding of Newton's theory can only arise from deep involvement with the theory—from thinking about it a lot and from various angles, from using it in various contexts and for various purposes.

In short, we can rather satisfactorily characterize understanding of Newton's theory in the same way that we can characterize understanding a person or a pipe wrench. There are differences, but the differences tend to place Newton's theory between the other two rather than off on a different dimension. That is, there are more varied ways to understand a person than there are ways to understand Newton's theory, but more ways to understand Newton's theory than there are ways to understand a pipewrench. A similar ordering obtains with respect to potential for intelligent action, for interest and feelings, for complexity of relationships, for the value of explanation, and for levels of depth of understanding.

If we are to treat understanding a theory in the same way that we treat understanding material objects, we run into an ontological problem that continues to arouse a good deal of opinionated dissent among philosophers. A glance back over the preceding points will reveal that each of them implies that Newton's theory is a real thing enjoying an existence of its own independent of our cognition. For, otherwise, how could we stand in relation to it, have an interest in it, know things about it, and so on? Accordingly, this view of understanding requires accepting the the idea of *conceptual artifacts,* as put forth in the preceding chapter. With the help of this notion, the point I have been trying to make can be stated in a more general form as follows: *We should treat the understanding of conceptual artifacts in the same commonsense way that we treat the understanding of material objects—as inhering in the relationship between the person and the object rather than as a characteristic of a different kind of object located in the person's mind.*

Understanding and Intelligent Action

'Understand' is a transitive verb. When we speak of understanding, we always speak of understanding something. The thing that is understood—the object of understanding—may be, as in previous examples, a person, a mechanical device, or a theory. Or it could be something less easily specified, such as a situation or a whole field of study or an aspect of life, such as love. Regardless, understanding implies something about the relation between a person and an object. The point of the preceding discussion has been that instead of speculating about a mental state that constitutes understanding we should be looking at the nature of the relationship between the person and the object of understanding. What is it about such a relationship that makes it a relationship of understanding as distinct from a relationship of liking or disliking, owning, depending upon, admiring, having doubts about, being in control of, or any of the myriad other ways a person may relate to something else?

Many dog owners will insist and can marshal evidence to show that their dogs understand them. According to the 11 points by which we characterized understanding of a person, it is conceivable that a dog could qualify to some extent on all of them except those that involve explananation and narrative. But, still, how are we justified in calling it understanding when there is no verbal evidence? Migratory birds may act in accordance with principles of celestial navigation, and this implies that something in their nervous systems is designed to carry out computations according to those principles, but, no matter how impressed we may be by this feat, we are not tempted to credit birds with even the most limited understanding of the subject. We do not imagine that if the birds could talk they would be able to expound on the principles of celestial navigation, but it is only fanciful and not blatantly silly to imagine that if the family dog could talk it would have some wise things to say about us. Why the difference?

When we attribute understanding to our dog it is on the basis of its whole relationship to us, which we take to be complex, sensitive, flexible—in a word, intelligent. We have no reason to suppose that such a relationship exists between a migratory bird and the stars (although we expect it to exist between an astronomer and the stars). The conventional, folk-theoretic way of construing understanding is that it is a property of the dog's mental makeup, which is the *cause* of the intelligent characteristics of the relationship or which is *manifested* in various aspects of its relationship to us. Picking up on the intuition of Brownell and Sims (1946), however, we can eliminate an unnecessary intermediary and simply declare that understanding is an *aspect* of the relationship. The relationship may also have many other aspects. Our dog's relationship to us may be characterized as affectionate, dependent, playful—characteristics that carry no implication of understanding.*Understanding refers to that aspect of a relationship which has to do with its potential to support intelligent action*.

Notice that I am not saying that understanding *is* intelligent action or the ability to act intelligently. Intelligent action, as I pointed out before, may depend on a number of factors in addition to understanding: opportunity, movitation, enabling skills, general mental abilities, and so on. What I am proposing is a relational conception of understanding. Understanding, accordingly, is not a thing—an object in the mind—nor an attribute of a mind. Understanding refers just to the relation between an actor and an object. It is the relationship viewed from the standpoint of its ability to support intelligent action.

In practical terms, this means that understanding is a precondition of intelligent action. No understanding means no intelligent action. You might immediately object that you do not need to understand how a microwave oven works, for instance, in order to use it intelligently. But of course you do need some understanding. If you had grown up in a remote village and had no acquaintance with either electrical applicances or ovens, a few how-to-do-it lessons would not be enough to enable you to use a microwave oven intelligently. The woman ahead of me at a take-out lunch counter, who protested in horror when the counterman proposed to warm her soup in a microwave, also lacked sufficient understanding to behave intelligently. And from my own experience as a heavy user of the microwave in cooking, I feel sure I could get better results if I understood its workings in greater depth. That is true of most complex artifacts, whether material or conceptual. But do you need to understand how a multidigit multiplication algorithm works in order to use multiplication intelligently? That is a live issue in mathematics education. I think the answer in this case is clearly no. You have to understand multiplication, but that is something different. The justification for trying to get children to understand the arithmetic algorithms they use is of a different order; it is to build up the conviction that mathematics in general makes sense, that it is not beyond understanding.

Feelings of Understanding

Understanding, or its opposite, is often accompanied by feelings. There is the exciting "Aha!" sensation of sudden illumination and the delightful feeling of 'getting it' when understanding develops smoothly. There are the feelings of bewilderment, confusion, and frustration or sometimes instead feelings of wonderment and anticipation when we are confronted with something we do not understand. These feelings no doubt have much to do with our commonsense assumption that understanding is something in the mind.

Wittgenstein (1980) examined at length the feeling of understanding and concluded that there is no particular feeling, such as there is with grief or joy or pain or warmth. "'Understand,'" he said, "just is not used like a word for sensation" (par. 311). Contemporary psychologists speak of 'the illusion of understanding.' Wittgenstein would probably have regarded this as a misuse of terms. True, we sometimes feel that we have understood something and it is later revealed that we did not, but our mistaken feeling is nothing like an optical illusion.

What is it, then, to *feel* that we understand? Taking a cue from Karl Popper (Popper & Eccles, 1975, p. 44), I would say that it is a feeling of confidence. It is confidence that we know how to proceed, how to deal as appropriate with the object of understanding. If someone has given us directions, we feel that we understand them when we feel confident that we know how to get to the destination, how to assemble the cabinet, or whatever. Feeling that we understand a theory, Popper said, means feeling confident that we can reconstruct it at will. In studying for an examination, we feel that we understand when we are confident that we will be able to answer the kinds of questions likely to be asked.

Like other kinds of confidence, confidence in our understanding may turn out to be ill-founded. The same is true of feelings of unconfidence. The directions we receive may seem unclear, but as we proceed they may prove to be perfectly adequate. But feelings of confidence and unconfidence are essential guides to our cognitive efforts. They constitute a form of what I will discuss in the next chapter as ‘impressionistic knowledge.’ It is feeling-based knowledge that plays a part along with other kinds of knowledge and is especially important in making judgments. Like other kinds of knowledge, it is improvable. Young children often claim that they understand a text as soon as they are able to say all the words. One thing that traditional classroom recitation does is build up a better anticipation of the challenges that text understanding may face. This is bound to have an effect on confidence, requiring more thorough study of a text before students feel sure of their understanding. An important part of becoming sophisticated in any scholarly discipline amounts to educating your feelings so that your confidence in understanding becomes an increasingly reliable guide.

This ‘education of feelings of confidence’ is an essential part of developing a close relationship with another person. Traditional courtship was carried out under such constraints that it hardly served much purpose in acquiring a deep and detailed understanding of the prospective mate. What it did enable a couple to do, however, was develop confidence in their feelings toward the other and the other’s feelings toward them. If the marriage went well, their confidence would become elaborated as new information emerged and would evolve into feelings of understanding—confidence in their ability to predict and interpret the actions of the other and repair and improve their relationship. We may liken this to the early stages of learning a discipline. You cannot know enough at an early stage to have a warranted feeling of understanding, but you should begin to acquire confidence that the discipline is understandable and that you have the capacity to understand it. As learning proceeds, this feeling is not supplanted but evolves into confidence that you can explain, apply, find out what is needed, correct errors—thus, improve and repair your relation to the objects of understanding.

Only occasionally would confidence in understanding be a result of conscious analysis. More often it appears as a spontaneous feeling, which is why we experience it as a direct feeling of understanding and why, accordingly, we locate understanding as something in our minds. But most feelings of confidence are spontaneous. If you have to jump over a ditch you may feel confident you can do it or you may be in some doubt; only in the latter case will you consciously gauge the distance against an estimate of your jumping ability. Otherwise you just feel confident and jump. The same is true of more complex psychological matters. For instance, your confidence that you will be able to handle yourself in an anticipated difficult social situation is likely to include confidence in understanding the situation but also confidence in your social skills, resourcefulness, emotional equilibrium, or even strength of character. In short, there is nothing very special about a feeling of understanding. It is similar to and often part of many other feelings of confidence or lack of confidence that are common in experience. All that is distinctive is what the confidence is about. It is confidence about knowledge. Specifically, it is confidence about the adequacy of our knowledge to support intelligent behavior in foreseeable situations—confidence that we will know how to proceed, that we will be able to treat the object of understanding in a suitable way.

Although feelings of understanding are important, as a kind of knowledge that guides action, we should not over-rate them and above all should not make the mistake of equating understanding with feeling. Most of our understanding is not accompanied by any feeling. If the ditch is only a foot wide, we do not feel confident that we can get over it; we simply step over it and think nothing of it. Similarly, when reading easy material we have no feeling of understanding it, we simply pick up the information. Our thoughts are on what the text is about, not on our understanding of it. Understanding, as I have argued, is a relation between the knower and the object of understanding. As with other relationships that shape our lives, we may only occasionally be conscious of it and have feelings associated with it. Most of the time we just go about our business and our feelings are about the things and events we encounter, not about our relations to them. But there are times when our relations are problematic or when they undergo a sudden change. At those times feelings arise that importantly affect our actions (Oatley, 1992, Ch. 4). It is no different with feelings of understanding.

Inert Knowledge

"Inert knowledge," as it is often referred to these days, is knowledge that is of little or no help in supporting any kind of intelligent action. A better way to conceive of it, however—and the way that I understand Whitehead (1929) to have conceived of it when he introduced the term—is as *theoretical knowledge that supports intelligent action only within the situations in which it was taught.* It is not limited to rote knowledge, but it does tend to be limited to knowledge helpful in performing intelligently in examinations. Examinations might well be intended to test understanding, and so might contain items at all levels of Bloom's taxonomy. Doing well on such examinations requires intelligent action, yet Whitehead was saying that there are students who can perform that kind of intelligent action in a knowledge domain while never taking advantage of the knowledge in any other context. Whitehead blamed this phenomenon on the encapsulation of the disciplines, each a world unto itself, taught as if the multifarious real world did not exist. I suspect that a more general kind of encapsulation takes place in schooling. Feelings of understanding or nonunderstanding, as I remarked in the preceding section, are conditioned by what we have experienced as the challenges that understanding may face. Over the course of years of schooling, successful students develop a fine sense of when their understanding is up to the mark—when they have studied enough and have pursued the right question in sufficient depth. That sense is derived from the recitations and examinations that have constituted the main challenges to understanding. No matter how seriously teachers may strive to pose questions and examinations that force fuller understanding, what they instill is inevitably a poor and narrow sense of what constitutes adequate understanding. It stops well short of the kind of relation to knowledge objects that supports intelligent action in the world at large.

I am not sure whether what I am about to say is in disagreement with Whitehead or merely an elaboration. It seems to me that the most serious encapsulation is not at the level of the disciplines but is at the level of the conceptual artifacts brought forth in the teaching of those disciplines. These tend to be terribly atrophied versions of the objects actually created and worked with in the living disciplines. Take theories, for example. In textbooks, a theory is treated as a usually quite small set of propositions or formulae, which the student is expected to learn in a literal fashion and then apply to set problems. But there is much more to a theory. Mario Bunge (1977-79, p. 162) described the multifaceted character of theories and the need for multiple viewpoints on them—asserting, among other things, that the set of propositions constituting a theory is infinite:

> Theories are variously described as ideal objects, systems of changeable meaning and truth value, growing bodies of knowledge, or prescriptions for doing things. These characterizations, though very different, are mutually compatible. In fact from a logical point of view a theory is an ideal object of a certain kind, namely an infinite set of propositions closed under deduction. That one and the same theory may be variously interpreted, and may be attributed different truth values on the strength of different bodies of evidence, is obvious. From an epistemological point of view theories are not static but they grow in certain directions (as more theorems are proved and more applications discovered) and shrink in others (as regions of falsity appear). Finally, from a pragmatic viewpoint theories, even mathematical ones, can be used as rules or prescriptions for computing or for designing experiments, much as musical scores are used as instructions for performing. Since theories are all four—logical objects, semantic systems, growing bodies of knowledge, and prescriptions—according to the viewpoint that is chosen, there is no incompatibility among the four descriptions.[10]

Inert knowledge is typically revealed by the failure of students to apply it outside of class, even to problems similar to those worked on in class. Accordingly, it is linked to the hoary problem of transfer, which suggests that the antidote is to be found in varied applications in ‘authentic’ or real-world situations. From the perspective offered here, however, this should not be seen as the cure for a specific problem but rather as one important aspect of the much larger educational problem of establishing intimate working relations between students and conceptual artifacts.

Going back to our previous analogies between understanding physical objects and understanding abstract ones, we may liken theories to the most complex of manufactured objects, which at this time are computers. If computers were presented in textbooks the way theories are, we could expect inert knowledge of computers to be the result. We would not expect this knowledge to play much of a part in intelligent action. Having varied authentic experiences in the use of computers is obviously an important part of acquiring a more useful knowledge of them. But it is only a part. Many of us with ample and varied experience in computer use remain severely limited when it comes to coping with anything outside the routine. The people we look to for help are ones who have a deeper understanding than we do (although not necessarily a deep theoretical understanding) . They qualify on all the eleven points discussed earlier. Compared to the complexity and many-facetedness of their relationship to computers, ours is one-dimensional. This complex relationship was developed and continues to grow through intense involvement. They read, they tinker, they experiment, they think a lot, they talk to people who are similarly involved. A theory, I am saying, is also a complex object. In order for it to play a significant role in our behavior—that is, to support intelligent behavior in spheres of activity that are important to us—we need to develop the kind of complex, intense relationship with the theory that we would have to develop with a machine or with a human co-worker under similar circumstances.

Naturalizing Abstract Knowledge Objects

Discourses about understanding commonly start off with an admission that understanding itself is not well understood (cf. Brownell & Sims, 1946; Nickerson, 1985). I have been arguing, on the contrary, that we already understand a great deal about understanding as it applies to the ordinary affairs of life; it is just that this knowledge tends to get left outside the school door. We know how to acquire an understanding of the various living and nonliving material things that we deal with in everyday life—our dogs, cars, neighbors, television sets and VCRs, espresso machines, African violets, pomagranites, income tax forms, cold tablets, martial arts academies, and so on. There is no single recipe for achieving an understanding of all these things. Each is distinctive, and yet we would seldom find ourselves at a loss in pursuing understanding or in guiding someone else in the pursuit of it. Not that we would care to pursue understanding of all of them: understanding takes time and effort, and so we have to allocate our resources. But if we did decide to pursue understanding of any of these, we would tackle it as a problem. What is it about the object in question that we trying to understand and to what purpose? Having framed the problem—perhaps initially in rather vague terms—we would set about solving it. In the process we would likely reframe the problem (as understanding grows, so does our recognition of what needs understanding). We would never ask, “Have I finally achieved understanding?” No such end is in view. But we would ask whether we had achieved understanding sufficient for our purposes.

At the center of our attention always would be the object itself, what we are trying to understand. Its context would be the spheres of activity in which the object matters to us. This could be a very limited sphere of practical activity, thus calling for a very limited understanding (such might be the case with television sets, for most people), or it could be a very general effort to make sense of the world. Together, the object and the situation or situations in which it figures for us are all we need when it comes to framing problems of understanding. We do not need to occupy ourselves with structures in our minds. It would not even occur to us. We may have subjective impressions of understanding or bewilderment and will pay attention to those; but we would treat them as no different from the feelings of confidence or uncertainty, satisfaction or discomfort, that we feel in any kind of problematic endeavor. They would not be definitive of the problem we are trying to solve or direct indicators of our progress in solving it.

We are accomplished in pursuing understanding of things in the material world because we have been working at it since infancy and in many different situations. But when it comes to understanding abstract objects such as theories and scientific, mathematical, or literary concepts, our experience is of later origin and much more limited. Even if we are fortunate enough to encounter schooling that emphasizes understanding, it is likely throughout elementary school, and perhaps through high school as well, that the focus will be entirely on understanding the natural world directly and not on understanding theories that purport to explain the natural world.

For most people, accordingly, the world of abstract knowledge objects—of conceptual artifacts—is an unreal world. In contrast, mental content does exist, or seems to; we seem to experience it and it seems to have a location, in our heads. Feelings of understanding and not understanding are undeniably real. And so, instead of attending to abstract objects we attend to mental states. This is satisfactory for many purposes, but it means that the whole complex of skills that we acquired for dealing with problems of understanding no longer applies because there is no longer a *thing* to apply those skills to.

My proposal is simply that we naturalize abstract objects, treat them as real things that we can nudge about and look at from different sides, take apart, try out, become fascinated with, discuss with our friends, and try to reinvent. We already do this with very familiar abstract objects, such as numbers. Unless we stop to worry about their ontological status, we will almost surely treat the whole numbers as enjoying an existence of their own, apart from the various ways of representing them concretely and symbolically. The number 25 has properties of its own, such as being odd and the square of 5 and the square root of 625, which are not properties of the Arabic or Roman numerals or of a collection of 25 bananas or of a 25-cent piece. Understanding numbers means understanding and being able to make intelligent use of these properties. Children brought up in a numerate environment have already acquired a considerable understanding of these abstract objects well before their first encounters with school arithmetic (Griffin, Case,& Siegler, 1994). Numbers can be fun to play with, and playing-with is a child's way of developing a relationship. Ironically, however, when children enter school they are liable to fall into the hands of an adult who believes that they must be shielded from abstraction and so they will spend their time working out puzzles with blocks and tokens, and later with numerals, without any discussion of what these exercises reveal about numbers in their own right. In the more reflective approaches to primary mathematics, however (e.g., Cobb, Gravmeijer, Yackel, McClain, & Whitenack, 1997), even though concrete representations may be used, children are encouraged to think about numbers as such—to establish meaningful relationships with these abstract objects in the same way that they establish meaningful relationships with the people, animals, and implements in their environment.

Other abstract objects that have managed to secure a place in the natural order include stories, poems, jokes, and songs. With respect to these, understanding also has a natural meaning that does not involve suppositions about mental representations. To 'get' a joke is to *see* what is funny about it. You may *see* it and yet not *like* it. *See* and *like* are relational terms. They do not refer to something in your head but to something that obtains between you and something else. Extending that relational view of understanding to the objects dealt with in formal education is the nut of what I am arguing for in this chapter.

Education for Understanding

'Teaching for understanding' has emerged as one of the banner principles of educational reform (Cohen, McLaughlin, & Talbert, 1993). Probably few teachers at any time would have declared that they were not teaching for understanding, but its salience as a regulatory principle waxes and wanes. The current rise in concern has been helped along by mounting evidence of students' lack of understanding in academic areas, especially science. I have already alluded to the research on misconceptions, but ordinary factual knowledge testing also points to deficient understanding. For instance, in a recent survey of Canadians aged 16 and over, almost a quarter indicated that they believed the sun goes around the earth and half the respondents failed to choose one year as the time it takes the earth to go around the sun. These cannot be written off as isolated bits of missing or wrong information, like not being able to name the 10 provinces (another task that Canadians often fail). They bespeak profound lack of understanding of the solar system.

A great deal of instructional research has dealt in one way or other with promoting understanding. This has been one of the main concerns of the enterprise known as 'Instructional Design Theory' (Reigeluth, 1983). Despite many variations, there seem to be three basic approaches to teaching for understanding. The time-honored, commonsense way is by direct explanation. It is the approach we spontaneously adopt in everyday life when we are trying to get someone to understand something. We try to explain it as clearly as possible and in terms the listener will be able to understand. We may make use of diagrams or analogies, anything that will get the idea across. Instructional Design Theory has mainly dealt with refinements of direct explanation. There is clearly much of value in the direct approach (Hirsch, 1996). In fact, it is a pretty sure bet that civilization would collapse without it.

The second approach to teaching for understanding may be broadly characterized as a 'process' approach. It is reflected in Bloom's *Taxonomy* and in Perkins's ''performance perspective.'It relegates knowledge to the status of reference material in a mental filing cabinet. All the intellectual richness resides in the cognitive skills that operate on this material. Such a standpoint offers no clear prescription for teaching for understanding, but it suggests that, besides mastering items of declarative knowledge, students ought to be engaged in a lot of interpreting, applying, analyzing, synthesizing, and evaluating. *Science: A Process Approach* (AAAS, 1967) exemplifies this approach; there is little direct focus on conceptual content but instead a focus on carrying out the various processes thought to constitute the 'doing' of science.

The third approach, associated with contemporary cognitive science, is focused on students' mental models. 'Conceptual change teaching' (Anderson & Roth, 1989) is perhaps the most straightforward exemplification. It involves trying to determine the student's present understanding and to devise some kind of experience that will change the student's mental model. There is no commitment to any particular kind of educational experience. It could involve direct explanation, problem solving, experiments or demonstrations, or discussions or debates among students upholding different ideas. The approach is essentially constructivist, however. Accordingly, it is presumed that conceptual change must take place through cognitive activity of the learner, that it cannot be merely conveyed.

The mental models approach is in the ascendancy, and it has the virtue of being able to incorporate advantages of the other two approaches as well. Still, something important is lacking, which is implicit in the direct explanation approach. It is close attention to the things to be understood. By casting everything in terms of mental content, the mental models approach focuses attention on the shortcomings of students' mental models rather than on the richness of what there is to be learned. Major theories have great depth and wide implications. Coming to understand a living theory means establishing a many-faceted relationship and one that will keep developing as one's experience grows and as the theory itself evolves. This grand sense of *what is there to be understood* —a sense that has guided traditionalists of the classic variety for centuries—tends to get lost in what amounts to a mental remodeling operation.

Yet a constructivist view of learning is in no way inconsistent with a regard for the abstract objects of understanding. This is most clearly shown in the writings of Karl Popper. On the one hand, Popper advanced the idea of a world of abstract objects existing in a way indpendent of human cognition—World 3, as discussed in the preceding chapter. On the other hand, he regarded coming to understand such objects as a highly constructive mental process:

> According to my view, we may understand the grasping of a World 3 object as an active process. We have to explain it as the making, the re-creation, of that object. In order to understand a difficult Latin sentence, we have to construe it: to see how it is made, and to re-construct it, to re-make it. In order to understand a problem, we have to try at least some of the more obvious solutions, and to discover that they fail; thus we rediscover that there is a difficulty—a problem. In order to understand a theory, we have first to understand the problem which the theory was designed to solve, and to see whether the theory does better than do any of the more obvious solutions. In order to understand a somewhat difficult argument like Euclid's proof of the theorem of Pythagoras (there are simpler proofs of this theorem), we have to do the work ourselves, taking full note of what is assumed without proof. In all these cases the understanding becomes “intuitive” when we have acquired the feeling that we can do the work of reconstruction at will, at any time. (In Popper & Eccles, 1977, p. 44.)

According to the view I am trying to advance, educators ought to be less occupied with what is in students’ minds and more concerned with what kinds of relationships are developing between the students and those conceptual artifacts that find their way into the curriculum. Formal education is inevitably much involved with conceptual artifacts. They might be and often are kept at bay in the elementary school, but it would make a travesty of education to keep them out of the picture indefinitely. They are what the sciences and learned disciplines produce; and formal education, regardless of the philosophy it ostensibly follows, must bring about some kind of relationship between students and such abstract objects. I do not see how anything less than that can be regarded as ‘teaching for understanding.’

Teaching for understanding entails more than that, however. Students are not likely to find most abstract knowledge objects very engaging in their own right. As Popper said, the first step in understanding a theory is to understand the problem it is intended to solve. Understanding the problem may require understanding other abstract objects but it also requires some engagement with the real world phenomena that give rise to the theoretical problem. It is the same as with understanding a tool or a machine. Students cannot be expected to understand a clock until they understand the problem of metering time and why it matters. The problem has to be a problem to the students, in other words. And so teaching for understanding, in all its creditable versions, occurs within the context of students' efforts to understand their world.

The several approaches to teaching for understanding that I have touched on in this section represent a descending scale of definiteness about the nature of the teacher's job. In the didactic approach, the job is clear-cut. There are certain concepts and principles to be taught and the teacher's job is to present these in a lucid enough way that they will be understood. In the 'process' approach this clarity of purpose is lost, but in its place are activities and skills that lend themselves to definite plans and to concrete action on the teacher's part. The conceptual change approach, however, asks teachers to get inside students' heads and to adjust their actions according to what is found there. As Roth (1992) observes, this task is so elusive that many teachers who declare in favor of a conceptual change approach are found actually to be pursuing one of the preceding two. But what, then, is a teacher to do on being told that understanding is a relation between the knower and the object of understanding and that the educational task is to make this relationship a rich and functional one? Shake one's head and walk away sounds like a reasonable response.

All I can do in defense of this seemingly outlandish fourth approach is to hark back to what was said earlier about everyday understanding. If your task as a teacher was to develop deep understanding of a radial arm saw, for instance, you might use elements of the first three approaches. There would be some stating of facts and explaining, but it would be limited to what seemed to need stating and explaining. Needed in what sense? Well, needed in order to proceed or needed because it is something that puzzles the learner or is likely to be overlooked. You might, for instance, feel it necessary to explain why, in doing a rip cut, you position the saw so that the teeth are coming up instead of down against the wood. Or you might even need to teach some trigonometry in order to get the student to understand the making of angle cuts. You might prescribe some experiments and skill practice, but these would be directly relevant to using the saw as a tool. You would not fancy that you were developing spatial reasoning skills or a woodworking schema by this means. And you would be on the watch for misconceptions, such as might arise from faulty generalization from other kinds of saws. But all of these would be subordinated to the purpose of helping the learner to become an intelligent and versatile user of the saw, aware of its capabilities, limitations, and dangers, familiar enough with its mechanisms to make necessary adjustments, recognize when something was out of whack, and improvise when novel problems were encountered. Theories are tools, too. Their main use is in explaining. Why should we take a more simplistic approach to understanding a scientific theory, something that represents a culminating achievement of civilization, than we do to understanding a piece of shop equipment?

Forms of Understanding

Through a series of brilliantly reasoned and erudite books, Kieran Egan (e.g., 1979, 1988, 1997) has developed an educational theory that centers on understanding but that seems at first sight to stand on an island of its own apart from all of the educational thought that has swirled around this topic and that I have advanced or criticized here. It is indeed original but I think it can be brought on to the same continent as the ideas I have been discussing, where it can be shown that on one hand it shares some of the same weaknesses as other theories grounded in folk psychology and epistemology while on the other hand it it is compatible with and adds significant dimensionality to the view of understanding I am advocating.

The central idea, which Egan has derived mainly from cultural history. is that there are different kinds or forms of understanding, each one growing out of its predecessor. There is at bottom a Somatic form of understanding, which comes with the equipment, so to speak, and is not unique to human beings. Then there is Mythic understanding, the characteristic way that preliterate societies make sense of the world. Literacy makes possible Romantic understanding, which has the manifold characteristics commonly associated with the Romantic Age. It is a richly textured understanding that always has human beings and their feelings as integral to the understanding, unlike Philosophic Understanding, which has the detached character of science, mathematics, and the more rigorous forms of philosophizing. Finally there is Ironic Understanding, which might also be called wisdom; it recognizes the limitations of all forms of understanding.

To the extent that it is concerned with understanding at all, formal education may be seen to focus almost exclusively on Philosophic Understanding. According to Egan this is poor practice on two counts: First, it doesn't work. Children cannot go straightaway to Philosophic Understanding. Egan argues that they need to in effect recapitulate cultural history by developing first Mythic and then Romantic Understanding. In fact, Egan argues, this is what children normally do in their own efforts to make sense of the world, and so schooling would do well to take advantage of this naturally occurring evolution rather trying to override it. Second (and this is a much deeper argument) Philosophic Understanding by itself is sterile, bloodless, and overweening—as Romantic critics have often alleged. It needs to incorporate or preserve elements of Somantic, Mythic, and Romantic Understanding in order to be fully human (and in order, furthermore, to provide the groundwork for later development of Ironic Understanding).

With some reservations about the recapitulation notion, I find all the rest of this to be good medicine. From Egan's vantage point, my treatment of understanding will be seen to focus too narrowly on Philosophic Understanding—particularly its subtype, Theoretical Understanding. The reason for this emphasis, however, has to do with the purpose of this book, which is not to deal with everything that matters in education but instead to deal with the shortcomings of folk theory of mind for addressing the needs of a knowledge society. Those shortcomings are most pronounced in matters having to do with conceptual artifacts, which in turn are assuming more widespread importance in the emerging knowledge society. Although folk theory has weaknesses for dealing with all forms of understanding, it is, as I have already tried to show, at its worst when it comes to theoretical understanding. If someone tells you that understanding Keats's *Ode on a Grecian Urn* means having something in your mind that matches what is in the mind of an expert, you will immediately recognize that such a criterion is at least narrow and perhaps rather screwy. You do not need an enhanced theory of mind to realize that there is something personal and contingent about understanding a poem or a work of art. It is not much of a stretch to think of understanding in these cases as involving the whole relationship between the knower and the object of understanding. But if someone tells you that understanding plate tectonics means having something in your mind that matches what is in the mind of a geologist, you may well accept that as reasonable. Many cognitive psychologists apparently do. I expect Egan would agree with me that such a conception of understanding is also narrow and rather screwy. Philosophic Understanding should involve the whole person as well. You should have feelings about ideas, intuitions that guide your thought. You should be excited by discoveries and insights; offended by distortion, oversimplification, and sloppy thinking; responsive to beauty in ideas; courageous in the face of difficulties and disapproval. (These are notions I will develop more fully in the next chapter.)

For Egan, the job of acquiring the various kinds of understanding consists of learning to use the intellectual tools that make such understanding possible. Mythic Understanding depends on the tools of oral language, especially story telling. Romantic Understanding depends on the expanded representational capabilities of written language. Philosophic Understanding uses the tools of logic, mathematics, the hypothetico-deductive method, and so on. Ironic Understanding involves the imaginative use of all of the tools "for putting into language meanings that the literal forms of language cannot contain" (Egan, 1997, p. 171). Although this treatment of intellectual tools may seem to resonate with what I have been saying about conceptual artifacts as tools, I see it as closer to Bloom's *Taxonomy* than to what I have been arguing for. Although it grows out of a different tradition, Egan's approach to teaching for understanding is, like that of Bloom and his committee, a process approach. Understanding, for Egan, consists of being able to use the appropriate 'intellectual tools'; for Bloom, it consists of being able to apply the appropriate 'intellectual skills.' Egan's 'tools' may be superior to Bloom's 'skills,' but they are alike in relegating content, the stuff to be understood, to the bottom of the ladder. Mind you, I don't think Egan wants knowledge to be down there. He shows far too much fondness for it himself for that to be a plausible surmise. But he is stuck with a folk theoretic framework, as were the authors of the *Taxonomy*, that almost forces knowledge to be put there. To bring knowledge out of the cellar and into the light, you need to be able to recognize that concepts, theories, aphorisms, and the like can also be tools. They are constructed with tools of the kind Egan recognizes, but they can function in turn as tools in their own right, as things we can use in making sense of the world.

It is important to realize, however, that conceptual artifacts are not *only* tools. What Mario Bunge said of theories can be extended to many other kinds of conceptual artifacts: the same conceptual artifacts may be variously regarded as “ideal objects, systems of changeable meaning and truth value, growing bodies of knowledge, or prescriptions for doing things.” Thus the relationship between a student and an idea has possibilities of great complexity and manifold values. In order to nurture such a relationship, educators need an epistemology that allows them to give full due both to the potentialities that inhere in the learner and the potentialities that inhere in the idea, the thing to be understood. If they see the idea as only an opinion, then they will see the student’s relationship to the idea as only some degree of belief or disbelief. If they see the idea as only an item of information, to be stored in the mind and acted upon with intellectual skills as the occasion arises, then they will see the student’s relationship to the idea as similar to the relation of file clerk to the items stored in files: Has the item in question been retained and can it be retrieved when needed? If they see the idea as only a tool, they will be far in advance of most other educators, but they will tend to see the student’s relationship to the idea in purely utilitarian terms, overlooking its more personal and emotional aspects and the possibility that the idea may change the way the student experiences the world.

Conclusion: What Does It Mean In Practice?

When I have tried out on educators the relational view of understanding developed in this chapter, they are generally sympathetic but then they ask, in a not too hopeful way, “What is a teacher supposed to do to produce this kind of understanding?” I think I have already offered the only answer that can be given for the general case: We all have a pretty good idea of how to teach for understanding when the thing to be understood is a material living or nonliving thing or an observable phenomenon. Just apply the same strategies when the thing to be understood is abstract, some kind of conceptual artifact.

Beyond that, we have to get down to particular cases. Do you want students to understand acceleration due to gravity? Then let’s start with what this thing is that is to be understood. It is an idea for which Galileo gets first credit. People long before Gailileo knew that things fall to earth at a fast clip, but they probably never thought much about the velocity. The implicit assumption was probably that things fall at a constant rate. Visualize a coconut falling from a tall tree. Now visualize this happening in slow motion. I'll bet that if you don't think much about it your mental image will be that of something falling at a constant rate. I suspect that we are not wired to perceive or visualize accelerating motion. If that is true, then the educational job set out for us is harder than we would have imagined from just examining the formula, $g = 32.174\ ft/sec^2$. We have got to establish a good working relation between the students and a scientific notion that is quite foreign to the way they naturally make sense of the world.

A good starting point—the point from which much of scientific inquiry starts—is things that people recognize they can't explain. Most of the questions students come up with spontaneously will have to do with outer space, so you may have to step in with questions of your own to get them to realize that there are things much closer to home that they don't understand. Why is a glass more likely to break if you drop it from three feet than if you drop it from one foot? There may be a variety of answers, but the most cogent one is likely to be that it takes a while for a falling object to get up to full speed. What do we mean full speed? Why is a glass dropped from a hundred feet on to a hard surface dead certain to break? Could it be that there isn't a 'full speed,' that things keep traveling faster and faster the farther they fall? Although acceleration is hard to measure with free fall, students can get measurable results with balls rolling down a long incline or things sinking in a tall container of water. Can something falling from a greater height ever overtake something falling from a lesser height? That is an interesting experiment that clearly demonstrates acceleration in falling objects.[11]

Up to this point, what I have described falls well within the range of things that inquiry-oriented elementary school teachers already do. Such teachers could adopt, criticize, or elaborate on these ideas without having to take on an esoteric new conception of understanding. Does this mean, then, that in practice the relational concept of understanding reduces to just another label for what several generations of educators have known by the name of 'learning by discovery'?

My first response to this question is to point out that what I have described so far does not involve the understanding of a conceptual artifact—the concept of *g*, for instance. Teacher and students are dealing entirely with World 1, the physical, observable world. World 3 is not in the picture. A teacher could do a reasonably effective job of shepherding children through the learning experience without himself or herself having much if any theoretical understanding. The understanding in question is understanding the phenomenon of things falling. As I have suggested, common sense is pretty good at handling this kind of understanding, as long as it doesn't get befuddled by abstract ideas. As long as teachers steer clear of Newton and concentrate on enriching students' understanding of concrete things and phenomena, common sense can guide them. But life since Newton is not that simple. The exploratory approach I have suggested will take substantial time and effort. To understand why it is worth the investment, why they should not dispense with the topic in ten minutes and move on to something else, teachers need to appreciate the significance of acceleration due to gravity. Understanding it is not important for daily life. Students already know enough not to jump off cliffs, even though they may not understand the concept. Its main value is as a start on understanding acceleration in its general, Newtonian sense—one of the most powerful concepts in physics yet one that eludes many students even after they have taken university-level physics.

Inevitably, however, the word 'gravity' will find its way into discussions of falling objects: What makes it happen? Gravity. But gravity is neither explained nor is it an explanation; it is more in the nature of a lexical filler. At the elementary school level that is all right. There are lots of things that children do not understand but that they still need words to talk about. However, especially if the word appears in curriculum standards, teachers are liable to feel that they have to teach it. It may even appear on a test. And so what might have been a meaningful inquiry into how things in the world behave when they fall risks degenerating into pseudoscientific verbalism. To avoid that, teachers need a clearer idea of what they are about. It needs to start with having a clear idea of what the *it* is that students are supposed to understand. A relational approach to understanding centers on building a sound and fruitful relation between the students and *it*. Pedagogy grounded in folk theory vacillates. Sometimes it goes all out for activities, which may include experiments of the kind I have suggested, except that there is no *it* there, no object of understanding that is the focus of all those activities. Other times it goes mental and concentrates on getting things into the students' heads or correcting what is already there. Again, the thing to be understood is lost sight of. In the first case it is replaced by activities. In the second case it is replaced by words. But in both cases understanding gets left behind.

No intellectually alive group of children or adults is going to remain satisfied with just describing the behavior of falling bodies. They are going to ask "Why?" Note that in the example we are working through, teaching for understanding does not start with that question. It begins with getting students on sufficiently intimate terms with the object to be understood that they can ask *why* questions of some meat. "Why do things fall down?" or "What is gravity?"—which are the naive questions people will ask—are not promising questions to pursue. However, "Why do things keep falling faster the longer they fall?" is a question that can lead right into the heart of Newtonian physics. The purpose of the preliminary exploration was to make that a meaningful question. If it is not, if the students do not find themselves wondering about it, the exploration has failed.

The question "Why do things keep falling faster the longer they fall?" poses a theoretical problem and it wants a theoretical solution. So this is the point at which instruction enters World 3 and where, accordingly, it is liable to lose hold of common sense. We are still interested in World 1, the world of real physical things falling from real places, but our interest also includes conceptual artifacts created to explain observations of such World 1 phenomena. The obvious target, the *it* with which to establish an understanding relationship, is Newton's laws. If we simply give students the forumula for acceleration due to gravity and teach them to make calculations with it, however, we will have given them a tool to be used without understanding. They will not understand where the number 32 comes from, and if we point out that the formula is an idealization and only applies in a vacuum, they will be justfied in asking why we have not given them a more useful formula, one that applies to the world they happen to live in. What can we say to that?

A little reflection will make several things clear. First, students are not going to be able to understand acceleration due to gravity in isolation. They are going to have to get into force and motion and something about friction. So we are really talking about understanding a domain, not understanding a single object. Second, students' "Why?" questions are not really going to be answered. The kind of “here’s how it works” answer they are looking for does not exist, and the latest scientific thinking on the matter is too far beyond them. Understanding gravitation, therefore, is rather like understanding another person. You cannot hope to understand everything.

My own inclination, drawn from some years of experimentation with students’ efforts to construct their own conceptual artifacts (Bereiter et al., 1997; Scardamalia & Bereiter, 1996; Scardamalia, Bereiter, & Lamon, 1994) is not to introduce Newton’s laws directly. Instead, have students construct their own theories, their own conceptual artifacts, and work to refine them. In the process they will encounter Newton, see that his laws are vital to their enterprise, and will start coming to terms with them in some way. Then, depending on the maturity of the students and the resources available, we can lead them into a thorough study of force and motion or we can allow them to make do with a partial theory that leaves unanswered questions. For instance, students might conclude that gravity keeps pulling on falling objects and the longer it pulls the faster they go. That is not a bad explanation for ten-year-olds to produce, although it leaves unexplained why the same rule does not apply to towing a wagon.

Understanding Newton’s laws, according to the argument of this chapter, means having a relationship to them that supports intelligent action. Leaving aside astronauts and engineers, the kind of intelligent action that an understanding of Newton’s laws might support is not likely to be of a practical nature. The action it can support is the pursuit of further understanding. This might range over understanding newspaper stories about space travel, understanding what your physics textbook means when it says there is no such thing as centrifugal force, and understanding how the gadget on your desk works—the one that has six steel balls hanging by strings and when you swing the first one the last one moves but the ones in between do not. Accordingly, instead of the usual textbook problems, which present a particular case and call for arriving at a numerical answer, instruction ought to focus on solving problems of understanding: What keeps a communications satellite in a stationary orbit? Why does an object thrown horizontally off a cliff follow the path it does? Would its path be any different in a vacuum? Why do light objects fall at the same speed as heavy ones? Can air resistance actually slow an object’s fall or only slow its acceleration? If gravity affects air, why don't all the air molecules end up flat on the earth? These may not be the most fruitful problems of understanding to pursue. Once they get into it, the students are likely to come up with better ones themselves. The purpose, however, is not to get answers to any particular questions and it is certainly not to exercise 'explaining' skills. Then what else could it be? Folk theory offers no significant third possibility.

The third possibility I am trying to illustrate here is developing a relationship between the learners and the object of understanding that is relevant to the kind of intelligent action that understanding is expected to support. Newton's laws are in one respect typical of the conceptual artifacts encountered in every school subject—historical concepts, literary concepts, social and economic concepts, and concepts from all areas of biology. The intelligent action that understanding them supports is primarily the pursuit of further understanding. This is not always the case. When we teach health concepts, for instance, the kind of action that understanding should support is intelligent eating, drinking, recreation, and so on. Accordingly—and this is not easy to do in school—the students' relationship to these concepts ought to develop within the context of lifestyle decisions and not be mainly theoretical. Learning your way around in the domain of health knowledge means learning to make lifestyle decisions that are appropriately constrained by this knowledge. The obvious way to do this—by analogy with learning your way around in a large corporate headquarters, for instance—is by pursuing varied goals that bring you into contact with the resources and constraints.

With health knowledge, as with every other domain of practical knowledge, however, there is always a secondary purpose in acquiring understanding, and that is to support further understanding. New health-related information comes with the daily news and poses new challenges to understanding. But even without new information, there is need for understanding to go on developing. A healthy relationship is a growing relationship, and this applies as much to our relationships with bodies of knowledge as to our relationships with other people. For students in a technical college, Newton's laws may be tools for the practical execution of their work, and so it makes sense that understanding should develop around applications rather than mainly around explanation, as I am proposing it should do at school. But it will be a poor education if the students' understanding of Newton's laws is so narrowly technical that it does not help them understand the world better and master new tools.

Thus, solving problems of understanding is not just one among a number of approaches to teaching for understanding. It ought to be part of every approach, and for most school subjects it ought to be central. That does not necessarily mean an inquiry approach such as I have sketched here. Socratic questioning is another way to pursue problems of understanding (Collins & Stevens, 1982). There are also inquiry approaches to understanding that are more structured and that give the teacher a more leading role (Brown & Campione, 1994; Lampert, 1990; Minstrell, 1989). Finally, a capable teacher who cares about understanding can engage students in thinking through problems of understanding using the old-fashioned methods of lecture, demonstration, and teacher-led discussion. My high school physics teacher did that. There was no lab, the textbook was dreary, and he was not an eloquent or inspiring lecturer. But he knew his physics, invited questions and took them seriously, and had an obvious regard for the worth of his subject matter. The result was that every day I left his class feeling that I understood the world a little better than when I walked in. Most of the problems of understanding, the things that didn't quite make sense and that led to deeper understanding, only occurred to me years later; but at least I had some understanding to build on, not just an assortment of half-remembered facts, formulas, and experiments.

In a better designed educational program, those problems would have had a better chance to arise during the course. But teachers back then didn't understand understanding in a way that would have suggested a different approach. To them, teaching for understanding meant making things as clear as possible, going slowly, and encouraging students to ask questions whenever something was not clear. To many teachers that is what it still means. To others it means letting students experiment and helping them arrive at scientific principles themselves. That, too, misses the point.

---

[1] The *Taxonomy* was actually the work of a blue-ribbon committee sponsored by the American Educational Research Association and chaired by Benjamin Bloom. Officially titled the *Taxonomy of Educational Objectives (Handbook I: Cognitive Domain)*, it was one of three taxonomies commissioned by the AERA, the others covering the affective and motor domains. These others never achieved the eminence of the taxonomy for the cognitive domain,

however, probably because achievement testing in these areas never got much backing.

[2] The *Taxonomy* draws on material from *The Measurement of Understanding* (Brownell, 1946), in which items spanning all the *Taxonomy*'s levels are related to measuring understanding. The fact that the *Taxonomy* authors refer to the Comprehension level as representing "the lowest level of understanding" (Bloom, 1956, p. 204) more explicitly indicates that they saw the higher levels as also representing understanding.

[3] Considerable opposition to use of the word 'misconception' has arisen among educators. Preferred terms include 'alternative conceptions' and 'experience-based conceptions.' Partly this reflects the same euphemistic etiquette that rejects 'blind' and 'visually handicapped' in favor of 'differently sighted.' But there is also a philosophical basis that I believe to represent a misconception in itself. Calling something a misconception is thought to imply a positivist epistemology; in order to call something a misconception, this reasoning goes, you have to believe there exists a true conception and perhaps even that you know what it is. That is nonsense. You can be perfectly confident that a certain key will not fit the lock without persuming to know what key will fit the lock or even whether such a key exists. If you want to be cautious, you should say "I *think* such-and-such is a misconception," just as you might hedge any other factual assertion you make. There is nothing *especially* risky or presumptuous about declaring something a misconception—compared, say, to declaring that it is time for lunch. Not only is the epistemology askew but so is the tender-mindedness. There is nothing shameful about having misconceptions. The best minds in history not only had them but articulated them brilliantly—and, of course, some of what brilliant people are articulating today will someday be judged to be misconception. One of the most helpful kinds of criticism you can ever receive is to have it pointed out that you are harboring a misconception. That can make a big difference to your understanding, compared to the little differences that accrue from the ordinary uptake of facts and ideas. I think misconceptions should be openly discussed with students. Presented in the right spirit, they are fascinating. We spoil the fun and misdirect learning when we call them things like 'alternative conceptions.'

[4] The correspondence view of understanding parallels the correspondence view of truth, which is a version of positivism—the notion that truth is a correspondence between what is in the mind and an external reality.

[5] *Pace* Homer Simpson, who said, "Just because I'm not interested doesn't mean I don't understand."

[6] An exception would occur if, for example, a discussion concerned with understanding some commonplace occurence took a bizarre turn that led us to question the mental state of the speaker. But in such a case, our task would have shifted from understanding the subject of discussion to understanding the discussant.

[7] Willie's analysis is that, in trying to keep the file level, people press down with equal force on both ends. But as the file moves across the blade, the leverage changes, so that the effect of equal force is that the front end goes down. To keep the file level, one must hold the file lightly and change the balance of pressure between the two hands as the file moves, so as to compensate for the change in leverage.

[8] Note that although the first belief is a reasonable inference based on experience, the second is not. Naive beliefs generally show this mixture of commonsense generalization and fanciful idealization.

[9] See "Translating formal knowledge into informal knowledge and skill" in Bereiter & Scardamalia, 1993, pp. 65-72.

[10] In this passage, you may observe, Bunge out-Poppers Popper in treating theories as real things having properties of their own quite distinct from what is in the heads of people who know them. The "infinite set of propositions closed under deduction" could not exist in any individual mind or indeed in all the minds of humanity put together. It is, as Bunge says, the property of an "ideal object." Yet, in another part of his magnum opus, Bunge ridicules Popper for advancing this very idea, archly asking what would happen to one of Popper's World 3 objects if all the people who knew of it had their heads chopped off. This goes to show how much trouble even accomplished philosophers have in facing up to an idea that is implicit in everything they do.

[11] If they are dropped at the same moment, then of course the one dropped from higher up will never overtake the other. But if there is a delay, so that the lower one is realeased when the upper one is part-way to it, then you may see the upper one zoom past the lower.