

Chapter 3

Knowledge Outside the Mind

The history of philosophy shows us very clearly that the full determination of a concept is very rarely the work of that thinker who first introduced that concept. For a philosophical concept is, generally speaking, rather a problem than the solution of a problem—and the full significance of this problem cannot be understood so long as it is still in its first implicit state. It must become explicit in order to be comprehended in its true meaning, and this transition from an implicit to an explicit state is the work of the future.

—Ernst Cassirer, *An Essay on Mind*
(1944)

One sign that we are ill-prepared for the knowledge age is that we do not have enough words. Modern languages have a vast number of words for describing human traits and feelings, for instance—far more words than most people can find a use for. And our vocabulary of technical terms expands as science and technology require. But we find words like ‘learning,’ ‘information,’ ‘teaching,’ and ‘knowledge’ each holding down two or more jobs, not doing any of them very well, and getting in each other’s way to boot.

There is an obvious reason for this discrepancy. People have probably been discussing each other’s traits and feelings since the dawn of civilization, whereas the need to talk about things like learning, information, and knowledge is a recent development, coinciding roughly with the institutionalization of these concerns. Within their own groups, educators, philosophers, information technologists, cognitive psychologists, and others may use such words with fair confidence that they know what their fellows are talking about. But now there is a need for these groups to talk to one another, and for this the old implicit meanings cannot be trusted.

The problem is not one that can be solved by getting people to agree on what each term means, however. That would leave some concepts with no words to represent them. Furthermore, it is not simply a matter of people being clear about what they mean; it is a matter of meaning the right things. That is, we need concepts, suitably labeled, that help rather than hinder the formulating and solving of problems and (I can’t help using one of the terms already) the advancement of knowledge.

When people are speaking casually about knowledge, they make fairly good sense. It is only when they are trying to be careful or systematic that they

get into tangles. In casual talk, people blithely treat knowledge as having a dual character. On one hand, they treat it as something real in itself. They talk about knowledge being contained in books, about the state of knowledge in a field, about the advance of knowledge, about the sharing and dissemination of knowledge, and so on. On the other hand, they treat it in the way that was discussed in the preceding chapter, as content in individual minds. That is how people talk about acquiring knowledge, about people having much or little knowledge of a subject, and so on. When they are writing articles for prestigious journals, however, people seem to become uncomfortable with this dualism and so they avoid talking about knowledge as having a life of its own and stick to the things-in-the-mind conception of knowledge, believing that to be the more defensible one.

If forced to choose, people dealing with social or economic issues would probably be better off making the opposite choice, treating knowledge as stuff out there in cyberspace and ignoring the psychological aspect. That is what epistemologists and historians of knowledge have generally done. But ignoring the human element is always dangerous, and especially so when the point of deliberations is to reform human institutions such as industries, governments, public media, and schools. Thus we need either to face up to the dualism or to find a satisfactory way of resolving it.

In this chapter I will discuss two approaches to resolving the dualism. One way comes from what is variously known as sociocultural theory, situated cognition, or situativity theory. This approach avoids both sides of the dualism: Knowledge does not exist either in a world of its own or as stuff in individual minds but is an aspect of cultural practice. Knowledge is not either produced or acquired but is *constituted* in communities of practice and *embodied* in the tools of such practice. This approach has proved illuminating in many respects, but in my view it is most applicable to traditional, slowly evolving communities of practice and is not an approach that will carry us very far into the Knowledge Age.

The alternative approach that I shall develop centers around the idea of *conceptual artifacts*, which are human constructions like other artifacts, except that they are immaterial and, instead of serving purposes such as cutting, lifting, and inscribing, they serve purposes such as explaining and predicting. These conceptual artifacts, in turn, become part of the vast array of things that we can become knowledgeable about. Thus, both of the commonsense ways of conceiving of knowledge are maintained but in a more coherent and—I hope to show—usable way.

Knowledge Embedded in Cultural Practice

‘Sociocultural’ describes a variety of theoretical positions that attribute an inherently social character to knowledge and learning. Some versions remain implicitly based on folk theory of mind, extending it to the social plane with concepts such as *intersubjectivity* (Bruner, 1990; Olson & Bruner, 1996). Others, however--those going by the names of ‘activity theory and ‘situated

cognition' (Davydov & Radzikhovskii, 1985; Suchman, 1987)--explicitly reject the mind-as-container metaphor.

Where is knowledge if it isn't contained in individual minds? The kind of answer coming from activity and situated cognition theorists runs along the following lines: Knowledge is not lodged in any physical or metaphorical organ. Rather, knowledge inheres in social practices and in the tools and artifacts used in those practices. Learning, as it pertains to individuals, is a process of moving from peripheral to full participation in cultural practices (Lave & Wenger, 1991). At another level, learning is the evolution of those practices.

This 'situated' view of cognition is very similar in form to the connectionist view of mind, although on a different plane. In both, knowledge is regarded as *distributed*. This does not mean merely that it is spread around, a bit here and a bit there, like the pieces of a treasure map. Knowledge does not consist of little bits. It has been said of connectionist models that "all the knowledge is in the connections" (Rumelhart, 1989, p. 135). Similarly, we may say of situated cognition that all the knowledge is in the relationships—relationships among the people engaged in an activity, the tools they use, and the material conditions of the environment in which the activity takes place. This does not mean that connectionism and situated cognition are essentially the same theory or even that they are compatible. It means that they employ the same abstract model (Ohlsson, 1993), so that if you can understand one you can probably understand the other, and if you cannot understand one you probably cannot understand the other.

Some of the most compelling arguments against the mind-as-container metaphor have come from the situated cognition literature (e.g, Agre, 1997; Lave & Wenger, 1991; Suchman, 1987). An important body of research motivated by the situated view of cognition has examined how learning goes on in nonformal settings, such as workplaces (Rogoff & Lave, 1984; Suchman, 1987). This research has revealed an important irony. Learning, as it goes on among 'just plain folks,' as Jean Lave calls them, is not at all congruent with folk theory. People do not acquire items of mental content, which they then take out on occasion and use. Instead, they learn how to take part in what is going on and how to function intelligently in the work environment. Folk theory of mind much better fits what is supposed to go on in schools, as far as the official curriculum is concerned, than what goes on in ordinary life. This is understandable, however, inasmuch as schools are contrived institutions designed according to what people believe about knowledge and learning (Olson and Bruner, 1996).

For understanding the nature of knowledge and its uses, one of the most instructive studies is one by Sylvia Scribner (1984) on dairy workers who load trucks or deliver dairy products to stores. The workers have to deal continually with quantitative relationships—assembling cases according to the number of pints of this and quarts of that in an order, computing costs, and so on; but little use of school mathematics was evident. Instead, the workers had clever ways of using the fixed row-by-column layouts of packing

cases to minimize mental calculation (and also to minimize bending and lifting). An observer could abstract the mathematical knowledge implicit in such activity and give it a formal description (which the workers might very likely have been unable to understand). But such formalized knowledge would have little value, because it would only apply to packing cases and items of certain dimensions. In effect, it would apply only to the observed situation and others much like it, and workers in those situations would not need such formalizations.

There is no question that a great deal of what we know is of this highly situated kind. One of my most routinized tasks is fixing two cups of cafe latte each morning to get my wife and me started on our day. General instructions for accomplishing this task—the kind that come with the espresso machine—cover only a small part of what actually goes on. Yet a more detailed description would for the most part be useless to people in other situations. The order in which things are done is adapted to the physical layout of the kitchen. That is why I discard the old coffee grounds on the way to picking up cups for the new batch. My coffee-making behavior is not rigidly fixed, however. There are all sorts of minor variations, some adapting to irregularities such as there being no clean coffee cups or the compost pail not being there to receive the coffee grounds, some having no apparent rationale at all; but they are all variations within the set of relationships connecting myself, the various tools of the task, and features of the kitchen environment. How long I run the espresso machine is not determined by time (which varies, depending on how tightly the coffee powder is packed) or by the measured volume of liquid that comes out. I just know what the right amount of coffee looks like when I peer into the cup, and this of course varies with the cup. One of the more eccentric features of my coffee-making practice is that I do push-ups while the espresso machine is running. This has nothing to do with making the coffee, but it nicely illustrates a point situativity researchers often make: That to understand situated action, we must not limit attention to discrete tasks and goals but must keep in mind that what people are mainly doing is living their lives, and this commonly means having multiple goals simultaneously operative. Coffee-making provides the right place and the right amount of otherwise uncommitted time for doing the kind of modified push-ups I do, and making them part of the coffee routine makes it much less likely that I will neglect to do them.

The distributed nature of situated knowledge becomes even more evident when we look at mass behavior, such as that of people walking along a crowded sidewalk. How everyone manages to keep moving at a fairly brisk pace without collisions, usually without even brushing up against one another, is quite amazing—although it is worth noting that fish are even better at getting about in dense traffic than we are. This is not just a matter of individual skills. There are evolved social practices, such as veering to the right, which make it possible. This becomes evident when the crowd contains people from different cultures with different right-left biases and no doubt

different signals, different body language. Then one sees people zigging when they should be zagging and the smooth flow of traffic being disrupted.

Such homespun examples are not far off from the kinds of activity situated cognition researchers have studied most closely. Situated cognition research has dealt mainly with traditional crafts, everyday activities, and low-status occupations (reviewed in Lave & Wenger, 1991). This raises what is for me the most serious question about situated cognition: How relevant is it to the modern world—to nontraditional occupations, to formal education, to rapidly advancing technologies that draw on abstract knowledge—in short, to the Knowledge Age?

On one hand, there is no doubt that the activities of workers in a software development shop or a biomedical research laboratory or an investment firm show the same kinds of situatedness as my coffee-making behavior or the behavior of shoppers in the Holiday Season rush. To make short work of it, let us simply acknowledge that all our behavior is situated, in the same way that all fish behavior and all squirrel behavior is situated (Bereiter, 1997). Should not this be the starting point rather than the ending point of our effort to understand human knowledge?

Scribner's account of dairy work included an interesting subplot. During the summer, high school students were hired to work along with the regular truck loaders. These novices, it was observed, did make use of school mathematics in figuring out how many containers of one product and another were required to fulfill an order. The point Scribner made was that the students' procedures were slower and more prone to error than those of the experienced workers. But looked at another way, school mathematics enabled the students to start right in doing the job without having to go through an apprenticeship or a long period of bumbling around.

School arithmetic has the advantage that it works in an unlimited range of situations, perhaps never as well as knowledge evolved in and adapted to each particular situation, but usually much better than quantity-manipulating knowledge evolved in one situation being transferred to another (Saxe, 1991). School mathematics seems to represent a different kind of knowledge, a kind that is in some sense less situated. That is, it is not knowledge embedded in some community of practice but rather is knowledge there for the taking, by anyone who has access to it and who can make something of it. This, I believe, is the knowledge we are talking about when we talk about the Knowledge Age and knowledge-creating companies. Without denying the importance of situated knowledge, we need to work out a clearer understanding of this other kind of knowledge.

Knowledge as Product

The most fundamental confusion that I see standing in the way of understanding what the Knowledge Age is about is a confusion between the knowledge used in productive work and the knowledge that is the object of such work. In traditional manufacturing and service industries the distinction is not important, because only the first kind of knowledge is

involved. If a company is in the business of making ice cream or cleaning furnaces, then, as situated cognition theorists would put it, the company comprises a community of practice in which is constituted a knowledge of ice cream making or furnace cleaning. In the course of productive work, new knowledge is generated, but it is further knowledge of making ice cream or cleaning furnaces. Thus the new knowledge remains constituted within the community of practice. The knowledge that comes out of the work remains part of the work. The product of the work is something entirely different.

But let us now consider a company whose business is producing knowledge. An opinion research organization will serve as an example. Such an organization carries out polls or interview studies to find out what the public or some segment of it thinks about one thing or another. Its clientele may include political parties, manufacturers of consumer goods, or anyone who is prepared to put up money to gain some particular knowledge of current public opinion. The knowledge that goes into opinion research is quite various. It includes formal knowledge of sampling statistics but also a great deal of informal knowledge having to do with how questions are phrased and posed and with the interpretation of results. Much of this is situated knowledge along the lines discussed in the preceding section. However, this is not the knowledge that is delivered to customers. The knowledge delivered to them is knowledge about the public, its opinions, attitudes, and beliefs relative to some object or issue of interest. The work of the opinion research organization thus generates knowledge of two kinds. One kind is knowledge that is inseparable from the work itself, because it is constituted by the evolving skills and practices of the people who compose the organization. The other kind is knowledge that is of no value unless it can be separated from the community that produced it. It is knowledge that the company produces and sells, much as a bootmaker produces and sells boots.

'Intellectual property' is a legal term applied to this merchantable kind of knowledge. The concept of intellectual property has limited value in characterizing knowledge, but it gives us a foothold. For intellectual property implies something that has existence in its own right. Otherwise, how could there be a dispute about who owns it? But intellectual property is not something material. It is not the bound report that the opinion research company delivers to the client, it is the knowledge contained in that report.

By expanding the concept of intellectual property beyond its legal boundaries, we can identify more broadly a type of knowledge whose distinguishing features are:

- (1) that it consists of immaterial things, not physical objects, and
- (2) that these things have an objective, out-in-the-world existence—they are not the contents of people's minds.

This broader category of intellectual property is already widely recognized in science and scholarship. Professional ethics typically dictate acknowledging the source of ideas, even when they have not been published, and of recognizing priority when several people have independently proposed the

same idea. But professional ethics stop short of giving credit for mental content. Your claim to have thought of a proof for Fermat's last theorem years ago, although you never publicized it or even wrote it down, will not get you much sympathy, let alone honor. As some writer I can no longer identify remarked, every good idea has already been thought of by somebody else, who did not appreciate its significance.

But it is not just making the private public or making the implicit explicit that is at issue here. Folk theory and situated cognition theory can handle those aspects of knowledge (cf. Nonaka & Takeuchi, 1995). What folk theory does not handle well is the status of knowledge once it is pruned loose from the mind of the individual knower. Correspondingly, situated cognition theory falls short when it comes to handling the status of knowledge once it is pruned loose from practice. What happens with such prying loose, I believe, is a level shift that is as radical as the shift from a barter to a money economy. Indeed, the shift is very similar to that from barter to money and it is what makes a knowledge-based economy a realistic possibility rather than merely a figure of speech.

Money, in fact, provides an instructive analogy for understanding this idea of knowledge as immaterial but real. We grow up thinking of money as material stuff—coins and bills. But as the movement advances toward a cashless economy, money becomes increasingly abstract. Getting paid for your work once meant being handed an envelope with cash in it; then it became receiving a check; now it is likely to mean a transfer between your employer's bank account and yours, with no material object of any kind changing hands. But money is no less real for being abstract, a fact that people with wallets full of credit cards reportedly sometimes fail to appreciate. And in a deeper sense money has always been abstract. That is what makes a dime a dime—a tenth of a dollar—and not just a small metal disk. An argument sometimes brought against treating knowledge as real is that it would cease to exist if all the people who know of it were to die. But the same is true of money. Roman coins are not money any longer, they are just coins. The money vanished with the people for whom those coins were its embodiments.

Ideas, facts, theories, algorithms, designs, problem formulations and problem solutions—these are all abstract. Commonly they are represented in some material form, such as a book, just as money is commonly represented in the material form of coins and banknotes. But the representation or concrete embodiment is not the knowledge. When we argue about a theory or design we are not arguing about a piece of paper or about the particular words or lines on the piece of paper. We are arguing about the abstract knowledge object of which those words or lines are but one possible representation. Just as electronic media may be getting rid of coins and banknotes, they may be getting rid of fixed representations of knowledge in tangible media. This may have important consequences, but it will not alter the basic nature of knowledge.

The Idea of Conceptual Artifacts

Philosophers have long struggled with the problem of how to characterize a kind of knowledge that is, so to speak, “out there,” removed from individual minds and particular contexts. Plato’s effort in this vein is the most famous but it is quite tangential to our present concerns. Plato posited a realm of pure ideas separate from human knowledge. Human knowledge represents our imperfect though improvable grasp of those pure ideas. This line of thought was refined by Bolzano (1972) in his conception of “propositions in themselves,” a domain that consists not only of all the propositions that have been uttered but of all the propositions, true and false, that might ever be uttered. That idea has been further refined by Erich Reck (in press) into the notion of “conceptual possibilities.” Actual human thought, on this view, attends to or develops or comes to believe in certain conceptual possibilities, which then form part of human knowledge. Reck’s formulation comes closer to what I believe we are looking for as a view of knowledge for the Knowledge Age. It is a view of knowledge as the active realization of certain possibilities inherent in the way the world is.

One reason none of these formulations quite does the job, however, is that they were not created with practical purposes in mind. They were created to deal with traditional epistemological issues, in particular the issue of whether a proposition can be true or false irrespective of what people happen to think. Common sense would say yes, but it has proved difficult to mount a coherent argument to that effect. These postulations of knowledge existing independent of people’s minds are all parts of efforts to do so. Interesting as the issue may be, it does not need to be decided for present purposes. We are not trying to establish what knowledge *really* is. Knowledge, like mind, can be pretty much what we choose to make it. We can tolerate some logical incoherence, uncertain boundaries, incompleteness, provided we are able to get on better with our work.

A philosopher whose efforts took on this practical character was Karl Popper. His distinction between three worlds—the physical world, the subjective world, and the world of ideas—has technical difficulties, as he recognized himself, but it is extremely useful for making distinctions that count in the affairs of education and knowledge work. Moreover, Popper stressed three points about his third world that seem to me crucial for getting straight about the role of knowledge in the knowledge age. First, the contents of World 3 are entirely human creations; they are not like Plato’s pure ideas or the universe of possible propositions of Bolzano and Frege. Second, these human creations, like other human creations, are fallible but improvable. Thus knowledge becomes, in Popper’s view, something you can work with. Finally, and most controversially, these human creations take on a life of their own, independent of their creators. They can be found to have characteristics, virtues and faults, implications and applications, that their creators could not have foreseen.

Popper made the unfortunate choice of calling this kind of knowledge “objective,” a term that has become a hot button for antifoundationalists and postmodernists. His essential points can be better captured, and

misunderstandings better avoided, if we adopt instead the term *conceptual artifacts*. I will elaborate on the meaning of the term later. For a first pass, *conceptual* may be understood to refer to discussible ideas, ranging from theories, designs, and plans down to concepts, like unemployment and gravity. *Artifact* conveys that these are human creations and that they are created to some purpose. However, being conceptual, they are not concrete artifacts, either, as are books, statues, and fire hydrants.

Conceptual artifacts share many of the characteristics of material artifacts. Consider the concept of natural selection and how it compares to a material artifact like an automobile:

- They both have origins and histories
- They can be described
- They can be compared with other artifacts of their type
- They may be valued or judged worthless
- They have varied uses
- They may be modified and improved upon
- They may be a subject of discussion
- New attributes, uses, or defects may be discovered that were not foreseen when they were created
- People differ in how well they understand them and in how skillful they are in using them

Especially in practical affairs, conceptual artifacts often have a close relationship to material things, but the distinction remains important. Consider the relationship between the automobile, a material artifact, and the design for the automobile, which is a conceptual one. Suppose your car begins making alarming noises whenever you turn a sharp corner. Your mechanic will regard this as a problem with the physical artifact and will treat it accordingly. But suppose it becomes known that the same problem is occurring regularly with cars of the same model as yours. Engineers trying to determine the cause will consider the possibility of defective materials or manufacturing procedures, which are still material problems, but they will also consider the possibility of a design defect, and this is something wrong with the conceptual object.

Much of what is meant by the shift from an industrial to a knowledge society is that increasing amounts of work are being done on conceptual objects rather than on the physical objects to which they are related. A widely reported story is that of the Boeing 767 aircraft, for which no physical models were built. Instead all the testing of the design was done computationally. Thomas Stewart's book, *Intellectual Capital* (1997), provides many more examples of knowledge taking over the role of physical artifacts.

Like most writers of futuristic business literature, Stewart is quite loose about what he treats as knowledge. It is almost anything of value that is intangible, ranging from customer relations and worker morale to trade secrets. This looseness is probably justified at present. With people just starting to appreciate the full significance of knowledge in practical affairs, we should not be in a hurry to put limits on it. The traditional philosophers'

definition of knowledge as true or warranted belief is obviously much too restrictive to cover what figures as knowledge in a knowledge-based economy, but once you abandon that definition it is not clear anymore what the limits should be. My purpose here is not to pin down what knowledge is but to argue for distinguishing knowledge in the form of identifiable conceptual artifacts from kinds of knowledge that are implicit or embedded in individual minds, in social practices, or in physical tools. Those other kinds of knowledge are important, too, and I will expand on some of them in Chapter 5. But they do not have the distinctive properties that conceptual artifacts have, properties which make them amenable to deliberate programs of knowledge creation and improvement.

Some intangibles are less intangible than others, we might say. Surely theories and plans are less intangible than intuitions and hopes. There might be no theories without intuitions or plans without hopes, but theories and plans represent a step forward that is more than a gain in clarity. Whereas intuitions and hopes are mental states that always remain to some extent private, theories and plans have the thing-like characteristics of artifacts. This makes it possible to treat them as objects of study and discussion. We can criticize them and propose improvements, experiment with applications—in short, make them part of the collective human enterprise in ways that mental states never can be.

What fundamentally sets conceptual artifacts apart, however, is not that they are immaterial. There are other immaterial artifacts, such as musical compositions, which are not conceptual. What sets conceptual artifacts apart is that they can be treated as knowledge. But what does that mean?

Because knowledge can play a variety of roles in human affairs, there is not one universal test for distinguishing conceptual artifacts from nonconceptual ones. However, there are two tests that between them cover a wide range of cases:

1. Can it serve as a tool?
2. Can it be used to rationalize behavior?

There is a third test which, in my first musings on this subject, seemed preeminent: Can there be profitable argument about whether it is true or false? However, that test gets so quickly mired in all of the bogs of epistemology that I prefer to skip it. Furthermore, it applies to only one kind of conceptual artifact, albeit a very important kind—what I will discuss later in this chapter as ‘assertive artifacts,’ with theories being prime examples. The other tests apply more broadly. Theories can be used as tools, but so can concepts, problem formulations, systems such as Euclidian geometry, histories, recipes, and aphorisms. But sonatas and poems cannot, at least not without violence to their normal cultural roles. Similarly, any of the conceptual artifacts mentioned can be used after the fact to justify actions, to defend them as reasonable. This does not mean that

conceptual artifacts cause behavior. It means that they come into play when we try to construct a rationale for our own or someone else's behavior. The statement, "Women are fickle," could be used in constructing a rationale for philandering. Verdi's aria, *La Donna è Mobile*, could not, even though its title expresses the same thought. The statement is a conceptual artifact; the aria is a musical one. The statement can be discussed as true or false, criticized as an over-generalization, modified so as to be more accurate. Verdi's aria cannot be treated in those ways, unless one ceases to regard it as an aria and makes it out to be a sort of theory set to music. There are ways in which the conceptual artifact and the musical artifact may be treated the same. They, could, for instance, both be treated as expressions of a masculine tendency to trivialize female emotionality. But in doing so we move from Popper's World 3 to World 2. We are no longer talking about ideas as putative knowledge but rather about the subjective states that give rise to certain ideas.

Let us pursue further what it means for a conceptual artifact to have a role in rationalizing behavior. If we ask people to draw the path of a ball thrown horizontally off a cliff and to explain their reasoning, we shall probably find that only a minority of them draw the curved downward path that can easily be observed by simulating the event with an eraser thrust off a table. Some will draw a path that goes horizontally for a while and then drops precipitously. Their explanations may refer to the ball slowing down and gravity "taking over." Science educators will explain this kind of response by saying that the students are applying the preNewtonian concept of *impetus*. Of course, the students are unlikely to have used that term and it is questionable whether anything went on in their heads that corresponds to recalling and applying such a concept. But *impetus* is clearly a concept that can be used to rationalize the observed behavior; i.e., if it were the case that propelled objects acquire an impetus which they gradually lose then these students' responses would be perfectly rational. Thus, *impetus* qualifies as a conceptual artifact. Other students will draw a straight angular path, and their explanations will suggest a role for the concept of *resultant force*. But again we cannot assume that students actually recall and judge the applicability of that concept. Possibly all that came to mind was the recollection of diagrams in a physics book that showed two arrows converging on a point and another arrow leading off at a peculiar angle. The diagram happened to be a depiction of resultant force, but the student may never have grasped that point or may have forgotten it. Even among those who draw a correct path, however, you are likely to find that their explanations refer to such concepts as *air resistance* rather than *acceleration due to gravity*. *Impetus*, *resultant force*, *air resistance*, and *acceleration due to gravity* are all conceptual artifacts. They may be combined with others to produce a theory or explanation of propelled motion, which is also a conceptual artifact.

Note the status that has been given to conceptual artifacts in the preceding examples. They are treated as things that exist and which a person may or may not know of or think about. When we give these artifacts a role in rationalizing behavior, this is only by way of distinguishing them from

nonconceptual artifacts on one hand (listening to *La Donna è Mobile* might influence a person's attitude toward women but it could not figure in a rationale for that attitude) and from mental states on the other. Mental states, we suppose, do cause behavior but they cannot constitute a rationale either.

Folk theory typically recognizes only two of Popper's three worlds. We can have knowledge of the physical world and we can have knowledge of our own (and to some extent other people's) subjective states. If we accept the idea of conceptual artifacts, however, then we should recognize that it is also possible to have knowledge of them. In the folk or two-worlds view, understanding acceleration due to gravity means understanding something about the physical world—understanding that the speed of a falling object increases in proportion to the time it has been falling. But then one must also understand that this is true only under ideal conditions and so it is only approximately true of the world we inhabit. In the three-worlds view, understanding acceleration due to gravity means understanding both the theoretical idea and how it applies to actual phenomena. Understanding the theoretical idea might include understanding that gravity is a force and knowing how force and time are related to acceleration. It could include understanding the difference between the concept of gravity and the older notion that things have an inherent tendency to fall downward, and it could include understanding the universality of gravity, its relation to tides, planetary orbits, and so on. Of course, it could include a great deal more, or understanding of the conceptual artifact could be spotty and inconsistent. For instance, one could think that gravity is universal and at the same time believe, as is sometimes stated in news reports, that there is no gravity in outer space.

Understanding acceleration due to gravity as a physical phenomenon might include knowing the gravitational constant on earth and being familiar with the factors that influence how closely falling bodies conform to theoretical values. Clearly, understanding the conceptual artifact and understanding the physical phenomena it applies to can be closely related, but they are not the same. You could understand Newton's laws well enough and yet have very naive ideas about how things work in the physical world. That, according to research on naive conceptions in physics, seems to characterize the typical university student (Clement, 1982). Or, you could have quite a good qualitative¹ understanding of forces and motion as they appear in

¹ Qualitative physics (De Kleer & Brown, 1985) refers to physics that can be expressed in terms such as "the more this the more that," "this causes that," and functions whose shape can be described but not their mathematical form. According to De Kleer and Brown, the expert has not only a grasp of formal laws but a very large amount of qualitative knowledge, which is particularized to concrete objects and phenomena. Thus, in Popper's terms, knowing formal laws would be knowledge of World 3; qualitative physics would be knowledge of World 1.

everyday phenomena and yet have little grasp of Newton's laws at a mathematical level.

In many situations, it will be impossible to make a clear distinction between knowing or thinking about the conceptual artifact and knowing or thinking about the material world the conceptual artifact applies to. For that reason, purists may believe the distinction should be abolished. If so, they ought also to argue against making a distinction between cars and trucks, now that manufacturers are producing a variety of vehicles that are not clearly one or the other. Motor vehicle offices still make the distinction and issue different license plates, based on use; that is how we ought to be treating the distinction at issue here.

Finally, however, it does not matter much whether we can cleanly separate knowledge of conceptual artifacts from knowledge of the world they pertain to. What matters is that we recognize conceptual artifacts as real things, recognize creating and improving them as real work, and recognize understanding them as real understanding. Failure to do so, I am trying to argue, means playing the game of life in the knowledge age with a short hand.

Without the idea of conceptual artifacts, knowledge creation becomes merely a glorifying figure of speech. Suppose that, through varied experience as a bombardier, a basketball player, and a hurler of stones off cliffs, I develop an excellent sense of trajectory. I have acquired knowledge but it is pointless to say that I have created knowledge. I have not produced a conceptual artifact. Suppose that, in order to impart my knowledge to others, I produce a training videotape that reveals in slow motion the trajectories of objects propelled under various conditions; or suppose that I am clever enough to construct a simulation device, a sort of Link trainer, that will enable people to acquire in short order the skill and intuitive knowledge that I possess. I have still not produced a conceptual artifact. I have not produced anything that would pass the test of serving to rationalize behavior. To do that, I would have to produce something like a theory of ballistics. People might then continue to admire my ability to predict where propelled objects will fall but they might say that my theory was wrong or too simple or too complicated or impossible to understand—or that it was brilliant and bound to put Newton's laws out of business. In producing the theory I have created knowledge, whereas in the other cases I may be said to have *acquired* knowledge (possibly through some internal constructive process) and created ways of instilling similar knowledge in others. Thus, creating a knowledge object is one thing whereas knowing and helping others to know is another. The first I shall call *knowledge building*, the second, *learning*. Everyone learns and in some fashion everyone teaches. But not everyone produces or works with or is even aware of conceptual artifacts.

Working in Three Worlds

One virtue of the idea of conceptual artifacts is that it offers us a straightforward way of defining knowledge work. If knowledge work does not

seem to need defining, consider how you would distinguish it from white collar work in general. You may find that the only way you can distinguish it is that it is somehow of a higher level than some white collar work and accordingly more highly paid. But that only raises the question of what is higher-level about it and why it is more highly paid. If you try to define knowledge work as work that requires or uses a great deal of knowledge, then you surely have to include brain surgeons and jewel thieves as knowledge workers, and the term quickly loses its point. Some people try to avail themselves of a distinction between information and knowledge. If you work with information you are a white collar flunky and if you work with knowledge you belong to the elite class of knowledge workers. But what is the difference between information and knowledge? About all you can say is that information is at a lower level than knowledge, which puts us back in the same loop as before. Once we have established a notion of conceptual artifacts, however, we can say quite simply that knowledge work is work with conceptual artifacts.

More specifically, knowledge work is work that creates or adds value to conceptual artifacts, in much the way that metalcrafting creates or adds value to metal artifacts. Many people may use a conceptual artifact, but that does not make them knowledge workers, anymore than driving a taxicab makes one an auto worker. Knowledge work seldom goes on in isolation from other kinds of work, however. Even scholars in the most rarified disciplines also do concrete things like writing papers and giving lectures, not to mention sharpening pencils and trekking to the library—all of which have something to do with knowledge work but are not themselves operations carried out on conceptual artifacts.

I think it helps in grasping the nature of work with conceptual artifacts to put it in a historical context, even though some of the history is rather speculative. World 1, we may suppose, exists for all animals whose nervous systems have some requisite degree of complexity. It is the world of practical action. To say that it exists for us and our nonhuman relatives is to say that our nervous systems have some way of representing it so that our actions exhibit a degree of intelligence. World 2, however, if not strictly confined to the human species, is surely not so widely distributed. There is evidence that apes have capacities to deceive one another on purpose and perform actions such as deliberate teaching, which seem to require a representation of one another's beliefs and intentions (Premack & Premack, 1996). But that does not mean they have any concept of an inner world of experience. Julian Jaynes (1976), using evidence from mythology, has argued in effect that World 2 did not come into existence for human beings until after the time of Homer. In the ancient myths there is nothing to suggest inner experience. Everything is overt. What we would interpret as thoughts and feelings appear there as voices of the gods or of the dead. Jaynes's radical idea is that that was not a mere literary convention but was how the ancients actually experienced cognitive events. Be that as it may, it is clear that the 'inside view' in literature developed very gradually over a span of centuries. Only with the

Romantic Movement, which straddled the 18th and 19th centuries, did people's inner lives come to be represented in the ways that we now regard as natural (see Oatley, 1992, especially Ch. 5). And even today the experiencing of an inner life appears to vary considerably as between introverts and extraverts, women and men, the educated and the uneducated. Talk therapies, which rely on exploration of World 2, are widely recognized as having little promise with people far out on the extraverted, masculine, uneducated dimensions. They are more like Achilles than they are like Anna Karenina. In practical terms, access to World 2 equips us to take initiative in improving our own knowledge and competence—to become 'active learners' in the current educational jargon—and to recognize our prejudices, resistances, and susceptibilities. Less obviously, but very importantly, it enables us to cultivate and make use of our feelings as a form of knowledge—what typically goes by the name of 'intuition' and what I will discuss in Chapter 5 as 'impressionistic knowledge.'

World 3 also shows evidence of gradual emergence, although it may appear at first glance that Plato and Aristotle had it all. They certainly operated consciously in a realm of conjectures, explanations, proofs, and arguments. The world of ideas was probably more immediate to them than it is to 90 percent of the world's present population. Still, there are differences. The most obvious is that they did not have the constructivist view that is essential to distinguishing World 3 from Worlds 1 and 2. Knowledge creation and the idea of conceptual artifacts would have made little sense to them. Truth existed and was to be discovered and apprehended, not constructed by human enterprise. Of course, that view persists today, but throughout most of recorded history it seems to have been universal. There is a famous quotation attributed to Newton. Whether apocryphal or not, it nicely conveys this premodern conception of knowledge:

I do not know what I may appear to the world, but to myself I seem to have been only a boy playing on the seashore, diverting myself in now and then finding a smoother pebble or a prettier shell than ordinary, whilst the great ocean of truth lay all undiscovered before me.

What I see here is not a conception of World 3 but rather belief in a mysterious and superior World 1 lying beyond the World 1 available to our senses. This is the essence of mysticism. It is evident in Plato's image of the cave, and it is there in the "great ocean of truth." In a fascinating essay titled "Newton, the Man," John Maynard Keynes (1956, p. 277) said,

Newton was not the first of the age of reason. He was the last of the magicians, the last of the Babylonians and Sumerians, the last great mind which looked out on the visible and intellectual world with the same eyes as those who began to build our intellectual inheritance rather less than 10,000 years ago.

Relying on Newton's private documents, Keynes inferred that for Newton the universe was a gigantic riddle set by God, with clues distributed (among the pebbles and shells, as it might be said) for the benefit of the solver. On this view, which certainly did not disappear with the Age of Reason, there is no

World 3 of conceptual artifacts. There is World 1, there is World 2, and then there is this ultra-World 1, the world of truth, which it is the thinker's job to puzzle out, using clues from the physical world.²

My hunch, supported by Alfred North Whitehead's sketch of the history of knowledge (1925/1948), is that this mystical view did not really give way until the Industrial Revolution began to join science with invention. Inventions are obviously human creations (except to some die-hards who insist on construing inventions as gifts from God or the work of the devil). But if the steam engine is a human creation, why not the differential calculus?³ And if the differential calculus, why not the kinetic theory of gases, and so on through all of theoretical knowledge? Where this leads is to what Whitehead called the greatest invention of the 19th century, the "method of invention":

the full self-conscious realization of the power of professionalism in knowledge in all its departments, and of the way to produce professionals, and of the methods by which abstract knowledge can be connected with technology, and of the boundless possibilities of technological advance. (Whitehead (1925/1948, p. 92)

This made possible what Whitehead called "disciplined progress," progress achieved through the deliberate and orderly pursuit of solutions to theoretical and technical problems. It also, I suggest, marked the emergence of World 3 as a part of the real and knowable world and as a world in which productive work could be done.

Popper, to the best of my knowledge, never went so far as to characterize World 3 as a workspace, as a sphere of activity, and so I am stretching the concept to give it that character. He did, however, come fairly close to it in saying (1972, p. 156), "I suggest that one day we will have to revolutionize psychology by looking at the human mind as an organ for interacting with the objects of the third world; for understanding them, contributing to them,

² There are other historical interpretations of Newton according to which he was the first of the moderns and a true Baconian (e.g., Bronowski & Mazlish, 1960). The important point for the present discussion, however, is the distinction between World 3 and ultra-World 1. Plato clearly believed in the latter and not the former. Whether Platonic mysticism carried as far as Newton and beyond is an interesting question, but what is more important is the failure of most philosophers and other thinkers to appreciate that there is nothing Platonic about World 3. It is a thoroughly modern and, I would say, post-postivist concept.

³ Whether mathematical reality is constructed or discovered remains a disputed issue (Tait, 1986). Belief that mathematical truths were revelations of the mind of God persisted into the 19th century (White, 1956). All I wish to argue is that with the rise of technology it became more plausible and natural to treat mathematics as a product of invention.

participating in them; and for bringing them to bear on the first world." Popper, thus, was describing what we would now call knowledge work.

An Ordinary Example of Conceptual Artifacts: Kitchen Recipes

Although World 3 is usually exemplified by such high-status cultural objects as theories and histories, a humbler example will serve better to illustrate basic ideas (mainly because it is less likely to stir up strong feelings about the nature of science or about cultural canons). Such a conceptual artifact is the kitchen recipe. Recipes have a life outside the minds of people who know them and outside their particular embodiments in printed form. We speak of recipes being handed down from generation to generation, undergoing modifications, splitting into various versions. A recipe may be tested, criticized, compared to others, kept secret or publicized, awarded prizes. It may become obsolete, or impossible to implement because of the unavailability of necessary ingredients; it may be adapted to new conditions, such as microwave cookery. New knowledge may reveal nutritional virtues or liabilities of the recipe that could not have been known to its originators.

What I shall have to say about recipes as conceptual artifacts is obvious and commonsensical, and so the conceptual distinctions I draw are likely to seem unnecessary and pedantic. The reason for making them is that when we move to higher-status conceptual artifacts, the kinds that make up the subject-matter of formal education, the same points apply but they are no longer part of conventional wisdom. A recipe might be regarded as a representation of something in a person's mind, but this is an inaccurate and unnecessary supposition. My mother-in-law, Bianca Contrucci, was noted for various succulent Italian dishes, which her offspring tried to induce her to prepare whenever they visited her or she them. From time to time one or another of the offspring got her to write out or dictate the recipe for one of these dishes. A comparison of different versions of the recipe for supposedly the same dish showed them to vary markedly on significant points. One explanation would be that she did not prepare the dish the same way every time. Another is that she did not measure ingredients or time things precisely, instead going by impression and various subtle clues that she herself was probably not aware of—so that the recipe represented only a guess about these variables. Folk theory, or its modern equivalents in cognitive psychology, would try to explain these phenomena as discrepancies between the recipe in the mind and the recipe as expressed in words. The gnocchi recipe in Bianca's mind, such theorizing might say, was complex and variable and, furthermore, not fully available to consciousness. So when it came to producing a verbal recipe, she had to improvise and fill in with inferences, possibly with remembered bits from other recipes.

Suppose we accept this last sentence as approximately right. The question I then ask is, what is the relevance of the preceding conjectures? What do we gain by dragging in the notion of a recipe (or a production system) in the mind? It is not that these conjectures are wrong or implausible (which is something we have already discussed). It is that they are on the wrong track. Making gnocchi and producing a recipe for gnocchi are two different tasks.

Some of the same knowledge and skills are involved, but succeeding at one does not ensure success at the other. Producing a recipe for gnocchi means creating a conceptual artifact that can serve as a guide to people in making the dish so that it will turn out a certain way. It also involves producing a representation of that recipe—normally a text of a certain conventional form (although it could be something more esoteric, such as a flow chart). To maintain clarity with regard to conceptual artifacts, we need to keep several things in mind at this point:

1. The written text is not, strictly speaking, the recipe. It is a representation of the recipe. If the written instructions are ambiguous, this does not make it a bad recipe; it means that we cannot be sure what the recipe is.

2. The recipe, thus, is an abstract object of which there may be various representations. When people criticize, argue about, or try to improve upon the recipe, it is this abstract object that they are attempting to deal with, not its representation.

3. This is not to imply that the recipe is first conceived as an abstract conceptual artifact and then translated into language or some other representational form. Everything we know about written composition implies that the mind does not work that way (Bereiter & Scardamalia, 1987). But in the case of recipes, theories, and other conceptual artifacts, we try to think about the content separately from its vehicle.

Suppose the publishers of a well-established recipe book decide it is time for a thorough-going revision. In addition to setting editors to work revising the text and illustrations, they also set a team of chefs to work trying out and making improvements in the recipes. The editors and the chefs would have to work in close communication, of course, but their jobs would be distinguishable. Both are functioning as knowledge workers, but they are working with different conceptual artifacts. The editors are working with texts. These are immaterial objects, we should note; the same text might appear in different editions in different typography and might also appear in electronic forms that readers could configure in any fonts and sizes they wished. In their normal work, chefs are not knowledge workers; they are food workers. They *use* knowledge, as all workers do, and they also *learn* or develop new knowledge in the course of their work. But their normal work is not producing or adding value to knowledge; it is producing and adding value to food. In the present instance, however, they have been seconded as knowledge workers. They are not getting paid for the food they produce but for the recipe evaluations and improvements they produce. If they take artistic license with the recipes and do various subtle things that cannot be translated into recipes, they may produce wonderful dishes but they will not be doing their job as knowledge workers.

Obvious as all this is, I believe it is something that neither folk theory of mind nor situated cognition theory can handle well. A folk psychological account of revising the recipe book would have knowledge in the minds of the chefs getting communicated to readers, with help from the editors. A

situated cognition account would have knowledge constituted in the practice of the chefs getting externalized in the form of written recipes as tools. Both kinds of accounts capture something of significance but both seem to miss the central point, which is that what these chefs are working with are not the contents of their minds or of their practice but are things that already existed before they came on to the scene and that they are charged to do something with—namely, the abstract objects known as recipes.

Common sense wants to reduce everything to Worlds 1 and 2. Baking a cake is activity in World 1, the material world. A recipe may therefore be thought of as a representation of that activity—like a movie or a scenario, but stripped down to essentials. Alternatively it may, as already suggested, be thought of as a representation of something in the mind of the chef—accordingly, a part of World 2. Common sense, not being greatly concerned with conceptual precision, may well regard a recipe as being both at the same time. This may be inelegant, but it does not actually lead to bad results in the kitchen. Cooking is a well-established traditional craft, and recipes play a limited but valuable role in contemporary versions of it. But what if, somehow, creating new and improved recipes became a national priority, with a premium on innovation, usability, and quality? What if the need exceeded the supply of top-level talent and the enterprise faced stiff competition from abroad? We might then find that commonsense muddling through was no longer sufficient. We might even find that we need a ‘theory of recipes,’ realizing that theories of nutrition, of cookery, and of food science, while all relevant, are not the same thing and do not quite fill the bill. A theory of recipes would be a theory of knowledge, albeit of a rather specific type of knowledge. It seems almost inevitable that a useful theory of recipes would have to make something like Popper’s distinction between three worlds: World 1, the world of actual cookery, where actual people do actual things with actual utensils and materials in actual kitchens; World 3, the world of recipes that are being created, tested, discussed, and modified; and World 2, the world of mental activity that mediates between the other two worlds.

Tools and Other Artifacts

The Russian activity theorists have made tool use a fundamental characteristic of human action. Animals and human infants act directly on objects in their environments. Mature human action, however, is always, as the activity theorists say, *mediated* by something external to both the actor and the object. The mediator may be a physical tool such as a hammer or shovel or it may be a nonphysical aid such as language or mathematics. Research by Vygotsky and Luria, from which activity theory grew, focused on young children learning to use language to regulate their own actions. There is also social mediation, as when parents help children with their homework. On this view, cognition, like every other form of human action, encompasses external mediation and therefore cannot be adequately characterized by internal mental operations performed on mental content.

On one construal, conceptual artifacts would count among the tools that mediate human action. On another construal, which seems to be the more accepted one among Western thinkers who have taken up action theory, there are no abstract tools such as Popper proposes. Tools are either material things, such as books and instruments, or they are practices, such as writing and mathematical notation. A good reason for preferring the first construal is that it allows for treating the same object—a theory or an algorithm, say—as a tool in one situation and as an object of inquiry or improvement in another.

One way or the other, a broadened conception of tools is useful for thinking about human intelligence. I have never understood why activity theorists are so vociferous about it, however, as if in proclaiming the ubiquity of tools in human activity they have cleared away a dense cloud. It seems to me that, like many another contribution to human thought, this one merely highlights something we already know and makes it more available for use but does not produce any rumblings of a paradigm shift.⁴

In advancing the idea of *distributed intelligence*, however, Roy Pea (1993) has carried activity theory a step or two beyond where common sense easily follows. “When I say that intelligence is distributed,” Pea explains (p. 50), “I mean that the resources that shape and enable activity are distributed in configuration across people, environments, and situations.” Distributed characteristics of any kind are a bewilderment to common sense, but distributed *intelligence* adds a further twist. In Pea’s usage, this means that tools as well as people embody intelligence—and he is not referring particularly to so-called intelligent machines. His examples are “jogger pulse meters, automatic street locators, currency exchange calculators, world-time clocks, and weight-loss calculators” (p. 53). Of such devices, he says:

These tools literally carry intelligence *in* them, in that they represent some individual’s or some community’s decision that the means thus offered should be reified, made stable, as a quasi-permanent form, for use by others. In terms of cultural history, these tools and the practices of the user community that accompany them are major carriers of patterns of previous reasoning.... They may now be used by a new generation with little or no awareness of the struggle that went into defining them and to adapting their characteristics to the tasks for which they were created. (Pea, 1993, p. 53)

Pea’s position is to be contrasted with the more ordinary view, which sees tools as parts of the real world that may be treated intelligently or not, according to one’s lights, the same as anything else. People who will use wood chisels as screwdrivers (and according to my experience the majority of human beings fall into this class) are behaving in a way that they may consider intelligent, inasmuch as they are exploiting affordances of the tool, but that anyone who owns a prized set of wood chisels will consider to fall

⁴By this I mean, following Thagard (1992), that the tools idea does not call for any drastic restructuring of the tree of concepts that we apply to human knowledge and intelligence.

below the normal range for human intelligence. A lot of intelligence goes into the design and manufacture of a good wood chisel, that we may agree. The chisel constitutes a partial solution to problems of removing precisely delimited portions of a substance (wood) which is nonhomogeneous in a particularly refractory way. (Wood seems to have a mind of its own, as woodworkers sometimes remark.) As Pea would have it, the intelligence that went into design of the wood chisel is now embodied in the tool itself and thus becomes a part of the intelligence that goes into solving the problems attendant on, for instance, cutting a mortise. You could cut a mortise (a squarish hole, for those unfamiliar with the term) using a sharp pocket knife, but that would take more intelligence on the part of the actor, to compensate for less of the relevant intelligence being embodied in the tool.

If our concern is with what it takes to produce a particular intended result (as in cutting a neat mortise or reducing the federal deficit), then Pea's idea of distributed intelligence provides a nice basis for considering the full range of possible contributors. If, for instance, the goal is to produce a sharp photograph of a spouting whale and tourist A has a state-of-the-art automatic camera whereas tourist B has a cheap camera with a few manual controls, we can set out necessary and sufficient conditions in ways that will assign most of the conditions to the camera in the case of tourist A and most of them to the person in the case of tourist B, although the total set of conditions is the same.

Notice, however, that in adopting this viewpoint, we are setting ourselves apart from the activity and are asking how a certain task may be accomplished, given certain tourists and certain cameras. From this standpoint, the human actor is also a tool. This is the standpoint of the manager, who decides whether a job should be done by a person or by a machine or by a low-paid worker and an expensive machine or by a higher-paid worker and a not-so-expensive machine. The decision may or may not involve humane considerations. That is not the point here. The point is that to view intelligence as distributed across people and artifacts is to view the people as artifacts. Such a viewpoint may be appropriate, may even be essential, where the design and management of work are concerned, but it is not the viewpoint of the worker or actor. Neither, I must emphasize, is it the viewpoint of the educator—unless, of course, the purpose of education is to produce human tools, as cynics of a Marxist persuasion like to allege.

Educators need a viewpoint on action that enables them to focus on what is learnable (by individuals or a collective) that will contribute to success. Intelligence comes into the picture in two forms. One is intelligence involved in the learning itself. This is how intelligence is usually considered in education—as a prerequisite to learning. The other is intelligence as the objective. This is the intelligence at issue here. We are interested in people's learning things that will contribute to intelligent action. In this regard, it seems to me that the commonsense approach is preferable to Pea's admittedly very appealing approach. We should attend to what it takes to use tools intelligently. Any intelligence that we attribute to the tool itself becomes simply part of what must be taken into account in using it intelligently. This

commonsense approach is the one Pea himself reverts to when he addresses educational implications:

We should reorient the educational emphasis from individual, tool-free cognition to facilitating individuals' responsive and novel uses of resources for creative and intelligent activity alone and in collaboration. Such an education would encourage and refine the natural tendency for people to continually re-create their own world as a scaffold for their activities. (Pea, 1993, p. 81)

There is, however, a more radical version of the idea of artifacts invested with intelligence. It comes from the Russian philosopher, Ilyenkov (Bakhurst, 1991), and it has something important and in a way disturbing to offer our conception of what it means to relate intelligently to things in the world. Underlying the whole treatment of tools and mediated action that we have been considering is the idea of artifacts—human creations that become part of the environment, part of the situation to which situated activity is attuned. Ilyenkov went farther, however, and argued in effect that to human beings *everything* is an artifact. Everything is invested with meaning, purpose, roles, affordances and resistances to our activity. We have no cognizance of things in the world except as they enter into our activity, and as we incorporate them, they become in a sense our own creation. We did not, of course, make the stars; but, depending on our knowledge and our purposes, we may know them as guides to navigation or as scenic backdrop to a nightscape. Indeed, it is only through their incorporation into human practice that things in the world can become objects of thought (Bakhurst, 1991, p. 201).

This sounds like idealism, but Ilyenkov was trying to find a place for ideas within the world of Marxist materialism, and so his idealism has a practical edge to it: To the person with a job to do or a problem to solve, everything becomes a tool. All of our past history of productive effort, all of our internalization of cultural practice—all of our practical intelligence, if you will—finds expression in the tool-like characteristics with which we invest the objects about us. This is the basis of human beings' boundless capacity to create technology. Everything, from a practical standpoint, already is technology; we just keep improving it.

Although there is a strong family resemblance between Pea's view and Ilyenkov's, Ilyenkov's is broader. The *artifactualization of the natural world*, as I shall call it,⁵ is an all-encompassing cultural phenomenon. We not only

⁵ Ilyenkov's terms, as translated by Bakhurst, include "idealization" and "spiritual culture"—terms that require considerable explanation in order to separate Ilyenkov's meaning from conventional meanings. To avoid such excursions, I have substituted terms that are already part of the present discussion. This means leaving out much of what is most distinctive in Ilenkov's ideas and no doubt distorting what remains, but I am not trying to give an exposition or critique (see Bakhurst for that) and have only taken pains to avoid making Ilyenkov out to say what I want him to say. In

invest objects with the properties of tools, we invest them with esthetic and moral properties as well. A tree is a radically different kind of artifact for the lumber baron, the romantic poet, and the ecowarrior (treating these as stereotypes). The first invests the tree with instrumental properties, the second with esthetic and the third with moral properties. Education's aim would be to see that all these are developed and harmonized. A too exclusive emphasis of tool-like properties leads toward dehumanization and alienation. If everything is a tool, then our fellow human beings are tools. When you are tying a bow in a tightly-drawn ribbon, there is no better tool than somebody else's finger. But carry the tool use of human beings much farther than that and you are into exploitation, manipulation, 'objectification' in its most negative sense, and ultimately slavery. But a too exclusive emphasis of esthetic properties can also be inhumane, and impractical to boot. And people who fasten on moral or political aspects to the exclusion of everything else tend to become fanatics and beyond reason. (Schools of education seem to have more than their share of such people.)

Regarding everything as an artifact, we can nevertheless see artifacts as arranged on a continuum. At one end are things that we artifactualize but have no control over: stars, weather phenomena, and so on. Then there are natural things that we do not create but that we can arrange or modify to serve our purposes: Trees can be planted so as to serve as windbreaks or providers of shade, for instance; animals can be trained to do work. Then come all the material artifacts that people produce, and finally abstract artifacts—pure knowledge objects. There are not many distinct gaps in this continuum, and technology is busy filling those in.

There are many different kinds of abstract artifacts, and nothing to stop people from creating more. Any nonmaterial creation will count: a song, a crossword puzzle, a paradox, a travelogue, a prediction, a historical account, an aphorism, a number system, a movie critique, a statute (but not a statue). All share the characteristics listed earlier, which may be boiled down to three fundamental ones: discussability, modifiability, and autonomy. Only the last is controversial, and to Popper it was the key:

The idea of autonomy is central to my theory of the third world: although the third world is a human product, a human creation, it creates in its turn, as do other animal products, its own domain of autonomy. (Popper, 1972, p. 118)

His prime example of autonomy is natural numbers: an ages-old human creation, yet mathematicians are continuing to find out new things about them and they keep impinging on our lives in new ways. Similar autonomy is exhibited by material products of human enterprise, such as aspirin.

The autonomy of physical things has given philosophers trouble: the idea that things exist and have properties independent of our cognition.

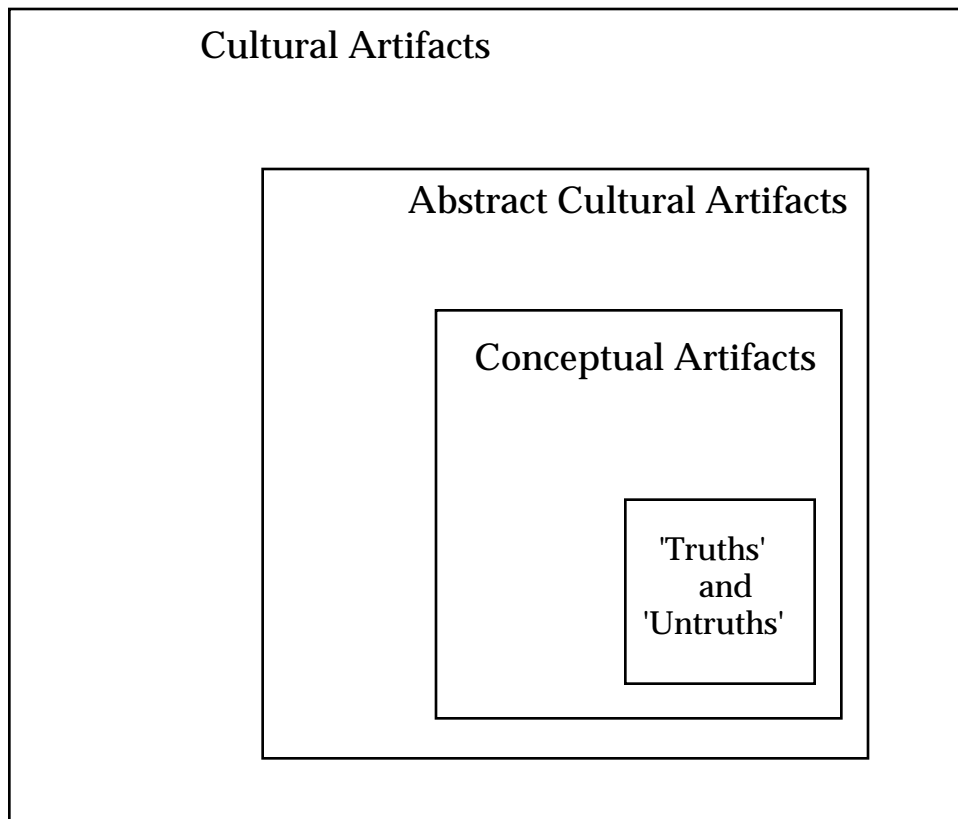
particular I am not trying to make Ilyenkov into an advocate for Popper's World 3, although some of his quoted utterances might be construed that way.

Understandably, then, the notion that mental constructions could enjoy similar autonomy has proved a great deal harder to swallow. Keeping in mind that our project here extends only to trying to produce a more *useful* theory of knowledge and mind, however, we can be pragmatic about such metaphysical issues. I do not see how discourse above the “Me Tarzan, you Jane” level is possible without treating ideas as things. Everybody does it. Radical behaviorists, while using laborious circumlocutions to avoid mentalistic language, never hesitated to treat theories, hypotheses, arguments, and the like as real things. Situated cognition theorists do it, even in the process of arguing against doing so. Jean Lave (1988, p. 91), for instance, criticizes “[t]he view that culture is the evolutionary accumulation of knowledge.” This *view* she says has *prevailed* through most of the last century. But a *problem* has *emerged* from this *view*. After considering two possible ways of dealing with the problem, she asserts that neither one “appears to offer a *solution*.” The italics, of course, are all mine and are meant to draw attention to Lave’s treating a *view* as a thing having identifiable properties and implications.

Certain kinds of intellectual products are treated as real, autonomous objects, without any hesitation or thought to alternatives. A sonata is an example. A sonata is not regarded as something in the composer’s mind—at least, not once it has been composed. But it is not the notes on paper, either. The notation is fallible. An editor may revise it to eliminate errors or to make it easier for a performer to follow. A sonata is not any particular performance or a totality of performances. It is the abstract thing that those are performances of. When people discuss a sonata they are discussing the abstract object, although what they have to say about it will undoubtedly be influenced by the performances they have heard. It takes some rather fancy theorizing, to raise any serious doubt about treating musical compositions as autonomous objects. This is because they do not *refer* to anything. It is when conceptual artifacts refer to things outside World 3 that troubles begin to arise.

Knowledge in Relation to Other Cultural Artifacts

Once we agree to recognize the existence of conceptual artifacts as distinct from the content of people’s minds, there remains the problem of distinguishing them from other kinds of artifacts such as totem poles, symphonies, and electric toasters. My efforts to do so have continually been defeated by sharp critics, and so I am prepared to concede that rigorous distinctions are impossible. But it is still useful to have a perspective on the issue, and so I offer a set of nested definitions, based on the following diagram.



Starting with the broadest category, we may say that *cultural artifacts* include everything that is judged worth preserving because of its meaning. This doesn't include everything in the world, but anything in the world, concrete or abstract, could potentially become a cultural artifact. The electric toaster in your kitchen is not a cultural artifact. When it ceases to toast bread you throw it away. But if someone retrieved it from your trash can and it now appears in an art museum or a historical museum, it has become a cultural artifact. It is now being preserved because of its esthetic or historical significance. The tree in your back yard is not a cultural artifact or even an artifact at all. But if it should happen that the agreement that finally establishes world peace is signed under that tree, it will probably be preserved, having become a monument.

Some cultural artifacts are *abstract*. These include myths, folk tales and literary works, folk songs and musical compositions, as well as the various kinds of conceptual artifacts we have been discussing. Whereas concrete cultural artifacts such as paintings and monuments are preserved as such, and secondarily in the form of reproductions and representations, abstract cultural artifacts can only be preserved in such secondary ways. So we have

books, tape recordings, and other objects that are preserved not because of their singular value but because of their content, which is abstract.

Among abstract cultural artifacts, some can be treated as knowledge and some cannot. Here, drawing a sharp distinction is difficult and perhaps impossible, because there are always intermediate instances. I have suggested that *conceptual artifacts* can be distinguished by the fact that they can be used to rationalize behavior, but this is more a rule of thumb than a definition. We can surely agree that it is possible to argue about a theory in a way that we cannot argue about a sonata, but it is not so obvious how to regard a poem; and poems vary in the extent to which they contain discussible ideas. As I read Popper, I find him wavering between including all abstract cultural artifacts in World 3 and limiting it to ones that constitute knowledge. Trying to be practical about it, I maintain that for purposes of discussing issues of education, knowledge work, and knowledge creation, we need to distinguish as best we can those cultural objects that we want to treat as knowledge: hence the category of *conceptual artifacts*.

Finally, there is an important subclass of conceptual artifacts that I will discuss in the next section as ‘assertive’ artifacts. These are conceptual artifacts that can be argued about as to whether they are true or not. Theories, hypotheses, and factual claims fall into this class. This doesn’t mean that it is fruitful to argue about their truth, only that it is something people find it possible to do, whereas they would not find it possible to argue about the truth of an icon or a business plan. (They might, however, argue about the truth of some of the statements in a business plan or of beliefs about the meaning or power of an icon.) Epistemology’s traditional concern has been confined to this little box. Education is necessarily concerned with it too, but not to the obsessive degree that many science educators are concerned with it. Education is really concerned with the whole range of cultural artifacts. That is what liberal education has been about, and contemporary arguments about multiculturalism, Eurocentrism, and the like reflect this broad concern. What I am trying to do in this chapter is pick out the subcategory of *conceptual artifacts* as one that deserves special attention as knowledge assumes a more salient role in society.

Theories: Assertive Artifacts

By far the most troublesome conceptual artifacts have been those that assert something about World 1. Included in this category are the theories of empirical sciences, conjectures, facts, histories, and observational reports (“I saw the defendant leaving the building at 9 o’clock”). These objects have been almost the sole concern of epistemology. Traditionally, the term ‘knowledge’ is reserved for objects of this sort, or, more precisely, for that subset of such objects that happen to be true—and in this last qualification lies the constitutive problem of epistemology.

The last half of the 20th century has seen a great deal of ferment in epistemology and the closely related field of philosophy of science. It is no longer quite so obvious or easy to treat knowledge of the truth-asserting kind

as consisting of autonomous objects. Antifoundationalism, postpositivism, and lately postmodernism, along with many less massive assaults on the ivory tower, must at the least be credited with making convincing cases for viewing knowledge within an extended network of relations. The following points, distilled and simplified from a variety of contemporary utterances, seem, on the surface at least, to argue against the Popperian approach:

- There are no isolated facts. All supposed facts have theoretical presuppositions and thus are not fundamentally different from theoretical propositions.
- Theories do not exist in isolation either, but are embedded in paradigms. And paradigms are generally not fully articulated. They are more like traditions than they are like supertheories.
- What we call knowledge is merely belief that has gained acceptance in some group. Thus knowledge cannot be separated from the people who uphold it.
- There is no value-free knowledge. The beliefs that a group upholds as knowledge or truth are ones that subserve its interests.

Without presuming to judge the merit of these assertions, I think it can be safely concluded that they do not constitute an argument or even part of an argument against treating knowledge as objects. They are statements about *how* we should treat these truth-assertive objects: We should treat them as theoretical, interconnected, historically and socially situated. We should be attentive to the motives—in ourselves as well as in others—that lie behind the endorsement of certain conceptual artifacts and the rejection of others. Stated in these general terms, the postpositivist position is far from radical and is in fact plain common sense. In our daily lives we have to take enormous numbers of assertions on authority, having neither the time nor the resources to investigate them. Therefore, the credibility of the source is always important and open to question. How far one should go in this questioning—where reasonable skepticism leaves off and delusional suspiciousness begins—is an important matter, but I do not see that its consideration is in any way compromised by treating knowledge as consisting of real objects.

Knowledge Improvement

“[M]odes of social knowledge such as theology, science, and magic are different, not inferior or superior” (Stanfield, 1985, p. 392). This statement sums up an extreme but increasingly popular outcome of what has been one of the most vigorously disputed issues in contemporary thought. Does science have a privileged claim on our belief or is it only one among a number of ways that different peoples have tried to make sense of the world? Cautious university students these days will not use the word ‘science’ by itself. They will preface it with ‘Eurocentric,’ thus making clear their awareness that there is not just one kind of science and that the belief that there is such a privileged route to knowledge is ethnocentric and due for its comeuppance.

There are issues of lasting human concern here: What should we believe? Can we be certain of anything? What do we do about the fact that so much has to be taken on authority? Except for some minor attention to the last, I am not going to deal with those issues. Instead, I want to call attention to one obvious and important distinction between science and myth. People are constantly at work trying to improve the content of science. No such efforts at improvement are to be found amongst the priests, shamans, elders, divines and the like who have custodial responsibility for the myths, religions, or magic of a people.

Does this mean that science is more deserving of our credence? That is part of the issue I am trying to avoid. It surely depends on the criteria one applies, and problems multiply as soon as one tries to apply criteria of believability. What seems to me to be the important point for a new theory of mind and knowledge is that deliberate effort to improve a body of knowledge is something quite rare in the world. Relatively few of the world's many societies have undertaken it, and in those societies that have done so, it has been the occupation of a small minority.

Myths, rituals, magic, folk medicine, and the like undergo change over time, of course, and some of these changes, especially in practical arts, may be counted as improvements. But such change is an evolutionary process. Some cultural anthropologists treat it literally as such. Dawkins (1976) defined *memes* as cultural units comparable to genes. New ones arise as mutations; they proliferate and spread about through cultural processes and they are subject to selection pressures of various kinds. A culture as a whole may thus be characterized as a 'meme pool,' whose composition changes over time through the entry of new memes and through changes in the frequencies of existing ones. Like biological evolution, cultural evolution may appear in retrospect to have been rationally guided. The evolution of the English language, for instance, appears in retrospect to have been directed toward eliminating unnecessary inflections and toward minimizing effort in vocalization; but it is doubtful if either of these was ever a consciously entertained purpose. Had it been, we should probably have a record of groups of scholars commissioned to design improved versions of the language.

There have probably always been individuals who, in solving a particular problem, came up with something of general significance, thus contributing a mutant meme that might survive and spread. But knowledge improvement as a social enterprise has been much more limited. Earlier I offered a speculative sketch of the evolution of knowledge work from a sort of mystical quest for truth, thought to reside in a world beyond the visible world, to the constructive, inventive activity that yields the "disciplined progress" Whitehead talked about. Cultural evolution was not replaced, but it was greatly accelerated by the introduction of designer memes (cf. Dennett, 1995).

It is often said that science, for the modern masses, serves as religion or myth, and that technology is for them equivalent to magic. If so, however, part of the religion is belief in progress, and this sets it apart from older systems of belief. People expect our modern myths and magic to keep getting

better—better in the sense that more and more will become explainable, curable, achievable. This belief that progress just happens and is inevitable has the irresponsible character of all sorts of fatalism, and is in no way to be advocated (cf. Lasch, 1991). But we should recognize that it is an altogether different kind of fatalism from that of ancients and mystics.

Let us be clear that in talking about the improvement of knowledge we are not assuming that it leads to improvement in the human condition. We are not even assuming that the effort to improve knowledge is successful in any respect. The only point that bears on our overall purpose in this discussion is that knowledge improvement is a distinctive kind of effort, absent—often intentionally absent—from many spheres of human activity. Improvement of personal knowledge, acquiring a fuller grasp of the truths and skills already accessible and valued in the culture, is almost universally honored and of course figures prominently in all kinds of educational philosophies. But trying to improve knowledge itself, to extend the limits of the learnable, appears to be a relatively recent innovation and one that is far from firmly established in education, even in the most modern nations. It is what I have referred to as *knowledge building* and will later advocate as a proper activity for educational institutions at all levels.

Objections and Alternatives

The standpoint of women, which locates us in the particularities of our experience, is profoundly contradictory to objectified forms of knowledge.

—Smith, 1990

We need to be clear about where the sociocultural and the Popperian views of knowledge part company. It is not on the issue of abstractness. Sociocultural theorists commonly recognize knowledge that becomes abstracted from particular situations and available for more general use. (Arguments to the contrary would be self-refuting.) Where they draw the line is at detaching knowledge from the web of practices, motives, and power relations within which it has come to be treated as knowledge. As Dorothy Smith's remark at the opening of this section illustrates, the reason for drawing the line is often political rather than theoretical. Once knowledge is objectified, those who control the agencies through which it was created, propagandized, and bureaucratized—woven into the established social structure—are no longer accountable for it. It becomes a *fait accompli*, a part of the way things are, something to which people must now adjust. The authoritative status of textbook knowledge is just one, and by no means the most important, manifestation of how objectified knowledge can create subservience and alienation for those whose standpoint on reality is different.

I need hardly acknowledge that the issues here are vast and that they do not lend themselves to cursory treatment. Yet I must say that my own view is the opposite of Smith's. I do not see objectified knowledge as the problem or as underlying the problems she addresses—mainly problems having to do

with power and exclusion. It is the *unobjectified* knowledge of the powerful that the weak are helpless against—the unwritten laws, the unofficial practices, the unconscious ways of thinking and acting that are from childhood bred into them. These frustrate both revolutionaries and people who are trying to gain membership in the elite. Objectified knowledge, by contrast, can be argued against, sometimes with success; it can be accepted with reservations, or used for alternative purposes—even turned against the people who uphold it.

Basil Bernstein (1996) has criticized objectified knowledge on somewhat different grounds:

Once knowledge is separated from inwardness, from commitments, from personal dedication, from the deep structure of the self, then people may be moved about, substituted for one another and excluded from the market....

Now we have two independent markets, one of knowledge and one of potential creators and users of knowledge. (1996, p. 87)

Bernstein's fundamental concern here seems to be that objectification of knowledge licenses a kind of amoral detachment: no one takes responsibility for the knowledge or its uses. Again, it seems to me that the concern should be directed as much if not more to *unobjectified* knowledge. No one takes responsibility for the unwritten laws and beliefs, the tacit knowledge of their group, because no individual is in fact responsible. The ethos is what it is, whereas individuals can be called to account for the explicit beliefs by which they justify their actions.

Bernstein is certainly correct to this extent, however: People are valued for what cannot be detached from them, which includes their tacit knowledge and skills. If everything you know that is of value to your employer can be put on paper, then your employer can keep the paper and get rid of you. This has become an explicit concern in universities, where professors fear that if they put their courses online and the university claims ownership of the media objects thus produced, the next step will be for the university to get rid of the professor. But this is not a problem of epistemology, it is a problem of property rights. Defending those rights depends in fact on being able to distinguish knowledge from its particular representations and from the people who have custody of them. The professors must be able to claim that they own the knowledge, the intellectual property, even if someone represents it in a different medium and owns the hardware and software by which it is transmitted.

All this is quite obvious, however, so how has it gotten about that objectification is the culprit? The following is no doubt only a fragment of an

answer, but it is particularly relevant to the theme of education and mind in the Knowledge Age: What is exclusionary is *unobjectified knowledge of how to work with objectified knowledge*. That is what the 'haves' have in a knowledge-based economy that puts them at such an inapproachable distance from the 'have-nots.' Conceptual artifacts, detached from their personal and institutional sources, are assuming increasing importance in knowledge-based economies. Those who, for whatever reason, are unable or not disposed to deal with them are at an increasing disadvantage.

Much of what Smith and other critics say is consistent with this claim, but they becloud the issue by bringing in the traditional epistemological concern with truth. The validity of alternative truth claims has almost nothing to do with issues of equity. Cathleen Loving (1997) reports several instances at gatherings of science educators where it was declared that Australian Aboriginal or Alaskan Inuit myths should be given status as scientific explanations. Even if there were some merit in such proposals, they are retrograde because they so badly miss the point. The United States is full of Christian fundamentalists (polls put them at close to half the population) whose beliefs are as contrary to mainstream scientific knowledge as those rooted in hunter-gatherer cultures. Yet many of these people, I trust, are getting on nicely as knowledge workers; if they are not, it is not because their beliefs are being trampled upon. To put it crudely, one's view of the world is a handicap only if it interferes with one's work.

Instead of trying to win a place for myth and emotion in science class, those concerned with social justice ought to be trying to win a place for alternative talents in the market place. What we are seeing instead is a kind of epistemological Luddism. To battle against objectification is as futile as it was for the Luddites to battle against industrialization. The legitimate concerns of the Luddites eventually found effective expression in labor unionism, which, instead of battling against machines, battled for workers' rights and welfare. The rising economic value of work with objectified knowledge and the declining value of other kinds of work is producing grave inequities, but it is ridiculous to place the blame on objectification itself. In my view, educators who think that postmodern, antifoundationalist epistemology will lead to greater equity are being led down a path of pedantic irrelevancy. And that is the case irrespective of the theoretical value of their epistemology.

Postmodernists, especially in mathematics and science education, seem to be obsessed with truth. Of course, they believe there is no such thing, but they also believe this fact has overwhelming educational implications. In reality, very little of educational consequence follows from whether or not one believes there are truths waiting to be known (just as very little follows for practicing scientists). It does not follow that the believer in truth must be dogmatic, didactic, a suppressor of dissent—or an androcentric political conservative, for that matter. If a teacher believes that the

textbook utters absolute truth, that is a problem, but it is not an epistemological one. It is a matter of simple ignorance. The cure is not to read Thomas Kuhn or Sandra Harding but to become better versed in one's discipline.

Regardless of their own epistemologies, textbook authors have to decide whether to represent scientific knowledge as fact or as the provisional results of continuing inquiry. Presenting it as provisional is clearly superior if the intent is to develop an understanding of science and to start students on the road toward scientific inquiry themselves. But at the same time this approach is the one most devastating to traditional beliefs and most conducive to the objectification of knowledge. When scientific knowledge is presented as fact, people can find ways to compartmentalize and allow their traditional beliefs a separate existence. I imagine that is how most Christian fundamentalists handle it. But bring creation myths into a classroom where they are treated as knowledge objects—to be examined, criticized, used to explain phenomena, explained and considered for improvement—and they are going to take a beating. That is why it seems to me that the advocates of 'anti-racist' science and the like are living in a dream world. Their proposals can only produce the desired result through the vigorous suppression of inquiry and debate.

Theoretical Objections

The commonest theoretical objections to Popper's World 3 are of the "what if everybody died?" variety, which I have already tried to dismiss. A more interesting thought experiment along these lines is to imagine what would happen if everybody's working memory capacity—the number of ideas they can hold in mind and coordinate simultaneously—were to decline by half. People would no longer be able to grasp anything very complex. Most scientific ideas, all of mathematics beyond simple arithmetic, and any deep analyses of issues or events would lie beyond their comprehension. This would not be merely a matter of inability to comprehend certain texts. The complex ideas would be incomprehensible no matter how they were presented. But people would still be able to understand reports about these conceptual artifacts, narratives about their invention, and simplifications that gave an inkling of what they were about and what they were good for. There might, indeed, be continual valiant but futile efforts to understand the conceptual artifacts and to learn how to use them. So, would those conceptual artifacts have ceased to exist? You could say so, but for most purposes, it seems that this situation is best represented as one in which the conceptual artifacts continue to exist (for the time being, anyway) but are inaccessible, ungraspable. This way of viewing the situation is especially

appropriate for education, because what I have described is not some flight of science fictional imagination. It is the actual state of the normal eight-year-old child. Major portions of World 3 are inaccessible to them, because—for reasons that have been in dispute—they cannot perform mental gymnastics of an adult level of complexity.⁶ However, the situation for the normal educated adult is different only in degree. There are still large tracts of intellectual territory inaccessible to us because we lack the necessary background knowledge and skills or because we have not put forth the necessary effort to master them. But they nevertheless constitute parts of the known world for us, like countries we have never visited. We know that they exist, we know some things about them, and we may pursue understandings that point in their direction. Children, with help from their teachers, could do the same; but this requires treating conceptual artifacts as real things with properties of their own that make them worth knowing about. (This is the subject of Chapter 9).

A more serious theoretical objection to World 3 has its roots in linguistics and semiotics. The objection is that to posit a world of autonomous knowledge objects is to promote a manner of speaking into an ontological category. In the various registers of scholarly discourse, according to Wells (in press), “event-types that, in everyday speech, are realized in active clauses of doing, feeling and saying, are nominalized as abstract nouns and noun phrases; these abstractions are then treated as ‘things’ that can themselves enter into processes and relationships in the construction of descriptive and explanatory tests.” The fact that we talk this way—that our language may indeed force us to talk this way—does not mean, however, “that there is a corresponding immaterial object that then exists, independent of the linguistic formulation and argumentation through which it was constructed.” I do not see any reason to disagree with Wells on this point. It should be understood, however, that what he has offered is not an argument against treating conceptual artifacts as real things, it is an argument against claiming that we are logically compelled to do so. That clears the way for pragmatic arguments about the desirability of doing so, which I will turn to shortly.

⁶ Case (1985) has developed the theoretical and educational implications of this aspect of cognitive development, with educational implications further developed in Case (1992).

An important conceptual distinction has been elaborated by Philip Agre (1998). Agre distinguishes between objective concepts and deictic ones. Objective concepts attribute to objects an existence independent of particular situations. They include all the concrete artifacts and natural kinds that common sense recognizes. So in Agre's terms, I have been arguing that ideas, as conceptual artifacts, should also be treated as objective concepts. Deictic concepts, however, have meanings that are not constant but that vary depending on who is speaking and where and when and to whom. Examples are *the shirt I am wearing today's weather*, and *the place where we met*. To know what shirt, what weather, or what place, you have to know the context.⁷ A deictic concept particularly relevant to the present issue is *the idea we are talking about*. Agre does not argue that objective concepts should be dismissed in favor of deictic ones, but he does claim that deictic concepts are primary. They come first, and objective concepts are derived from them, whereas conventional wisdom would suggest that it must be the other way around.

I find Agre's argument especially convincing as it applies to abstract concepts like theories, plans, facts, proofs, and arguments. These probably originate as *what we are talking about* or, in the isolated case, *what I am thinking about*. Their role in conversation is to give it focus and help to move it along. These undefined deictic concepts begin to acquire some ontological status of their own, however, when *what we were talking about yesterday* or *what I was thinking about yesterday* is important enough to be brought up again and pursued further. If this revival of topics goes on, it soon becomes worthwhile to give *what we are talking about* a name—for instance, 'your theory.' But if 'your theory' is going to be discussed in other contexts where you are not present, it will eventually get a name such as 'Planck's law' or 'quantum theory.' It will be described in characteristic ways, will begin to accumulate a body of received opinion, will be discovered to have implications you were unaware of, and so on. Thus it evolves into an objective concept.

This does not mean that objective concepts are superior to their deictic predecessors. The benefits of having something relatively stable to refer to and work on are bought at the price of a certain rigidity and artificiality. For Agre, whose project is to bring computer technology into closer accord with human activity, deictic concepts are the more promising kind. *What we are doing now* is the meeting ground for human and machine intelligence, and so you do not want the machine running according to a lot of fixed concepts

⁷ Extending the idea, you can say that *shirt* in general is a deictic concept, because you cannot understand what a shirt is apart from the human purposes and human body configurations that make some things shirts and other things nonshirts. That line of reasoning, however, leads to the conclusion that all concepts are deictic, which is not what Agre argues and which, in my view, destroys the point of introducing the idea of deixis in the first place.

built into it, often unwittingly, by engineers far removed from the situation. You want the adaptability that comes from being able to conceptualize interactively and on the fly; the same of course is true of human beings. But one of the things human beings do is produce, reconstruct, and refine objective concepts or what I prefer to call conceptual artifacts. That is what education beyond the basic level and the more creative kinds of knowledge work are mainly about. In the discussions through which such activity goes on, deictic concepts will play a large part. But it is important for participants to keep in mind that some of the concepts brought into the discussion have histories outside the discussion and, moreover, that the object of the discussion is to produce ideas that can be carried away from the discussion and used in other contexts.

Boundary Problems

There are problems with conceptual artifacts. When are we talking about the same idea? When has a theory changed so much that it is no longer the same theory? If you are the kind of person who takes definitions and categories very seriously, these are likely to be vexing questions. But we should not be looking to epistemology for answers. Such problems are part of the work of any discipline, not something outside it. Drawing distinctions, showing that two hypotheses, problems, or proofs are really the same or that one is a special case of the other—these are part of the creative work of scientists and scholars, along with producing new conceptual artifacts or revising old ones.

Popper (1972, pp. 140-141) spelled out some of the things that a scientist, *S*, might do with respect to a knowledge object, *p*:

- S* tries to understand *p*.
- S* tries to think of alternatives to *p*.
- S* tries to think of criticisms of *p*.
- S* proposes an experimental test for *p*.
- S* tries to axiomatize *p*.
- S* tries to derive *p* from *q*.
- S* tries to show that *p* is not derivable from *q*.
- S* proposes a new problem *x* arising out of *p*.
- S* proposes a new solution of the problem *x* arising out of *p*.
- S* criticizes his latest solution of the problem *x*.

These are clearly actions, and they fall well within the normal scope of scientific activity; yet they are not actions carried out on the physical world (Popper's World 1). They are mental actions, but they are clearly actions carried out *on* or *with respect* to *something*. Furthermore, they are all carried out with respect to the same 'something,' and so to make sense of the scientists' actions and to see how they fit together, it is important that this 'something' have enough of a fixed identity that we can talk about different things being done to it, with cumulative results. This does not mean that we have to be able to say precisely what *p* is and where *p* leaves off and non-*p* begins. It just means that we have to identify something as the object that ties together a variety of actions. In the famous story of the blind men and the

elephant (where, feeling different parts of the elephant, one blind man concludes that it is a tree, another that it is a snake, and so on) there is the possibility that through discussion the blind men could work out a coherent description of the animal. But this possibility depends on their being aware that they are all talking about the same entity. Failing in that, they would have no reason to try to reconcile their descriptions.

To deal with constancy and shift in conceptual artifacts, Landman (1986) has proposed the concept of 'peg.' *Electron* is a peg to which we attach information, just as a clothes hook is a peg to which we attach clothes. You can add more information or you can take some away and replace it with other information, so that eventually what is hanging on the peg is entirely different from what was there originally, but the peg is continuous over time.

'Peg' is a makeshift notion, to be sure, but it serves to deal with what is really just a practical problem. When long time periods are involved or large conceptual distances, as between theoretical scientists and school children, then we need to consider seriously what we have to gain by treating these disparate groups as talking about the same thing. If we are writing history, then changes in meaning are part of the story and we must guard against producing a history that has today's theoretical pegs existing long in the past, with people in earlier times hanging quaint notions on them (Danzinger, 1990). Similarly, if we are writing child development, we need to consider that concepts to which children attach labels such as "heat," "energy," and "gravity" may be in entirely different ontological categories from the concepts scientists associate with those labels (Chi, Slotta, & deLeeuw, 1994). If, however, we contemplate having children and scientists enter into discourse together, then it would be good to work out some realistic sense in which they might see themselves as talking about the same conceptual artifacts. This would mean agreeing on some pegs on to which both groups could hang their ideas. These do not have to be rigorously defined; they just have to serve their purpose of enabling people to share ideas and to preserve some continuity as their ideas progress over the course of an extended conversation.

Potential Misuses

Ideas can be misused, and so a legitimate objection to the idea of conceptual artifacts may be that, although harmless when properly used, it is liable to be used to bad effect. Gordon Wells (in press) has raised two objections of this pragmatic sort: first, that treating knowledge as consisting of autonomous objects "misrepresents the way in which knowledge is constructed and used"; second, that "knowledge, by being reified, becomes a commodity to be transmitted to students and its possession subsequently assessed and quantified." The upshot is that strategically "there is more to be lost than gained by using Popper's three world model..."

Against these pragmatic objections it is not enough to show that, on the contrary, the three worlds model sets the stage for unrestricted investigation of how knowledge is constructed, of the role of discourse and the effects of various psychological and sociological factors on its development, and that it makes clear the artifactual character of textbook knowledge, its fallibility and improvability. Such virtues will be lost on educators who are liable to equate World 2 with students' naive beliefs and World 3 with the body of truths that are supposed to replace them. And people do tend to assimilate new ideas to their existing categories.

If we are going to be pragmatic about it, however, we have to consider alternatives on the same basis. The alternative that is probably favored by most advanced thinkers these days is one that treats knowledge as inseparable from the contexts of discourse and action in which it is constituted. As Wells (in press) puts it,

knowledge does not have an existence apart from the situated acts of knowing in which it is constructed, reconstructed and used. Theories and explanations only exist in the particular occasioned use of the semiotic representations of various kinds in which they are realized (or the reconstruction of these representations from memory) by specific individuals who are engaged in some activity in which these semiotic artifacts play a central role in the knowing of those involved.

With variations, that is the alternative favored by all the critics I have been citing—Smith, Bernstein, Wells—and behind them the growing number of people who adopt labels like 'semiotics,' 'socio-cultural,' 'situated cognition,' and 'action theory.' But this conception lends itself to misuse as well. At best it is hard to get hold of and to do much with. At worst, it becomes degraded into the anything-goes relativism of Stanfield (1985) and even such suicidal concoctions as 'race-based epistemologies' (Sheurich & Young, 1997).

The one idea I have been advocating that seems to be at least potentially acceptable to a wide range of critics is the idea of regarding theories and the like as tools. Wells explicitly recognizes conceptual artifacts in this sense, though adding the qualification that "these do not themselves constitute knowledge; they only play a role as tools that have a potential for facilitating problem solving and further knowledge building..." The tool idea is not hard to grasp and it would seem to be safe from the worst excesses of relativism. No one would claim that every tool is as good as every other. The

value of a tool is relative, but always relative to some purpose. So far so good. It may well be that pushing the notion of conceptual artifacts as tools is the entering wedge for breaking through folk theory's barriers to understanding knowledge. But it seems important to press on from there to a recognition that conceptual artifacts are not just tools. Some of them make assertions about the world that we may want to judge as true or false. Some of them function like recipes, which is not quite the same as functioning as a tool. Even as tools, they may themselves be objects of inquiry. We can investigate how they work, what their limitations are, and how they might be improved. Conceptual artifacts relate to one another in ways that physical tools cannot; one of them may imply or be implied by others. This creates the possibility of assembling them into larger and more integrated structures—which suggests treating them more as building materials than as tools.

There is no use looking for an idiot-proof way of treating knowledge. Into the foreseeable future, most people are going to continue thinking of knowledge as stuff variously contained in people's minds and in books and of education as a business of getting it moved from one container to another. There is no official theory to that effect and no group committed to upholding it. It is just a matter of the way we were born and bred. For that reason, however, I do not believe that it is a bad mark against any theory of knowledge that it can be distorted so as to support this unenlightened view. The unenlightened don't need any such support. Some advanced ideas can do harm. I suspect that relativism generally has had a bad effect on education because, regardless of the merits of its more sophisticated versions, in practice it has discouraged the serious pursuit of understanding. What would be the effect of trying to put into practice the ideas developed in this chapter? That should become clearer in later chapters. If its only effect is to provide spurious theoretical support for an educational approach that people were going to follow anyway, that is the same as having no effect at all. If, however, it can get educators starting to treat knowledge as something that can in fact be treated in different ways for different purposes, then it may have opened a door through which education could conceivably enter the Knowledge Age.

Progressive Discourse

Where Popper and contemporary socioculturalists are in solid agreement is on the importance of dialogue in the practice of science and other

disciplines. This applies to empirical sciences as well as to more speculative disciplines. Empirical observations by themselves do not constitute scientific knowledge. They influence the development of scientific knowledge by being brought into the discourse through which knowledge is constituted. One of the tenets of this discursive view of science is that science does not start with facts. Rather, statements become facts through a lengthy discursive practice, at the end of which people no longer see any virtue in contesting them (Latour, 1987). On that basis, we could say it is a fact that scientific knowledge is constituted through discourse. There is no longer any lively interest in disputing that statement. What this means about the epistemological status of scientific claims—whether they amount to anything more than consensus—is still a controversial matter. But regardless of your beliefs about truth and the nature of knowledge, you are likely to find little to gain by questioning the centrality of discourse.

To Popper, discourse was a sort of Darwinian process by which weaker theories are eliminated and the stronger survive—until replaced by yet harder ones. That is not much different from the view of some contemporary sociologists of knowledge, who see scientific dialogue mainly as competitive and the whole business a struggle for dominance (Latour, 1987; Latour & Woolgar, 1979). Popper saw the process as rather more high-minded than the sociologists do, but that could be the difference between idealization and description. In either case, their focus seems to be on knowledge claims that have already been formed—that already constitute World 3 objects, in Popper's terms—rather than on how such knowledge comes about in the first place. Kevin Dunbar (1993), by contrast, has studied the discourse that goes on in research laboratories as interdisciplinary teams try to work their way to the solution of scientific problems. Dunbar finds predominantly cooperative effort, with a great deal of explaining as opposed to argument. Woodruff and Meyer (1997) have suggested that both descriptions are valid but that they apply to different discourse communities. There is, on one hand, intragroup discourse in which the main work of knowledge construction is carried out and in which cooperation is the norm. On the other hand there intergroup discourse in which the products of different research groups are brought under consideration and in which criticism and controversy are normal. Both kinds of discourse have valuable roles to play in the development of knowledge.

Discourse plays an important part in all human communities, but its role is obviously going to be different in a community of historians from what it is in a community of body builders, for instance. Our particular concern here is with communities whose job is creating and improving knowledge. With the possible exception of philosophers, this job is never carried out through discourse alone. To understand a community of practice such as a scientific discipline, we need to understand the social and material circumstances in which it functions, the tools at its disposal, the division of labor within the community, and the norms regulating behavior in the community (Engestrom, 1987). Accordingly, there is some danger in focusing too much on

discourse, as if it were the community's preoccupation and reason for being. What I would say, rather, is that discourse is a screen onto which everything else is projected and from which the success of the community may be read.

What it means to create and advance knowledge is going to be understood differently in different communities, but whatever it means must be mirrored by a conception of progressive discourse. It would not make sense to claim that knowledge is advancing but that the discourse within a discipline is getting nowhere. When we talk about the centrality of discourse we are claiming more than that, however. We are saying that knowledge advancement and progress in knowledge-building discourse are not two things that mirror one another but are one and the same thing.

In psychology, which is the discipline I know best, the colloquium is a time-honored form of discourse. In the prototypic colloquium, a speaker presents a series of four to ten experiments. The first tests some theoretical notion. Each succeeding experiment overcomes a limitation or eliminates an alternative explanation of the results of its predecessor until, with the final experiment, every cloud of doubt has been removed, including clouds hardly anyone else would have thought of. There follows a question period during which the more clever and informed members of the audience try to show that in fact unnoticed limitations or alternative explanations still exist. It is difficult not to be impressed by such performances, even when one is not much interested in the point at issue. This colloquium form clearly embodies the Popperian notion that science advances through the elimination of hypotheses. Also, it is clearly not just talk. No amount of discursive skill will take the place of having planned and conducted a series of experiments that interlock in the particular way required to support the kind of story that a successful colloquium speaker aims to deliver. But, just as clearly, the design of the experiments is guided by the kind of story the psychologist aspires to tell. The criticisms that arise during the question period may determine the next experiment the psychologist carries out, and hence the next chapter in the lengthening story to be told in subsequent colloquia, or they may mark the beginning of someone else's research program that extends or runs counter to that of the speaker.

The colloquium form I have described is most characteristic of experimental psychology, which has been the dominant branch of North American psychology throughout most of its history. A somewhat different form has evolved in cognitive developmental psychology. It is still based on a series of experiments, but the experiments are more likely to consist of various ingenious demonstrations of the same point. Although this might seem like a minor stylistic difference, it entails a different way of going about experimental research and it implies a different sort of progression in knowledge: The fitness of a theory is not so much determined by survival of criticism and counterexplanations as by its success in generating captivating experimental demonstrations. Knowledge advances not so much through progressive refinements and paring down of alternatives as through periodic innovations, not unlike progress in the automobile industry. A

developmental psychologist I know, whose work is highly successful by this standard, applied for a position in a department dominated by the experimentalist paradigm. She was told that her research was brilliant but that it wasn't the way they did things there.

Traditionally, scientific method has meant empirical methods, and so-called 'methods' courses perpetuate this view. On this view, one would look for differences between experimental and developmental psychology in the ways data are obtained and analyzed. No very profound differences would appear. The same principles of experimental design and control would be honored and the same kinds of statistical reasoning applied.⁸ It would appear that the same basic methodology is followed, except that the experimental psychologists follow it more rigorously. I suspect that that is in fact how some experimental psychologists see the difference. I have tried to suggest that the differences run deeper than this, that they might even be thought to represent divergent paradigms; but, if so, where do the differences reside? They do not reside in any particular rules or procedures but rather in the whole systems of practice to be found in the two subdisciplines. However, these differences show up in, are indeed vividly projected onto the discourse of colloquia.

Just as discourse practices distinguish one discipline from another, they also show what disciplines have in common and what distinguishes knowledge-building communities in general from other kinds of discourse communities. Experimental psychologists and developmental psychologists can, after all, converse, criticize one another's ideas, suggest questions for further research, and—not least—argue constructively about their differences. On the other hand, there are adherents of cults whose principal tenets are psychological but with whom it is impossible to carry on a discussion of psychological issues. Many intellectuals these days question whether there is any fundamental difference between science and other ways of making sense of the world (e.g., Rorty, 1991). Any attempt to carry on a knowledge-building discussion with a cultist should make it evident, however, that important differences do exist, even if it is debatable whether the line of demarcation runs between science and nonscience or between some larger class of rational inquiry and sense-making activities that lie outside it.

The psychological cultist may tell you, for instance, that his leader has determined that there are six types of personality. Having heard what these are, you may engage with him in an entertaining hour of classifying mutual acquaintances and public figures. But the cultist has no interest in comparing his typology with other typologies, such as Jung's, that have gained some scientific standing—much less any interest in examining the value of type theories in general. The six types are among the fundamental truths that his cult has established beyond doubt. Only their application to particular cases is

⁸ Many psychologists these days adopt what are called 'qualitative' methods, which do deviate from the experimental psychology norm in how data are gathered and analyzed; but I have purposely chosen examples where the main differences do not take this form.

open to discussion. But how is this different from the scientist's acceptance of certain propositions as facts and no longer worth questioning?

I don't think the difference can be pinned down by examining what people take as disputable and indisputable. One must look instead at the dynamics of sense-making discourse, as carried out over sizeable stretches of time, and at the motives that drive and guide it. What becomes immediately apparent from this viewpoint is the scientists' commitment to knowledge improvement. This entails both a belief that all knowledge is potentially improvable and a moral conviction that one should keep trying to improve it. From this commitment to knowledge improvement flow a number of distinctive characteristics of what I have called progressive discourse (Bereiter, 1994):

1. *Focus on conceptual artifacts.* All serious discussions involve ideas, but it is not so common for the ideas themselves to become objects of discussion. Our psychological cultist is happy to apply the ideas represented in his six personality types, but he is not interested in discussing the typology itself or the ideas behind it. Without doing so, knowledge improvement cannot even get under way.
2. *Improvability as a positive attribute of conceptual artifacts.* I owe to Marlene Scardamalia the realization that not all conceptual artifacts are improvable. The clichés of education tend not to be: 'Every child is different.' 'Children's minds are not empty vessels waiting to be filled.' It is not just that such statements are irrefutable. They act as thought-stoppers. Having assented to them, one has nowhere to go. There is no forward direction. An improvable conceptual artifact, by contrast, is likely to strike us as interesting, at least somewhat unsettling to our existing beliefs, and as raising questions and having implications beyond those that are immediately apparent. An example is the idea that children first develop an awareness of other minds and from this become aware of their own (Gopnik, 1993). Encountering this idea for the first time, you may realize that you have unquestioningly assumed the opposite, that the idea nevertheless has some immediate plausibility, and that the whole question of how theories of mind are acquired may involve quite perplexing chicken-and-egg issues. A commitment to knowledge advancement accordingly entails a commitment to formulate conceptual artifacts that are improvable—hence, vulnerable to criticism and disconfirmation.
3. *Common understanding given priority over agreement.* In politics and other practical affairs, it is often advantageous to have vague principles that you can get people to agree to even though they understand them differently. In knowledge-building discourse, however, agreement without common understanding is pointless, and one of the more discouraging things that can happen is to find that people who endorsed your ideas turn out not to have understood them. On the other hand, when there is mutual understanding of the

ideas under discussion, there is a basis for productive discourse even if there are major differences of opinion.

4. *Commitment to expand the factual base.* Even though people may disagree seriously, there are bound to be some facts on which they do agree, and these form a necessary basis for progressive discourse. Facts constrain ideas. Opposing theories may be compared to the extent that they attempt to explain the same facts (Thagard, 1989). Accordingly, people committed to knowledge advancement have a stake in expanding the set of accepted facts, while recognizing that any factual statement is itself a conceptual artifact open to criticism and improvement. This is quite a different commitment from that which prevails in formal debates and in court trials, where part of the strategy for success is to undermine as many as possible of the opponent's factual claims.
5. *Selective criticism based on knowledge advancement goals.* Sophomoric criticism is criticism disconnected from purposes of advancing knowledge. In the behavioral sciences, students quickly learn how to criticize studies for not measuring certain variables, for not including this variation or that, for small sample sizes, for not meeting all the requirements of certain statistical tests, and so on down a list; but they are often dumfounded if asked why their criticism matters—what conclusion of the study is challenged, what of importance would be learned by elaborating the study in the ways they propose. A commitment to *constructive* criticism is not a commitment to being gentle or to making only positive suggestions. It is a commitment to using criticism in the service of knowledge advancement.
6. *Nonsectarianism.* No one expects complete impartiality, even in high court judges and Nobel Laureate scientists. What can be expected of those seeking to advance knowledge, however, is that they will not adhere to sects that actively limit the scope of ideas to be entertained or tested. Much of nineteenth century medicine consisted of a battle between sects, each adhering to a restrictive set of beliefs and practices. But as scientific medicine began to take shape, that changed. “Scientific medicine,” said William Rothstein (1985, p. 325), “was not a triumph of any sect; it was the death of all sects.” Of course, sects live on, in medicine as elsewhere, but a commitment to progressive discourse entails a commitment to their elimination.

These six commitments constitute virtues and imply a rather saintly standard of conduct. It is not for me to say how much the actual behavior of scientists and scholars conforms to them. What I am saying is that these commitments represent what it means to engage in knowledge-building discourse. Any group interested in advancing knowledge would want to do what it could to encourage these commitments. They also serve to distinguish knowledge-building discourse from other serious kinds of discourse. For instance, argumentation is a form of discourse often advocated

in education (e.g., Kuhn, 1993). Raise an issue, get students to commit to a position on it, and then let those with opposing positions argue it out. Yet argumentation can proceed in quite a rational and civilized way without honoring any of these commitments.⁹ (Which, to my mind, raises serious questions about its value as an educational activity.) Knowledge-building discourse may take various forms and serve various specific purposes, but there is always an expectation that knowledge-building discourse will be progressive. Unlike social smalltalk, for instance, the discourse is supposed to *get* somewhere. Participants need to feel that something has been accomplished, that the state of knowledge in their community—however small or grand they may conceive their community to be—is in better shape than it was before.

Conclusion: Why Two Worlds Are Not Enough

Every way of dividing up reality has its problems, and so a pragmatic choice has to weigh gains against losses. In the Western world, what has come to be the commonsense or intuitively obvious way is to divide the world into two parts, the external world of physical reality and the internal world of our thoughts and perceptions—what Popper called World 1 and World 2. Descartes declared such a dualistic view of the world back in the 17th century, and it has been the object of much philosophical controversy ever since. There have been repeated attempts to eliminate one or the other of these two worlds. Arguments for eliminating World 1 often turn on the claim that we can have no knowledge of it apart from our thoughts and perceptions, and so—like it or not—World 2 is all we have. Arguments for eliminating World 2, which I touched on in the preceding chapter, turn on the claim that we, along with our thoughts and perceptions, are part of the natural order and therefore ought to yield to the same explanatory efforts as other parts. But the dualism persists in ordinary thought, and for good reason. It is terribly important whether the gun is in fact loaded, irrespective of what we or anyone else happens to believe; and, in the event that the gun was loaded and discharged a bullet into an unfortunate visitor, the courts must do their best to determine whether the person who pulled the trigger believed the gun to be loaded. So, while the struggle against dualism continues in scholarly journals, everyday affairs—including those conducted by philosophers in their off hours—are bound to continue in the old dualistic mode.

An important departure from the continuing controversies has been the idea put forth by Bolzano, elaborated by Frege and then Popper, that two worlds are not enough. The reasons that these philosophers thought two worlds were not enough are not, however, the reasons why they are not enough for education. Or, at any rate, the educational reasons only partially reflect the underlying epistemological reasons. We might say simply that two

⁹ There are, however, approaches to argumentation that do place cooperative knowledge advancement above controversy—for instance, *Coalescent Argumentation* (Gilbert, 1997).

worlds do not allow enough degrees of freedom for education to proceed in a sufficiently flexible and adaptive way. Or to put it in more visualizable terms, they do not provide enough dimensionality. They force educational thought into a two-dimensional space when it needs at least three dimensions.

The typical teaching situation involves three main elements. Let us consider them as they appear in science teaching. There is World 1, which may be thought of as how the physical world really is, how it really works. (Ignore for now any objections you may have to acknowledging such a reality; all I am claiming is that the typical science teaching situation presupposes a physical reality which it is the object of the course to teach students about.) Then there is World 2, which for the teacher is the content of students' minds. The third element is scientific knowledge, the stuff represented in textbooks. If everything is either World 1 or World 2, then this third element must somehow be relegated to one of those. (There is another possibility, to ignore authoritative knowledge altogether and limit science education to what students can find out for themselves through experimentation; but for obvious reasons this cannot be considered a serious alternative.) The old-fashioned way is to treat the textbook as a stand-in for World 1, that is, as a record of "how it really is." Although this approach leans toward treating textbook information as infallible and immutable, it is possible to fudge a bit and allow skepticism to enter the picture. But the more uncertainty that is introduced, the more students are bound to question why they should be required to learn the stuff. The postmodern or 'postpositivist' way is to treat the textbook as a record of what scientists believe, thus relegating it to another World 2, the minds of those fellows in white coats. The question thus naturally arises of why those people's beliefs should be taken more seriously than anyone else's. There are sophisticated answers available to that question as well as various compromises (reviewed in Loving, 1997), but in practice answers will either tend back toward the old-fashioned position—what the scientists say is the best representation we have of "how it really is"; after all, they got us to the moon, didn't they?—or toward the conclusion that indeed scientists' beliefs should not receive preferential treatment in comparison to the students' own beliefs or those of folklore, religion, and supermarket tabloids. No matter whether you relegate textbook knowledge to World 1 or World 2, you run into serious problems of motivation and credibility. Leaf through issues of the journal, *Science Education*, and you will see that these problems occupy an inordinate amount of science educators' attention, attention that might better be directed toward remedying the failures of science education to get across an understanding of fundamental scientific ideas.

Something else that seems to get lost in these two-dimensional approaches to science education is any realistic experience of what it is like to do science. Doing science may be reduced to activity in World 1—hands-on work with physical and biological things—or to mucking around with personal beliefs in World 2. Both of these go on in science, of course, but even in combination they miss the main point, which is creating scientific

progress. There is no good way to represent this in the two-dimensional world of folk epistemology. There are ways to work around the limitations, especially if one is not concerned with logical consistency; but overall the impression is that of three dimensions having been squashed down to two. Many important relationships are no longer represented.

When we move to three worlds, the first two remain as before. There is the real physical world, which remains important because, inaccessible as it may be to direct knowledge, it is what science is trying to understand. And of course there are the mental states of students and others, also not directly accessible but unavoidably important to people in their roles as learners and teachers. But scientific knowledge need no longer be reduced to these two categories. It consists of conceptual artifacts that stand as a third element in their own right. There are now three dimensions, and these provide space for a number of relationships that are lost in the two-dimensional world of folk epistemology. There is the relation between theory and observation, between personal belief and observation, and between personal belief and the theory. And there are the relations between different theories, different phenomena, and different people's readings of the same phenomena. None of these relations are easy. They are all inferential and highly problematic. But they are what people work on when they are building scientific knowledge.

In other disciplines the conceptual artifacts are different. There are historical accounts and explanations, interpretations and criticisms of literary and artistic works, and so on. But in all cases these add an important third dimension, without which there is no way to give an adequate representation of the learner's relation to the content of the discipline.

The idea of World 3 has a strangeness to it, but I think that is one of its virtues. It is not actually a very radical notion at all. As noted earlier, treating knowledge as something real in itself is one of the ways commonsense deals with it—the other being treating it as mental content. All Popper has done is to sharpen the distinction and confront us with its implications. In order to do that with concepts that have characteristically gone unexamined and been muddled together, it is often a good idea to articulate them in some unusual, often metaphorical way—to make the familiar strange, so as to help people think about it in less habit-bound ways.