

Learning as Constructive Activity

Ernst von Glasersfeld

For about half a century behaviorists have worked hard to do away with “mentalist” notions such as meaning, representation, and thought. It is up to future historians to assess just how much damage this mindless fashion has wrought. Where education is concerned, the damage was formidable. Since behaviorism is by no means extinct, damage continues to be done, and it is done in many ways. One common root is the presumption that all that matters — perhaps even all there is — are observable stimuli and observable responses. This presumption has been appallingly successful in wiping out the distinction between training and education. Educators share the goal of generating knowledge in their students. But knowledge is not a transferable commodity and communication not a conveyance. Children are not repositories for adult “knowledge,” but organisms which, like all of us, are constantly trying to make sense of, to understand their experience. If we come to see knowledge and competence as products of the individual’s conceptual organization of the individual’s experience, the teacher’s role will no longer be to dispense “truth,” but rather to help and guide the student in the conceptual organization of certain areas of experience.

The general topic I was given for this chapter is “Research in Mathematics Education from an Epistemological Perspective.” That sounds no more dangerous than so many other academic topics. But don’t let the prosaic surface deceive you. To introduce epistemological considerations into a discussion of education has always been dynamite. Socrates did it, and he was promptly given hemlock. Giambattista Vico did it in the 18th century, and the philosophical establishment could not bury him fast enough. In our own time there was Jean Piaget. He really wanted to stay out of education but allowed himself to be drawn in — and we know what has happened to his epistemology at the hands of interpreters and translators. It seems that to discuss education from an epistemological point of view was a sure way of committing intellectual suicide. Recently, however, the world of education may have begun to change. At least the particular discipline that is represented in this meeting, the discipline that is concerned with numbers, with arithmetic, and ultimately with mathematics, is manifesting symptoms that indicate the will to change.

The rapid shifts in the methods of mathematics education that have taken place in the last few decades — from simplistic associationism to “New Math” and “Back to Basics” — did not work the miracles that were expected of them. Their failure has created a mood that no longer fosters enthusiasm for new gimmicks. Today, I think it is fair to

Originally published in *Proceedings of the 5th Annual Meeting of the North American Group of Psychology in Mathematics Education*, Vol. 1 (Montreal: PME-NA, 1983, pp. 41– 101), edited by J.C. Bergeron & N. Herscovics.

say, there is a more or less general disillusionment. This disillusionment is healthy and propitious because it pushes us closer to the point where we might be ready to review some of the fundamental presuppositions of the traditional theories of education. Among these presuppositions are our conception of teaching and learning and, most fundamental of all, the conception of what it is “to know.”

Ten or 15 years ago, it would have been all but inconceivable to subject educators or educational researchers to a talk that purported to deal with a theory of knowledge. Educators were concerned with getting knowledge into the heads of their students, and educational researchers were concerned with finding better ways of doing it. There was, then, little if any uncertainty as to what the knowledge was that students should acquire, and there was no doubt at all that, in one way or another, knowledge could be transferred from a teacher to a student. The only question was, which might be the best way to implement that transfer — and educational researchers, with their criterion-referenced tests and their sophisticated statistical methods, were going to provide the definitive answer.

Something, apparently, went wrong. Things did not work out as expected. Now there is disappointment, and this disappointment — I want to emphasize this — is not restricted to mathematics education but has come to involve teaching and the didactic methods in virtually all disciplines. To my knowledge, there is only one exception that forms a remarkable contrast: the teaching of physical and, especially, athletic skills. There is no cause for disappointment in that area. In those same 10 or 15 years in which the teaching of intellectual matters has somehow foundered, the teaching of skills such as tennis and skiing, pole jumping and javelin throwing, has advanced quite literally by leaps and bounds. The contrast is not only spectacular but it is also revealing. I shall return to this phenomenon at a later point when, I hope, we will be able to consider an analogy which, at this moment, might seem utterly absurd. If educational efforts are, indeed, failing, the presuppositions on which, implicitly or explicitly, these efforts have been founded must be questioned and it seems eminently reasonable to suggest, as did those who formulated the topic for this discussion, that we begin by inspecting the commodity that education claims to deal in, and that is “knowledge.”

This chapter is an attempt to do three things. First, I shall go back to what I consider the origin of the troubles we have had with the traditional conception of knowledge. This historical review will not only be sketchy, but it will also be quite biased, because I have rather strong views on the subject. However, considering the mess in which the theory of knowledge has been during the last 50 years in the “hard” sciences, my attempt will, I hope, not be deemed unjustified.

Second, I shall propose a conceptualization of “knowledge” that does not run into the same problem and that, moreover, provides another benefit in that it throws helpful light on the process of communication. As teachers, I said a moment ago, we are intent upon generating knowledge in students. That, after all, is what we are being paid for, and since the guided acquisition of knowledge, no matter how we look at it, seems predicated on a process of communication, we should take some interest in how that process might work. In my experience, this is an aspect that has not been given much thought. Educators have spent and are rightly spending much time and effort on curri-

culum. That is, they do their best to work out what to teach and the sequence in which it should be taught. The underlying process of linguistic communication, however, the process on which their teaching relies, is usually simply taken for granted. There has been a naive confidence in language and its efficacy. Although it does not take a good teacher very long to discover that saying things is not enough to “get them across,” there is little if any theoretical insight into why linguistic communication does not do all it is supposed to do. The theory of knowledge which I am proposing, though it certainly does not solve all problems, makes this particular problem very clear.

Lastly, having provided what I would like to call a model of “knowing” that incorporates a specific view of the process of imparting knowledge, I shall briefly explore a way to apply that model to the one thing all of us here are interested in: how to introduce children to the art, the mystery, and the marvelous satisfaction of numerical operations.

I.

The nature of knowledge was a hotly debated problem as far back as the 6th century B.C. The debate has been more or less continuous, and while in many ways it has been colorful, it has been remarkably monotonous in one respect. The central problem has remained unsolved throughout, and the arguments that created the major difficulty at the beginning are the very same that today still preclude any settlement of the question.

The story begins with the first documents on epistemology that have come down to us, the so-called “fragments” of the Pre-Socratics. The ideas these men struggled with and tried to clarify must have arisen some time before them, but since we have no earlier written records, that background is extremely hazy. The Pre-Socratics, at any rate, exhibit a degree of sophistication that is unlikely to have been acquired in one or two generations. Fragmentary though they are, their pronouncements leave no doubt that, towards the close of the 5th century B.C., the process of knowing had been conceptually framed in a relatively stable general scenario. By and large, the thinkers who concerned themselves with the cognizing activity tacitly accepted the scenario in which the knower and the things of which, or about which, he or she comes to know are, from the outset, separate and independent entities.

I want to stress that this dichotomy does not coincide with the split between the knowing subject and the subject’s knowledge. That second dichotomy appears whenever an actor becomes aware of his or her own activity or when a thinker begins to think about his or her own thinking. That second problem of self-consciousness is not identical with the problem of cognition. Though the two are related in that they interact (e.g., in an analysis of reflective thought, which will enter our discussion later), I here want to deal only with the first. The Pre-Socratics, in any case, took for granted the human ability to be aware of knowing. What they began to wonder about was how it was possible that one could come to know the world. It is in this quest that the cognitive scenario they accepted and that has been perpetuated by almost all epistemologists after them, is of decisive importance. Once it was chosen as the basis for the construction of a theory of knowledge, that construction was saddled with a paradox. The paradox is

inescapable and it has haunted philosophers incessantly in the 2,500 years since then.

The reason why that particular cognitive scenario was adopted is very simple. It reflects the situation as it initially appears to any experiencer. The question, how it comes about that we know anything, is not likely to be asked at the beginning of a prospective knower's development. A six-year-old who bicycles home from school would be a very peculiar six-year-old if she suddenly asked herself, "How on earth do I manage to find my way home?" or "What exactly happened when I learned to ride my bicycle?"

I am not suggesting that these are questions a six-year-old or, indeed, anyone should ask. I am merely saying that if we ever do ask them, it will be at a somewhat later age. The same goes for the question, "How is it that I can know what I do know?" Those who have felt such epistemological curiosity probably formulated their first relevant question in their middle teens or later. That is to say, they began to question their knowledge at a point in their cognitive career when they had already acquired an enormous amount of know-how and learning. Inevitably, nearly all they knew at that point was tacitly assumed to be knowledge about the environment, about the world in which they found themselves living. It is not surprising that this should be the case. Once one has learned to manage things, there is no reason to suspect that they might not be what they seem.

If a person whose knowledge has been growing and expanding over the years then raises questions about how one comes to have all that knowledge, it seems reasonable to postulate at the beginning an inexperienced and totally ignorant knower, who comes into the world, much as an explorer might come into a terra incognita, with both the need and the will to discover what that world is like. The first if not the only tools that seem to be available for such a task are obviously the senses. Therefore, the senses are at once categorized as organs, or channels, through which the experiencer receives messages from the environment. On the basis of these messages, the experiencer then must, and apparently can, build up a "picture" of the world. In our contemporary jargon, this is often expressed by saying that the senses convey information which enables the experiencing subject to form a representation of the world. Usually this seems to work quite well. Occasionally, of course, the senses turn out to be somewhat deceptive, but by and large they work well enough for us to build up a *modus vivendi*. Provided we remain patient and flexible, we will continue to make adjustments, and as long as things work moderately well, there will be no need to question the overall validity of whatever picture of the world we have built up.

The Pre-Socratics started out in this thoroughly normal way, but because there were some highly original thinkers among them, they came up with mutually incompatible pictures of the world [1]. Obviously, that was felt to be a problem and it led to two closely connected questions: One, how could anyone compose a picture of the world out of sensory messages and, two, how could one be certain that a particular picture of the world was "true"? Attempts to answer these questions soon ran into troubles, some of them so serious that they have not yet been overcome. Here I want to focus on the second problem because it is inherent and unavoidable in the discoverer's scenario. If experience is the only contact a knower can have with the world, there is no way of

comparing the products of experience with the reality from which whatever messages we receive are supposed to emanate. The question, how veridical the acquired knowledge might be, can therefore not be answered. To answer it, one would have to compare what one knows with what exists in the “real” world — and to do that, one would have to know what “exists.” The paradox, then, is this: to assess the truth of your knowledge you would have to know what you come to know before you come to know it.

The argument that the likeness or trustworthiness of a picture can be assessed only by looking at both the picture and what it is supposed to depict, was brought forth already at the time of the Pre-Socratics and it has been the mainstay of all skepticism ever since. The history of Western epistemology is the history of more or less imaginative attempts to circumvent it. None of these attempts was satisfactory. Plato’s poetic genius almost succeeded in eliminating the dilemma by undercutting the role of experience. He placed the real reality into the world of ideas and turned sensory experience into a secondary affair, murky, unreliable, and ultimately irrelevant to the quest for truth. Since the world of ideas was accessible only to the thinking mind, this arrangement bred the notion of solipsism, the notion that there is no real world beyond the world the mind creates for itself.

Alternatively one could, as Descartes suggested, place one’s faith in God and trust the divine maker not to have been so malicious as to have provided his creatures with deceptive senses.

Neither of the two alternatives provides a durable solution. Solipsism turns into absurdity whenever an idea we have conceived is shattered by experience. In fact, that is not a rare occurrence. We are constantly reminded that the world we live in is not quite the world we would like and that there is, indeed, a hard and unforgiving “reality” with which we have to cope. On the other hand, the Cartesian notion boils down to a simple injunction to believe, and that does not satisfy the philosopher’s need. If epistemology must be founded on the blind faith that our senses convey a true picture, it cannot accomplish what it sets out to do, namely, provide a rational basis for the generation and assessment of knowledge. Actually, Descartes’ injunction to trust our God-given senses merely shifts the problem. If the senses were thought to be trustworthy, the fact that we so often draw the wrong conclusion from their messages should show that there is a serious difficulty of interpretation; and if we cannot be sure how to interpret what the senses tell us, we again have to admit that we have no certain knowledge of the world and that the picture we come to have of it remains questionable.

The problem, as I suggested at the beginning, is intrinsic to the traditions scenario. It arises from the “iconic” conception of knowledge, a conception that requires a match or correspondence between the cognitive structures and what these structures are supposed to represent. Truth, in that conception, inevitably becomes the perfect match, the flawless representation. The moment we accept that scenario, we begin to feel the need to assess just how well our cognitive structures match what they are intended to represent. But that “reality” lies forever on the other side of our experiential interface. To make any such assessment of truth we should have to be able, as Hilary Putnam recently put it, to adopt a “God’s eye view” [2]. Since we are not, and logically

cannot be, in a position to have such a view of the “real” world and its presumed representation, there is no way out of the dilemma. What we need is a different scenario, a different conception of what it is “to know,” a conception in which the goodness of knowledge is not predicated on likeness or representation.

The first explicit proposal of a different approach originated in those quarters that were most concerned with faith and its preservation. When, for the first time, the revolutionary notion that the Earth might not be the center of the universe seriously threatened the picture of the world which the Church held to be unquestionable and sacred, it was the defenders of the faith who proposed an alternative scenario for the pursuit of scientific knowledge. In his preface to Copernicus’ treatise *De revolutionibus*, Osiander (1627) suggested:

There is no need for these hypotheses to be true or even to be at all like the truth; rather, one thing is sufficient for them — that they yield calculations which agree with the observations [3].

This introduces the notion of a second kind of knowledge, apart from faith and dogma, a knowledge that fits observations. It is knowledge that human reason derives from experience. It does not represent a picture of the “real” world but provides structure and organization to experience. As such it has an all-important function: it enables us to solve experiential problems.

In Descartes’ time, this instrumentalist theory of knowledge was formulated and developed by Marsenne and Gassendi [4]. It was then extended by Berkeley and Vico, given strong but unintended support by Hume and Kant; and at the end of the last century, it was applied to physics and science in general by Ernst Mach and to philosophy by Georg Simmel [5]. It was not and still is not a theory popular with traditional philosophers. The idea that knowledge is good knowledge if and when it solves our problems is not acceptable as criterion to those who continue to hope that knowledge, ultimately, will at least approximate a true picture of the “real” world. Karl Popper, who has given a lucid account of the beginnings of instrumentalism [6], has struggled hard to convince us that, though reasonable, it is an unsatisfactory theory. As he reiterates in his latest work:

What we are seeking in sciences are true theories — true statements, true descriptions of certain structural properties of the world we live in. These theories or systems of statements may have their instrumental use; yet what we are seeking in science is not so much usefulness as truth; approximations to truth; explanatory power, and the power of solving problems: and thus, understanding [7].

This suggests that “descriptions,” “explanations,” and “understanding” can indeed capture aspects of “the world we live in.” Whether we can or cannot agree with this statement will depend on how we define “the world we live in.” There is no doubt that Popper intended an objective world, i.e., a ready-made world into which we are born and which, as explorers, we are supposed to get to know. This is the traditional realist view, and Popper does his best to defend it, in spite of all arguments one can hold against it. The realists and the skeptics are once more in the familiar deadlock.

Yet, there is another possibility. “The world we live in” can be understood also as the world of our experience, the world as we see, hear, and feel it. This world does not con-

sist of “objective facts” or “things-in-themselves” but of such invariants and constancies as we are able to compute on the basis of our individual experience. To adopt this reading, however, is tantamount to adopting a radically different scenario for the activity of knowing. From an explorer who is condemned to seek “structural properties” of an inaccessible reality, the experiencing organism now turns into a builder of cognitive structures intended to solve such problems as the organism perceives or conceives. Fifty years ago, Piaget characterized this scenario as neatly as one could wish: “Intelligence organizes the world by organizing itself” [8]. What determines the value of the conceptual structures is their experiential adequacy, their goodness of fit with experience, their viability as means for the solving of problems, among which is, of course, the never-ending problem of consistent organization that we call understanding.

The world we live in, from the vantage point of this new perspective, is always and necessarily the world as we conceptualize it. “Facts,” as Vico saw long ago, are made by us and our way of experiencing, rather than given by an independently existing objective world. But that does not mean that we can make them as we like. They are viable facts as long they do not clash with experience, as long as they remain tenable in the sense that they continue to do what we expect them to do.

This view of knowledge, clearly, has serious consequences for our conceptualization of teaching and learning. Above all, it will shift the emphasis from the student’s “correct” replication of what the teacher does, to the student’s successful organization of his or her own experience. But before I expand on that I want to examine the widespread notion that knowledge is a commodity that can be communicated.

II.

The way we usually think of “meaning” is conditioned by centuries of written language. We are inclined to think of the meaning of words in a text rather than of the meaning a speaker intends when he or she is uttering linguistic sounds. Written language and printed texts have a physical persistence. They lie on our desks or can be taken from shelves, they can be handled and read. When we understand what we read, we gain the impression that we have “grasped” the meaning of the printed words, and we come to believe that this meaning was in the words and that we extracted it like kernels out of their shells. We may even say that a particular meaning is the “content” of a word or of a text. This notion of words as containers in which the writer or speaker “conveys” meaning to readers or listeners is extraordinarily strong and seems so natural that we are reluctant to question it. Yet, it is a misguided notion. To see this, we have to retrace our own steps and review how the meaning of words was acquired at the beginning of our linguistic career.

In order to attach any meaning to a word, a child must, first of all, learn to isolate that particular word as a recurrent sound pattern among the totality of available sensory signals. Next, she must isolate something else in her experiential field, something that recurs more or less regularly in conjunction with that sound pattern. Take an ordinary and relatively unproblematic word such as “apple.” Let us assume that a child has come to recognize it as a recurrent item in her auditory experience. Let us further assume

that the child already has a hunch that “apple” is the kind of sound pattern that should be associated with some other experiential item. Adults interested in the child’s linguistic progress can, of course, help in that process of association by specific actions and reactions, and they will consider their “teaching” successful when the child has come to isolate in her experiential field something that enables her to respond in a way which they consider appropriate. When this has been achieved, when the appropriate association has been formed, there is yet another step the child must make before she can be said to have acquired the meaning of the word “apple.” The child must learn to represent to herself the designated compound of experiences whenever the word is uttered, even when none of the elements of that compound are actually present in her experiential field. That is to say, the child must acquire the ability to imagine or visualize, for instance, what she has associated with the word “apple” whenever she hears the sound pattern of that word [9].

This analysis, detailed though it may seem, is still nothing but a gross summary of certain indispensable steps in a long procedure of interactions. In the present context, however, it should suffice to justify the conclusion that the compound of experiential elements that constitutes the concept an individual has associated with a word cannot be anything but compound of abstractions from that individual’s own experience. For each one of us, then, the meaning of the word “apple” is an abstraction which he or she has made individually from whatever apple-experiences he or she has had in the past. That is to say, it is subjective in origin and resides in a subject’s head, not in the word which, because of an association, has the power to call up, in each of us, our own subjective representation.

If you grant this inherent subjectivity of concepts and, therefore, of meaning, you are immediately up against a serious problem. If the meaning of words are, indeed, our own subjective construction, how can we possibly communicate? How could anyone be confident that the representations called up in the mind of the listener are at all like the representations the speaker had in mind when he or she uttered the particular words? This question goes to the very heart of the problem of communication. Unfortunately, the general conception of communication was derived from and shaped by the notion of words as containers of meaning. If that notion is inadequate, so must be the general conception of communication.

The trouble stems from the mistaken assumption that, in order to communicate, the representations associated with the words that are used must be the same for all communicators. For communication to be considered satisfactory and to lead to what we call “understanding,” it is quite sufficient that the communicators’ representations be compatible in the sense that they do not manifestly clash with the situational context or the speaker’s expectations.

A simple example may help to make this clear. Let us assume that, for the first time, Jimmy hears the word “mermaid.” He asks what it means and is told that a mermaid is a creature with a woman’s head and torso and the tail of a fish. Jimmy need not have met such a creature in actual experience to imagine her. He can construct a representation out of familiar elements, provided he is somewhat familiar with and has established associations to “woman,” “fish,” and the other words used in the explanation.

However, if Jimmy is not told that in mermaids the fish's tail replaces the woman's legs, he may construct a composite that is a fish-tailed biped and, therefore, rather unlike the intended creature of the seas. Jimmy might then read stories about mermaids and take part in conversations about them for quite some time without having to adjust his image. In fact, his deviant notion of a mermaid's physique could be corrected only if he got into a situation where the image of a creature with legs as well as a fish's tail comes into explicit conflict with a picture or with what speakers of the language say about mermaids. That is, Jimmy would modify the concept that is subjective meaning of the word only if some context forced him to do so.

How, you may ask, can a context force one to modify one's concepts? The question must be answered not only in a theory of communication but also in a theory of knowledge. The answer I am proposing is essentially the same in both.

The basic assumption is one that is familiar to you. Organisms live in a world of constraints. In order to survive, they must be "adapted" or, as I prefer to say, "viable." This means that they must be able to manage their living within the constraints of the world in which they live. This is a commonplace in the context of biology and evolution. In my view, the principle applies also to cognition — with one important difference. On the biological level, we are concerned with species, i.e., with collections of organisms which, individually, cannot modify their biological make-up. But since they are not all the same, the species "adapts" simply because all those individuals that are not viable are eliminated and do not reproduce. On the cognitive level, we are concerned with individuals and specifically with their "knowledge" which, fortunately, is not immutable and only rarely fatal. The cognitive organism tries to make sense of experience in order better to avoid clashing with the world's constraints. It can actively modify ways and means to achieve greater viability.

"To make sense" is the same activity and involves the same presuppositions whether the stuff we want to make sense of is experience in general or the particular kind of experience we call communication. The procedure is the same but the motivation, the reason why we want to make sense, may be different.

Let me begin with ordinary experience. No matter how one characterizes cognizing organisms, one of their salient features is that they are capable of learning. Basically, to have "learned" means to have drawn conclusions from experience and to act accordingly. To act accordingly, of course, implies that there are certain experiences which one would like to repeat rather than others which one would like to avoid. The expectation that any such control over experience can be gained must be founded on the assumptions that: (1) some regularities can be detected in the experiential sequence; and (2) future experience will, at least to some extent, conform to these regularities. These assumptions, as David Hume showed, are prerequisites for the inductive process and the knowledge that results from it.

In order to find regularities, we must segment our experience into separate pieces so that, after certain operations of recall and comparison, we may say of some of them that they recur. The segmenting and recalling, the assessing of similarities, and the decisions as to what is to be considered different, are all our doing. Yet, whenever some particular result of these activities turns out to be useful (in generating desirable or

avoiding undesirable experiences), we quickly forget that we could have segmented, considered, and assessed otherwise. When a scheme has worked several times, we come to believe, as Piaget has remarked, that it could not be otherwise and that we have discovered something about the real world. Actually we have merely found one viable way of organizing our experience. "To make sense" of a given collection of experiences, then, means to have organized them in a way that permits us to make more or less reliable predictions. In fact, it is almost universally the case that we interpret experience either in view of expectations or with a view to making predictions about experiences that are to come.

In contrast, "to make sense" of a piece of language does not usually involve the prediction of future non-linguistic experience. However, it does involve the forming of expectations concerning the remainder of the piece that we have not yet heard or read. These expectations concern words and concepts, not actions or other experiential events. The piece of language may, of course, be intended to express a prediction, e.g., "tomorrow it will rain," but the way in which that prediction is derived from the piece of language differs from the way in which it might be derived from, say, the observation of particular clouds in the sky. The difference comes out clearly when it is pointed out that, in order to make sense of the utterance "tomorrow it will rain" it is quite irrelevant whether or not there is any belief in the likelihood of rain. To "understand" the utterance it is sufficient that we come up with a conceptual structure which, given our past experience with words and the way they are combined, fits the piece of language in hand. The fact that, when tomorrow comes, it doesn't rain, in no way invalidates the interpretation of the utterance. On the other hand, if the prediction made from an observation of the sky is not confirmed by actual rain, we have to conclude that there was something wrong with our interpretation of the clouds.

In spite of this difference between the interpretation of experience and the interpretation of language, the two have one important feature in common. Both rely on the use of conceptual material which the interpreter must already have. "Making sense," in both cases, means finding a way of fitting available conceptual elements into a pattern that is circumscribed by specific constraints. In the one case, the constraints are inherent in the way in which we have come to segment and organize experience; in the other, the constraints are inherent in the way in which we have learned to use language. In neither case is it a matter of matching an original. If our interpretation of experience allows us to achieve our purpose, we are quite satisfied that we "know"; and if our interpretation of a communication is not countermanded by anything the communicator says or does, we are quite satisfied that we have "understood." The process of understanding in the context of communication is analogous to the process of coming to know in the context of experience. In both cases, it is a matter of building up, out of available elements, conceptual structures that fit into such space as is left unencumbered by constraints. Though this is, of course, a spatial metaphor, it illuminates the essential character of the notion of viability and it brings out another aspect that differentiates that notion from the traditional one of "truth"; having constructed a viable path of action, a viable solution to an experiential problem, or a viable interpretation of a piece language, there is never any reason to believe that this construction is the only one possible.

III.

When I began the section on communication by talking about the concept of meaning, it must have become apparent that I am not a behaviorist. For about half a century behaviorists have worked hard to do away with “mentalist” notions such as meaning, representation, and thought. It is up to future historians to assess just how much damage this mindless fashion has wrought. Where education is concerned, the damage was formidable. Since behaviorism is by no means extinct, damage continues to be done, and it is done in many ways. One common root, however, is the presumption that all that matters — perhaps even all there is — are observable stimuli and observable responses. This presumption has been appallingly successful in wiping out the distinction between training and education.

As I hope to have shown in the preceding section, a child must learn more than just to respond “apple” to instantiations of actual apple experiences. If that were all she could do, her linguistic proficiency would remain equivalent to that of a well-trained parrot. For the bird and its trainer to have come so far is a remarkable achievement. For a human child it is a starting point in the development of self-regulation, awareness, and rational control.

As mathematics educators, you know this better than most. To give correct answers to questions within the range of the multiplication table is no doubt a useful accomplishment, but it is, in itself, no demonstration of mathematical knowledge. Mathematical knowledge cannot be reduced to a stock of retrievable “facts” but concerns the ability to compute new results. To use Piaget’s terms, it is operative rather than figurative. It is the product of reflection and while reflection as such is not observable, its products may be inferred from observable responses.

I am using “reflection” in the sense in which it was originally introduced by Locke, i.e., for the ability of the mind to observe its own operations. Operative knowledge, therefore, is not associative retrieval of a particular answer but rather knowledge of what to do in order to produce answer. Operative knowledge is constructive and, consequently, is best demonstrated in situations where something new is generated, something that was not already available to the operator. The novelty that matters is, of course, novelty from the subject’s point of view. An observer, experimenter, or teacher can infer this subjective novelty, not from the correctness of a response but from the struggle that led to it. It is not the particular response that matters but the way in which it was arrived at.

In the preceding pages, I have several times used the term “interpretation.” I have done it deliberately, because it focuses attention on an activity that requires awareness and deliberate choice. Although all the material that is used in the process of interpreting may have been shaped and prepared by prior interaction with experiential things and with people, and although the validation of any particular interpretation does, as we have seen, require further interaction, the process of interpreting itself requires reflection. If an organism does no more than act and react, it would be misusing the word to say that the organism is interpreting. Interpretation implies awareness of more than one possibility, deliberation, and rationally controlled choice.

A student's ability to carry out certain activities is never more than part of what we call "competence." The other part is the ability to monitor the activities. To do the right thing is not enough; to be competent, one must also know what one is doing and why it is right. That is perhaps the most stringent reason why longitudinal observation and Piaget's clinical method are indispensable if we want to find out anything about the reflective thought of children, about their operative knowledge, and about how to teach them to make progress towards competence.

At the beginning of this talk, I mentioned that a useful analogy might be found in the teaching of athletic skills. What I was alluding to are the recently developed methods that make it possible for athletes to see what they are doing. Some of these methods involve tachistoscopes and are very sophisticated, others are as simple as the slow-motion replay of movies and videotapes. Their purpose is to give performers of intricate actions an opportunity to observe themselves act. This visual feedback is a far more powerful didactic tool than instructions that refer to details of the action which, normally, are dimly or not at all perceived by the actor himself.

The proficiency of good athletes springs, to a large extent, from the fact that they have, as it were, automated much of their action. As long as their way of acting is actually the most effective for the purpose, this automation is an advantage because it frees the conscious mind to focus on higher levels of control. When, however, something must be changed in the routine, this would be difficult, if not impossible, to achieve without awareness of the individual steps. Hence, the efficacy of visual feedback.

Although the speed of execution that comes with automation may be a characteristic of the expert calculator, the primary goal of mathematics instruction has to be the student's conscious understanding of what he or she is doing and why it is being done. This understanding is not unlike the self-awareness the athlete must acquire in order consciously to make an improvement in his physical routine. Unfortunately, we have no tachistoscope or camera that could capture the dynamics, the detailed progression of steps, of the mental operations that lead to the solution of a numerical problem. Yet, what the mathematics teacher is striving to instill into the student is ultimately the awareness of a dynamic program and its execution — and that awareness is in principle similar to what the athlete is able to glean from a slow-motion representation of his or her own performance. In the absence of any such technology to create self-reflection, the teacher must find other means to foster operative awareness. At the present state of the art, the method of the "teaching experiment" developed by Steffe seems to be the most hopeful step in that direction [10].

The term "teaching experiment" could easily be misunderstood. It is not intended to indicate an experiment in teaching an accepted way of operating, as for instance, the adult's way of adding and subtracting. Instead, it is primarily an exploratory tool, derived from Piaget's clinical interview and aimed at discovering what goes on in the student's head. To this it adds experimentation with ways and means of modifying the student's operating. The ways and means of bringing about such change are, in a sense, the opposite of what has become known as behavior modification.

A large part of education research has been employing a procedure that consists of setting tasks, recording solutions, and analyzing these solutions as though they resulted

from the child's fumbling efforts to carry out operations that constitute an adult's competence. The teaching experiment, instead, starts from the premise that the child cannot conceive of the task, the way to solve it, and the solution in terms other than those that are available at the particular point in the child's conceptual development. The child, to put it another way, must interpret the task and try to construct a solution by using material she already has. That material cannot be anything but the conceptual building blocks and operations that the particular child has assembled in her own prior experience.

Children, we must never forget, are not repositories for adult "knowledge," but organisms which, like all of us, are constantly trying to make sense of, to understand their experience [11].

Most traditional measurements of student learning in mathematics assume that individual differences in student's concepts either vary substantially or are unimportant in their influence on the mathematics studied. . . . In contrast if one assumes that there are a variety of ways of understanding a concept mathematically, individual differences in mathematics become the diversity in students' understandings of concepts or of mathematics itself. The clinical interview provides a means for searching for and exploring these individual understandings [12].

It is not in the least facetious to say that the interviewer's goal is to gain understanding of the child's understandings. The difference between the child interpreting (and trying to solve) a task in the given context, and the interviewer interpreting the child's responses and behavior in the context of the task, is that the interviewer can test his interpretation by deliberately modifying certain elements in the child's experiential field. The interviewer can also ask questions and see whether or not the responses are compatible with his or her conjectures about the child's conception of what is going on. Whenever an incompatibility crops up, the interviewer's conjectures have to be changed their replacements tested again, until at last they remain viable in whatever situations the interviewer can think of and create.

In short, the interviewer is constructing a model of the child's notions and operations. Inevitably, that model will be constructed, not out of the child's conceptual elements, but out of conceptual elements that are the interviewer's own. It is in this context that the epistemological principle of fit, rather than match, is of crucial importance. Just as cognitive organisms can never compare their conceptual organizations of experience with the structure of an independent objective reality, so the interviewer, experimenter, or teacher can never compare the model he or she has constructed of a child's conceptualizations with what actually goes on in the child's head. In the one case as in the other, the best that can be achieved is model that remains viable within the range of available experience.

The teaching experiment, as I suggested before, is, however, something more than a clinical interview. Whereas the interview aims at establishing "where the child is," the experiment aims at ways and means of "getting the child on." Having generated a viable model of the child's present concepts and operations, the experimenter hypothesizes pathways to guide the child's conceptualizations towards adult competence. In order to formulate any such hypothetical path, let alone implement it, the experimen-

ter/teacher must not only have a model of the student's present conceptual structures, but also an analytical model of the adult conceptualizations towards which his guidance is to lead.

The structure of mathematical concepts is still largely obscure [13]. This may seem a strange complaint, given the amount of work that has been done in the past 100 years to clarify and articulate the foundations of mathematics. As a result of that work, there is no shortage of definitions, but these definitions, for the most part, are formal rather than conceptual. That is, they merely substitute other signs or symbols for the definiendum. Rarely, if ever, is there a hint, let alone an indication, of what one must do in order to build up the conceptual structures that are to be associated with the symbols. Yet, that is, of course, what a child has to find out if it is to acquire a new concept.

Let me give you an example. A current definition of number, the sense of "positive integer" stays that it is "a symbol associated with a set and with all other sets which can be put into one-to-one correspondence with this set" [14]. The mention of "put" makes it sound like an instruction to act, a directive for construction, which is what it ought to be. But, in order to begin that construction, the student would have to have a clear understanding of the "set" and, more important still, of "one." Such understanding can be achieved only by reflecting on the operations of one's own mind and the realization that with these operations one can create units and sets anywhere and at any time, irrespective of any sensory signals. That means that, rather than speak of "sets" and "mathematical objects" as though they had an independent existence in some "objective" reality, teachers would have to foster the students' reflective awareness of his or her mental operations, because it is only from these that the required concepts can be abstracted.

The teaching experiment, at any rate, presupposes an explicit model of adult functioning. The experimental part of the method then consists in a form of "indirect guidance" aimed at modifying the child's present concepts and operations (which the experimenter "knows" in terms of the model constructed on the basis of observing the particular child) towards the adult concepts and operations (which the experimenter "knows" in terms of the model constructed on the basis of analyzing the adult procedures). Because the child necessarily interprets verbal instructions in terms of her own experience, the "guidance" must take the form either of questions or of changes in the experiential field that lead the child into situations where her present way of operating runs into obstacles and contradictions. Analogous to the adult who organizes general experience, the child is unlikely to modify a conceptual structure unless there is an experience of failure or, at least, surprise at something not working out in the expected fashion. Such failure or surprise, however, can be experienced only if there was an expectation — and that brings me to the last point I want to make. If I have had any success at all in presenting the constructivist epistemology as a possible basis for education and educational research, this last point will be easy to make and its importance should become obvious.

The more abstract the concepts and operations that are to be constituted, the more reflective activity will be needed. Reflection, however, does not happen without effort. The concepts and operations involved in mathematics are not merely abstractions, but

most of them are the product of several levels of abstraction. Hence, it is not just one act of reflection that is needed, but a succession of reflective efforts — and any succession of efforts requires solid motivation.

The need for motivation is certainly no news to anyone who has been teaching. How to foster motivation has been discussed for a long time. But here, again, I believe, the effect of behaviorism has been profoundly detrimental. The basic dogma of behaviorism merely says that behavior is determined by the consequences it has produced in the past (which is just another way of saying that organisms operate inductively). There is every reason to agree with that. The trouble arises from the usual interpretation of “reinforcement,” i.e., of the consequence that is rewarding and thus strengthens specific behaviors and increases the probability of their recurrence.

There is the widespread misconception that reinforcement is the effect of certain well-known commodities, such as cookies, money, and social approval. It is a misconception, not because organisms will not work quite hard to obtain these commodities, but because it obscures the one thing that is often by far the most reinforcing for a cognitive organism: to achieve a satisfactory organization, a viable way of dealing with some sector of experience. This fact adds a different dimension to the conception of reinforcement because whatever constitutes the rewarding consequence in these cases is generated wholly within the organism’s own system.

Self-generated reinforcement has an enormous potential in cognitive, reflective organisms. (All of us, I am sure, have spent precious time and sweat on puzzles whose solution brought neither cookies, nor money, and negligible social approval.) But this potential has to be developed and realized.

When children begin to play with wooden blocks, they sooner or later place one upon another. Whatever satisfaction they derive from the resulting structure, it provides sufficient incentive for them to repeat the act and to improve on it. They may, for instance, implicitly or explicitly set themselves the goal of building a tower that comprises all the blocks. If they succeed, they are manifestly satisfied, irrespective of tangible rewards or an adult’s comment, for they build towers also in the absence of observers. The reward springs from the achievement, from the successful, deliberate imposition of an order that is inherent in their own way of organizing. To repeat the feat, the tower has to be knocked down. That, too, turns out to be a source of satisfaction because it once more provides evidence of the experienter’s power over the structure of experience.

To some these observations may seem trivial. To me, they exemplify a basic feature of the model of the cognitive organism, a feature that must be taken into account if we want to educate.

From the constructivist point of view, it makes no sense to assume that any powerful cognitive satisfaction springs from simply being told that one has done something right, as long as “rightness” is assessed by someone else. To become a source of real satisfaction, “rightness” must be seen as the fit with an order one has established oneself. Teachers, as well as mathematicians, tend to assume that there exists in every particular case an objective problem and an objectively “true” solution. Children and

students of any age are, therefore, expected somehow to “see” the problem, its solution, and the necessity that links the two. But the necessity is conceptual and it can spring from nothing but the awareness of the structures and operations involved in the thinking subject’s conceptualization of the problem and its solution. Logical or mathematical necessity does not reside in any independent world — to see it and gain satisfaction from it, one must reflect on one’s own constructs and the way in which one has put them together.

Final Remarks

Educators share the goal of generating knowledge in their students. However, from the epistemological perspective I have outlined, it appears that knowledge is not a transferable commodity and communication not a conveyance.

If, then, we come to see knowledge and competence as products of the individual’s conceptual organization of the individual’s experience, the teacher’s role will no longer be to dispense “truth,” but rather to help and guide the student in the conceptual organization of certain areas of experience. Two things are required for the teacher to do this: on the one hand, an adequate idea of where the student is and, on the other, an adequate idea of the destination. Neither is accessible to direct observation. What the student says and does can be interpreted in terms of a hypothetical model — and this is one area of educational research that every good teacher since Socrates has done intuitively. Today, we are a good deal closer to providing the teacher with a set of relatively reliable diagnostic tools.

As for the helping and guiding, good teachers have always found ways and means of doing it because, consciously or unconsciously, they realized that, while one can point the way with words and symbols, it is the student who has to do the conceptualizing and the operating.

That leaves the destination, the way of operating that would be considered “right” from the expert’s point of view. As I have mentioned earlier, a conceptual model of the formation of the structures and the operations that constitute mathematical competence is essential because it, alone, could indicate the direction in which the student is to be guided. The kind of analysis, however, that would yield a step by step path for the construction of mathematical concepts has barely begun. It is in this area that, in my view, research could make advances that would immediately benefit educational practice. If the goal of the teacher’s guidance is to generate understanding, rather than train specific performance, his task will clearly be greatly facilitated if that goal can be represented by an explicit model of the concepts and operations that we assume to be the operative source of mathematical competence. More important still, if students are to taste something of the mathematician’s satisfaction in doing mathematics, they cannot be expected to find it in whatever rewards they might be given for their performance but only through becoming aware of the neatness of fit they have achieved in their own conceptual construction.

Notes & References

- [1] For one, Parmenides, the real world was an indivisible, static whole, for others, Leucippus and Democritus, the real world was a mass of constantly moving atoms.
- [2] Putnam, H. *Reason, Truth and History*. Cambridge, Mass.: Harvard University Press, 1982.
- [3] Translation from Popper, K.R., *Conjectures and Refutations*. New York: Harper Torchbooks.
- [4] An excellent exposition can be found in Popkin, R.H., *The History of Scepticism from Erasmus to Spinoza*. Berkeley: University of California Press, 1979.
- [5] Berkeley has been interpreted in this way by Popper (*op. cit.* Chapters 3&6). Vico's instrumentalism is explicit in his *De Antiquissima Italorum Sapientia* of 1710 (Naples: Stamperia de'Classici, 1856). Hume's analysis of induction is instrumentalist and Kant's program of "transcendental inquiry" into the workings of reason in *Critique* supplied powerful ammunition to the modern instrumentalists; cf. Mach, E., *The Guiding Principles of My Scientific Theory of knowledge* (1910) in S. Toulmin, *Physical Reality* (New York: Harper Torchbooks, 1970) and Simmel, G., *Hauptprobleme der Philosophie* (Berlin: de Gruyter, 7th edition, 1950).
- [6] Popper, *op. cit.*, chapter 3.
- [7] Popper K.R. *Quantum Theory and the Schism of Physics*. Rowan & Littlefield, 1982 (p. 42). (Popper's emphasis)
- [8] Piaget, J. *La Construction du Réel Chez l'Enfant*. Neuchâtel: Delachaux et Niestlé, 1937 (p. 400).
- [9] If this isolating of the named thing or "referent" is a demanding task with relatively simple perceptual compounds, such as an apple, it is obviously much more difficult when the meaning of the word is a concept that requires further abstraction from sensory experience or from mental operations. But since we want to maintain that words such as "all" and "some," "mine" and "yours," "cause" and "effect," "space" and "time," and scores of others have meaning, we must assume that these meanings, though they cannot be directly perceived, are nevertheless somehow isolated and made retrievable by every learner of the language.
- [10] Steffe, L.P. *Constructivist Models for Children 's Learning in Arithmetic*. Research Workshop on Learning Models. Durham, N.H., 1977.
- [11] Cobb, P. & Steffe, L.P. *The Constructivist Researcher as Teacher and Model-Building*. J.R.M.E. 1983, 14 p. 83-94.
- [12] Confrey, J. *Clinical Interviewing: Its Potential to Reveal Insights in Mathematics Education*. In R. Karplus (Ed), *Proceedings of the 4th International Conference for the Psychology of Mathematics Education*, Berkeley, CA. 1980 (p. 400).
- [13] For recent conceptual analyses see Steffe, L.P., von Glasersfeld, E., Richards, J., & Cobb, P. *Children's Counting Types: Philosophy, Theory and Application*. New York: Praeger, 1983.
- [14] James & James, *Mathematical Dictionary*, 3rd ed., Princeton, N.J.: Van Nostrand, 1968 (p.193).

This paper can also be downloaded from the Ernst von Glasersfeld Homepage  maintained by Alexander Riegler.

It is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/2.0/>  or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, CA 94305, USA.

(Preprint version of 12 Mar 2006)