Chapter 8

Putting Learning in Its Proper Place

When Marlene Scardamalia and I started using the term ‘knowledge building’ around 1987, we had never encountered it before.[1] A decade later, ‘knowledge building’ was in common use, but like many terms in popular educational discourse, its meaning had become degraded. It became merely a synonym for learning, a term to be used when one wants to add a constructivist flourish to whatever one is saying about that well-worn topic. We saw it as something different, although it took several years to work out the nature of the difference.

The main burden of this chapter is to establish a workable distinction between learning and knowledge building. This is not a hair-splitting distinction. It marks off two importantly different kinds of activities that can go on in schools. Ideally, they should be complementary. But by failing to distinguish between them, educators fail to do full justice to either one and get into needless turmoil through trying to do both in the same way. Or else they make a distinction but make it in wrongly, producing anomalies like the kind found in many American elementary schools, where the educational program shuts down for weeks on end in order to drill for upcoming tests.

To head off one potential misunderstanding, we must note at the start that learning accompanies all conscious activity. Therefore learning necessarily accompanies knowledge building. But this does not make them the same thing. Learning occurs while setting out garbage, too, but we do not conclude from this that learning and setting out garbage are synonymous. The learning that accompanies everything we do does not figure in the distinction I am proposing. Rather, the distinction is between activities carried out *for the purpose* of learning and activities carried out for this other purpose that we call knowledge building.

For those who have tracked the discussion in preceding chapters about World 2 and World 3, there is a quick way to distinguish learning from knowledge building. Learning is activity directed toward World 2. It is doing something to alter the state of your mind so as to achieve a gain in personal knowledge or competence. Knowledge building is activity directed toward World 3. It is doing something to a conceptual artifact. In Chapter 3 I cited Popper’s list of actions that a scientist might take toward such abstract objects (1972, pp. 140-141). It included thinking of alternatives, thinking of criticisms, proposing experimental tests, deriving one object from another, proposing a problem, proposing a solution, and criticizing the solution. Popper contrasted the kinds of items on this list with psychological states such as knowing, mistakenly believing, or doubting some knowledge object. Those are World 2 phenomena and belong with learning rather than knowledge building.

If you happened upon a scientist or scholar at work and asked, “What are you doing?” you might expect an answer along the lines of the statements in Popper’s list. On occasion, however, you might get a different kind of answer, something like the following:

“I’m trying to learn how to use this new stats program.”
“I’m brushing up on my Spanish for a conference in Caracas.”
“I’m rereading Dewey.”

In these instances, the scientist or scholar has taken time out from normal knowledge building work in order to learn something. It may be learning something essential to getting on with the work, such as using a new software application, or learning something incidental, like some useful foreign phrases, or acquiring background knowledge for no very specific purpose.

In the ordinary run of affairs, the distinction between learning activities and knowledge-building activities is not clear-cut. In reading a journal in your field, you may read some articles just for background knowledge—to “keep up with your field.” Others you may mine for specific information or ideas to use in the advancement of your work. You may start reading with a learning purpose in mind, then notice something significant that causes you to shift into a knowledge-building mode. This is no cause to reject the distinction. We often have mixed motives, and we often try to make one action serve more than one purpose. This does not detract from but rather adds to the value of distinguishing one purpose from another.

Distinguishing Knowledge Building from Learning in the Classroom

In the preceding section we distinguished learning from knowledge building in the activities of people whose main work is knowledge building. It becomes trickier but more important to our purposes to distinguish them in the activities of people whose main work is learning—namely, students. When we move into the schooling context, a further distinction becomes relevant: the distinction between being engaged in learning and being engaged in a learning activity.

Consider a grade 5 class discussing the brain, students having read something about it or seen a video. Certain propositions are advanced by students and contested by others. Occasionally a student raises a question. Sometimes the teacher will answer, but more often will turn it back to the class. The teacher does quite a bit of calling on students who have not contributed: "Mandy, what do you think about that?" "How many of you agree?" The teacher will also ask questions to turn the discussion to neglected points: "Does the brain have anything to do with breathing and heart beats and things like that?" If the point is a minor one, the teacher may accept the first more-or-less adequate answer or may quickly supply an answer if none is forthcoming. On other, more major points, the teacher may solicit a number of opinions and will not leave the point until an acceptable answer has been established and the dissenting opinions have been dealt with.

What is going on here? From the point of view of the teacher, this is what is customarily and appropriately called a ‘learning activity.’ Whatever other values might be attributed to it, the success of the discussion will ultimately be judged by its impact on the beliefs and knowledge of individual students. Thus, the locus of the discourse for the teacher is World 2, the contents of individual minds. Among the students, however, there are likely to be a variety of points of view on what is going on. For present purposes let us ignore those students who do not actually participate, who may be daydreaming or dozing, and those who see the activity as an opportunity for clowning around or bullying, and consider only those students who appear to be seriously engaged in the activity. Though outwardly similar, they may have importantly different points of view on the discussion. Research by Evelyn Ng (Ng & Bereiter, 1991)[2] distinguished three kinds of goals that distinguished three different kinds of student orientation:

1. Task-completion goals. Students focusing on task-completion are engaged in the learning activity at a behavioral level but are not cognitively engaged. They may enjoy the social give-and-take or they may just be trying to do what is expected of them. They are learning, of course (that goes without saying), and, if the activity has been well designed and conducted, they

may even be learning what the teacher has intended. But the learning will have been incidental. In Popperian terms, the focus of these students is World 1, the external world of people, things, and actions.

2. Learning goals. These goals distinguish students who are purposefully engaged in learning. They have an idea of the educational purpose of the activity and they adopt it. Their interest in learning may only extend to doing well on the next exam or it may have deeper roots; in either case, their focus is on Popper’s World 2. They are concerned with the content of the individual mind—specifically, their own. Thus the point of view of these students is complementary to that of the teacher. They are trying to learn what the teacher is trying to teach.
3. Knowledge building goals. These goals distinguish students who are actively engaged with problems beyond the immediate situation. For these students, the teacher-led discussion about the brain is a real discussion, whose purpose (as a fifth-grader is likely to conceive of it) is to arrive at the truth. In other words, this is World 3 discourse dealing with such World 3 objects as theories and explanations, carried out in a mode of critical inquiry. These students also learn, but like students of the first type, their learning is incidental—only, in this case, the learning is incidental to knowledge building.

Despite their outward similarities, these students may be said to live in three different worlds as far as the learning activity is concerned. Conveniently, their worlds have names: World 1, World 2, and World 3. Most teachers who conduct discussions of the kind described probably hope to have students engaged with the problem—hence, functioning in World 3. However, because conventional theory of mind does not distinguish between learning and knowledge building, they will not recognize any fundamental difference between the second and third types of students.[3] Because their professional orientation is to World 2, they design and manage activities so that they are most suited to students of the second type. Students of the third type, those with a knowledge-building orientation, are likely to find the teacher's approach somewhat frustrating. There is too much opinion sampling, too much wrapping things up into tidy bundles and moving on to the next topic, too much repetition and not enough working out of implications. In order to pursue their own World 3 interests, students may carry the discussion on after class, among themselves or with the teacher; and that is when the real discussion will take place—discussion that is less school-like and more like the discussions adults have who are engaged in efforts to advance knowledge.

Sharpening the Distinction Between Learning and Knowledge Building

The literature on instructional approaches is full of comparisons, often in tabular form, between the author's favored approach on one side and a caricature of conventional instruction on the other. They commonly go like this (I have put together bits of several tables so as not to impugn anyone in particular):

Old	New
Knowledge transmission	Knowledge construction
Memorization	Reasoning
Teacher-directed	Learner-centered
Competitive	Collaborative
Tightly scheduled	Opportunistic
Fact-centered	Idea-centered
Etc....	Etc....

Such comparisons are uniformly unenlightening. It is not that they are inaccurate or that the 'old' is a harmless straw man. It is just that such comparisons appeal only to those who are already sold on the 'new' and for them they obscure differences that matter and bury real problems under a fluff of self-congratulatory clichés.

Knowledge building, of course, belongs on the 'new' side of the table, but by no means does everything on the 'new' side involve knowledge building. In this section I want to compare two approaches to student inquiry in science that are very close in many respects, that may be said to share a common philosophy, but which nevertheless diverge in ways that help to sharpen the distinction between learning and knowledge building. More detailed accounts appear in the same issue of *The Elementary School Journal* (Anderson, Holland, & Palincsar, 1997; Bereiter, Scardamalia, Cassells, & Hewitt, 1997).

Anderson, Holland, and Palincsar (1997) analyzed the behavior of a small group of elementary school students over the course of four days during which they worked on the following task. The students had previously observed and taken notes on a demonstration illustrated in Figure xx. A clear liquid in the bottom part of the apparatus was heated, giving off an invisible gas that rose to the top of the apparatus, where it was cooled by ice and condensed back into a clear liquid. The task for the students was to explain the phenomenon and to present their explanation to the class in the form of a poster with drawings and text, accompanied by a demonstration using models of molecules. The pedagogy here included two elements that are common in science education at all levels:

1. The demonstration of a puzzling phenomenon. Real scientific research often starts with puzzling observations. Teachers use such demonstrations to arouse curiosity, to bring forth conjectures, and to stimulate argument and further inquiry (Hunt & Minstrell, 1994).
2. A material product—what is known in the world of research contracts as a "deliverable."In school, "deliverables" are likely to be called "projects." Typically they take the form of a report, but in this multimedia age many more dramatic kinds of presentation are often encouraged. The pedagogical assumption behind requiring them is that "[p]roducts and performances are important sources of evidence about whether students are helping each other to learn, and they play an essential role in building confidence and motivation, especially for students who have traditionally been unsuccessful in school science"(Anderson, Holland, & Palincsar, 1997, p. 379).

The story that Anderson et al. tell is one of science learning getting overwhelmed by social complexities. The students' concerns are almost wholly with their poster and its presentation, with who gets to perform, with looking good and not making fools of themselves. The teacher gets drawn into these concerns as well. There is the problem of a bright, high-status girl continually upstaging a lower-status boy who is actually the one most engaged with scientific issues. The authors' conclusion is that if there is to be scientific literacy for everyone, learning will have to be viewed within a sociocultural framework. It would be difficult to dispute that conclusion, but the story can equally well be understood as bringing into question the two pedagogical elements noted above.

1. The problem that the students were supposed to address did not arise from their own wondering and inquiry but was presented to them. Although the evaporation and condensation demonstration may have aroused some initial interest, the problem of explaining it clearly did not weigh heavily on their minds. When they embarked on the

assignment the first thing they did was allocate different tasks to different group members, and the task of producing the explanation fell to the high-status girl. The only discussion occurred when they were practicing their presentation and she coached the boy in presenting his part correctly. It seems fair to conclude that what was presented as a scientific problem was not really a scientific problem for the students but only part of a task to be accomplished.

2. The task of preparing the deliverables, far from serving as a vehicle for science learning, seems to have obliterated whatever scientific inclinations the students might have had. Although it may have played a role, as claimed, in building confidence and motivation, the confidence and motivation do not appear to have had anything to do with science. Teachers I have talked to report similar problems, but so deeply entrenched is belief in the necessity of a tangible product that it does not occur to most of them to look for alternatives. Anderson et al. anticipate the objection "that students should be held accountable for helping each other to learn rather than for products and performances" (p. 379), but they say that the two are indissociable. In the next example I will try to show that this is not true.

The preceding example is an unusually well worked out instance of a common approach to elementary science. Had the authors been less attentive to what actually happened and less candid in reporting it, they might have put this example forth as a model for teachers to follow. The next example is too unusual, too seredipitous to serve as a model for practice. Instead its value is conceptual. It enables us to imagine the possibility of science learning that is not driven by presented problems and by deliverables. This example is sufficiently important for giving a concrete idea of knowledge building in school that I will quote at length from the published description that appeared in Bereiter, Scardamalia, Cassells, & Hewitt (1997, pp. 334-337):

> [W]e will describe one discussion that went on over a period of almost 3 months among 17 sixth-grade students.... The discussion was carried on by means of notes entered into CSILE, a networked computer environment designed to support knowledge-building discourse (Scardamalia & Bereiter, 1994; Scardamalia, Bereiter, & Lamon, 1994). The format used in this case was the discussion note (Scardamalia, Bereiter, Hewitt, &

Webb, 1996), which is an extended note with entries by different students appearing one after the other in chronological order. Individual entries are labeled according to ‘thinking type’: P (Problem), MT (My Theory), INTU (I Need to Understand), NI (New Information), C (Comment), and WWHL (What We Have Learned). Any entry labeled INTU may be used as beginning of a subdiscussion. The work we shall discuss was actually a subdiscussion initiated by a student, branching off from a discussion that had been initiated by one of their teachers. The subdiscussion was titled “About Growing” and it comprised 179 entries.

Initially, the discussion consisted mainly of students expressing their personal interests and concerns with growing. The adolescent growth spurt was the dominant topic, with some students concerned about when it would start for them and others concerned about when it would end. “Well, as you know, I’m one of the shortest people on this team,” says one student. “I hardly grow any over a year’s time.... It actually feels like I’m shrinking. Everyone else is growing....”[4] “I know that I am going to be a tall person but in a way I don’t want to be a tall person,” says another. There is a pervasive sense of being under the control of forces they don’t understand and can do nothing about. At the same time, several of them wonder what it must be like to be done growing, and they ask their teacher whether it feels weird to be the same size year after year.

So far, it is a worthwhile discussion and the students are amazingly supportive of one another, but it is not a discussion that anyone would call scientific. A desire to move beyond expression toward understanding emerges quickly, however, and it takes the form of scientific conjectures. One student wonders if growth rate is hereditary, another wonders whether trees have growth spurts; one suggests that “people stop growing when the clock inside their head says they should,” another that growing stops when the body runs out of material to grow with.

The students are not merely expressing opinions for the sake of being heard. There is a clear expectation of working together to figure things out. Entry number 33 reads:

> Everyone:
> If you have a theory on mine and Jake’s INTU, please write it and maybe we can combine them and have more new learnings.

...Whereas, initially, growth was discussed only in terms of reaching adult height, new information brought into the discussion greatly extends the range of issues related to understanding growth. For example:

- Hair and nails keep growing after growth in height stops. So in some respects we never stop growing.
- We also keep growing in knowledge; but it is said that babies are born with all the brain cells they will ever have. So what kind of growth is this?
- Elderly people are said to get smaller. Is this a reversal of growing or something else?
- Growth is not uniform. The wrists grow more than the rest of the arm.
- Skin cells on humans and leaves on trees die and are replaced. Is this growth?
- Trees and other plants also appear to grow rapidly when they are young and then to slow down or stop growing. Is this the same process as in animals?

Throughout the discussion, there is a concern with... empirical testability: How do experts know such-and-such? How could we find out? One student raises an explicit question of this kind:

> I wonder how doctors or other people that have knowledge about growing know how tall you might be and what age you might stop growing. The doctor has predicted I'll be 5'4". He also thinks that my sister is finished growing and he used to think that she was going to be very tall. Do they look for patterns?

Heredity is an issue that comes up repeatedly. Will one end up being the same size as one's parents? One student reported having read that height of parents is not a major factor. The students decide to conduct a survey. A questionnaire is constructed, asking for students' and parents' heights. Grandparents are included as well, to check on the question of shrinkage in old age. Although the students voice concerns about the validity of the data (some of the reported heights are suspect), they conclude tentatively that there is not much relationship between the heights of students and their parents. Actually, with some students into their adolescent growth spurt and others not, there is too much random variability for any strong relationship to appear. (Height at ages 12-13 is not highly correlated with one's own adult height, let alone that of relatives.) Although, understandably, the students do not recognize this problem with their research, they do figure out how to look for a relationship—by seeing whether students who report above average height for their age report above average heights of their parents.

...The last entries in the discussion consist of students' reflections on what they have learned. These are entries under the WWHL (What We Have Learned) label, a standard part of the discussion note format that was used.The following are excerpts from one student’s reflective summary:

> We have worked incredibly hard on this note. It has been growing and growing.... Discussion and learning has changed greatly over the few months. We have had great ideas, some of which have been rather strange.... But, even though these were strange, they were interesting, and well thought about.
>
> At the start of this note, we talked about what it was like to stop growing. And, me being vertically challenged, as I prefer to be called, was asked a lot of questions. I was asked about my brother’s height, how much I grew a year, and things like that. We thought about what it would be like to stop growing.
>
> Many adults wrote in, and students asked their parents what they thought. But, as most learning goes, we weren’t happy. We wanted to know “why” we stopped growing. This led to many people studying things such as the pituitary glands, and hormones, and learned about all that stuff that we consider rather gross. But we learned a lot, which helped us understand why we stopped growing.

Students remark on the many facts and ideas they have acquired, but an underlying theme is that of conceptual change, a significant deepening and expanding of their conception of what growth means:

> Many of my theories have changed since I began working on this note.... I thought that you stopped growing. But, as my research continued, I learned that you never stop growing. Your hair, nails, and mind is always growing and expanding. You grow in maturity and of course you grow in weight. I thought that tendons just stretched and I didn’t know much about bones and what they are made of.

The students also seem somewhat awestruck by what they have accomplished, and they are eager to repeat or continue the experience. One student remarks, “It is amazing to think that such a complex discussion note evolved from a single INTU, about how people feel at different heights.” Another says, “I feel that we did a good job making theories but that we needed more time to get new learnings. I would still like to continue my research on this and write in next year with more new learnings. I think that we should learn stuff... over the summer and share it with the others in the fall. We can continue to write letters and make phone calls to experts though. I think that this research has been fun and we should pick it up next year or start a new

> one about the brain and its growth." Another concludes, "I hope we can all get together next year and try to get more people involved."

There are too many differences between this example and the first one to make detailed comparison profitable. The contrasts I want to highlight pertain to the two pedagogical features discussed earlier.

1. Whereas in the first example a problem was presented at the outset that continued to be the only scientific problem to be dealt with, in the second example the problem developed over time, became elaborated and increasingly scientific. A number of scientific subproblems emerged so that by the end a fairly comprehensive effort to understand biological growth was in evidence.
2. Whereas the focus of activity in the first example was on producing a poster and presenting it, there were no deliverables in the second example. The only tangible residue was the trace of the dialogue itself. Of the justifications that Anderson et al. offer for "products and performances," this computer trace provided one: a source of evidence about children's contributions to the collective effort. As for building confidence and motivation, I think any reader of the complete transcript would agree that such building did occur and in ways relevant to science. Motivation was, in fact, high enough that some students proposed continuing their inquiry over the summer and resuming the discussion in the fall. What kept motivation high, I believe, was the thrill of the chase—the excitement of making progress in pursuit of understanding a compelling phenomenon.

I think it is clear that the second example is one of knowledge building but that the first is not. But if the first had worked better, if the students had been more involved in explaining the demonstrated phenomenon and less in their presentation, would that not have been knowledge building too? To a very limited extent, perhaps. The only knowledge to be constructed was an explanation of the behavior of an outlandish piece of laboratory equipment. Surely, the value in trying to get the students do so is not in the knowledge they construct but in what they are expected to learn in the process of doing so. When the knowledge actually to be constructed is trivial or fantastic, we have the clearest indication we can hope to get that we are dealing with a learning activity that is not a knowledge building activity.

Nevertheless, when teachers, impressed by examples like the second one, decide that knowledge building is for them, their immediate inclination is to do something like the first example, perhaps augmented by having students state their theories or predictions. Partly it's that they do not trust that students can be sufficiently motivated just by the pursuit of understanding, and so they feel the need for concrete activity that will yield a tangible result. I believe it is also, however, a matter of not yet having a clear distinction between knowledge building and learning, regarding learning as merely an inferior form of knowledge building. Lacking the distinction, they do things that might make sense as part of a learning activity, but they do not go all the way with them, and so they end up with something that is neither good knowledge building nor good learning.

Puzzling and intriguing demonstrations can play a valuable part in teacher-led learning activities. Not long ago I saw a brilliant 20-minute lesson on air pressure put on at a science museum. The presenter went through half a dozen demonstrations with variations, showing things collapsing or exploding, liquid levels rising or falling, and I can't remember what else as air was introduced or removed. She kept firing questions at the audience of drop-in children, getting them intensely involved in predicting and explaining, and I am pretty sure that by the end of it the great majority of them understood that air has weight and that there is no such force as suction.

I think what happens when people try to adopt a constructivist pedagogy without a clear distinction between knowledge building and learning is that they try to take parts of what they recognize as good learning activities and turn them over to the students. Sometimes it works, but in science, especially, the students are in no position to assume responsibility for the stock-in-trade of old-style pedagogy, such as was on display in the science museum session. In the first place, no single demonstration can carry learning very far. There needs to be a series of demonstrations, and they need to be tied together so that they build up a general understanding that extends beyond the particulars. Students cannot be expected to do that. It requires somebody who already understands what the demonstrations are supposed to teach, who can draw out their implications and reveal the conceptual thread that ties them together.

This does not mean that experimental science has to be didactic. Jim Minstrell (Hunt & Minstrell, 1994) has developed an approach to teaching physics that relies heavily on the students’ own efforts to explain, to deal with inconsistencies, and to resolve conflicts among their theories. Minstrell starts with a problematic experimental situation, just as in the Anderson et al. example. In the cited account, a unit starts with students being challenged to predict and explain what will happen to the apparent weight of an object on a spring balance when a vacuum is created around it. This is followed by a series of related demonstrations or experiments, which require students to modify and elaborate their explanations. But the demonstrations are chosen by the teacher, who also plays a directive role in the students’ deliberations:

> The important thing is that the students are encouraged to extract general principles from a variety of specific contexts. They are continually asked, “What can each experiment tell us that might relate to all of the other situations, including the original benchmark problem?” New issues are opened up as well, touching on such concepts as the “stickiness” of water, and the apparent “sucking” by vacuums. In addition to encouraging additional investigation of issues, the teacher can help students note the analogical similarity between what happens to an object submerged in a container of water and what happens to an object submerged in the “ocean of air” around the earth. (Hunt & Minstrell, 1994, p. 60).

Is this knowledge building? I expect that when Minstrell is in charge it is, but that in other hands it might become only a learning activity. It might also fail to materialize as knowledge building if the students proved unable to generate the central ideas themselves and the teacher had to end up providing all the answers. The crucial issue is whether the students are working in World 3. If they are, then they are doing knowledge building, regardless of how active a role the teacher plays in their World 3 work. The students in the evaporation-condensation example were not working in World 3, although they were working fairly independently. They were working in World 1. That is how the activity was set up. Its central objects were the physical demonstration, the poster the students were producing, and the physical models of molecules that they were to use in their presentation. What happened in World 2, their minds, was an incidental result of their work in World 1. World 3, the world of ideas, did not figure in the activity. In the growth example, World 1 was continually referred to, and there was the overt collection of questionnaire information from other people; but the main focus of work was on ideas—on producing and refining explanations and bringing new information into the process. This they did with little adult help, whereas in the Minstrell example the teacher was continually involved; but the focus of the students' work in both cases was on ideas. World 1, in the form of external information sources, entered the picture only insofar as it was relevant to the work with ideas.

The Three Pillars of Conventional Pedagogy

Schools are often called upon to teach something new: environmental issues, diet, anti-racism, drugs, street survival skills, Afro-American history, computer literacy, phonics (what is old sometimes becomes new again), problem solving—the list grows endlessly. Yet schools manage to incorporate this endless stream of new requirements without disintegrating. How do they do it? As I make it out, they do this by relying on a standard set of reductive moves which convert the new challenge into something they already have the tools to handle. These three reductive moves are so heavily relied upon that it is not altogether fanciful to call them the pillars of conventional pedagogy. They are

- reduction to subject matter
- reduction to activities
- reduction to self-expression

Reduction to subject matter means converting the subject into propositions or procedures that can be directly taught and tested. Reduction to activities is self-explanatory and is far and away the method of choice among modern teachers, just as reduction to subject matter is the old-fashioned choice and the one teachers tend to revert to if they are held accountable for students' learning. Reduction to self-expression is more daring but has long been the ideal in early childhood education. The educational objective is translated into some regulated way of letting students follow their own inclinations—in writing, speaking, making, or doing. It requires structuring the situation, socially or by means of physical arrangements, so that students will spontaneously undertake activities that are expected to produce the desired learning (Weber, 1971, p. 109).

Of course, these three reductive approaches work much of the time. Otherwise schooling in the forms that we know it would be a complete failure. But they are also responsible for much of what is useless and silly in formal education. Here are some notable examples:

1. The reduction of phonics to subject matter. Phonics is a skill that all normal children acquire in the course of learning to read. It is the ability to 'sound out' unfamiliar words according to their spelling, hopefully getting close enough to the actual pronunciation to recognize the word. Its main value, obviously, is in the early stages of learning to read, when children know by sound many words that they do not yet know by sight. In the 1930s the teaching of phonics fell out of favor in the U.S., but in the 1960s public pressure forced textbook publishers to bring it back in. However, it was brought back as a separate body of subject matter. Instead of being taught as a way to identify words, it became a pointless kind of analysis carried out *after* words had been identified (circle the words that contain the 'sh' sound, for instance). Instead of being concentrated in the first weeks of learning to read, it was drawn out over six years! As subject matter, as a body of knowledge *about* written language, phonics is ludicrous, and linguists delighted in pointing out how it oversimplifies and distorts the relation between spelling and pronunciation (Smith, 1971). As a *way to help children get started in reading*, the teaching of phonics has demonstrable virtues (Adams, 1990), but that conception of it has been largely lost through reduction to subject matter.

2. The reduction of science to quiz games. Not one but two projects funded by the National Science Foundation purport to teach science by having students create their own science-related computer games (Fischer, 19xx; Yarnall & Kafai, 1995). The 'constructivist' appeal of this approach is obvious: Instead of having students play games concocted for them by science educators (of which there are some good ones), let them design and construct their own. The result, however, is that students create games that quiz one another on isolated facts. In one case, that is the only kind of game the software permits; in the other, students fall into it naturally because any other kind of game would far exceed their programming capabilities. Thus the kind of science teaching that constructivists universally condemn

becomes institutionalized in activities intended to promote constructivist learning. This is what happens not infrequently when ‘project-based learning’ takes the form of reduction to activities.

3. The reduction of literature study to self-expression. The teaching of literature often succumbs to reduction to subject matter. It is much easier to teach facts about authors, history, and literary features than it is to promote the kind of intense engagement with literary works that educated people recognize as the real point. There has been a strong reaction against this kind of reductionism, however, sparked by Louise Rosenblatt’s influential essay, “What Facts Does This Poem Teach You?” (1980), and sustained by a turn in literary theory that emphasizes what readers bring to literary experience (xx). Questions on the order of “Why did Amy pick the marigolds?” have begun to disappear from textbooks, replaced by questions of “How did it make you feel...?” The result, however, has been to replace one kind of reduction by another. Discussions carried out under the aegis of ‘reader response theory’ are no longer about the book but about personal experiences and concerns brought to mind by the book. Among other things, this makes it possible for students to participate without having actually read the work. They need only have picked up a few clues as to what it is about. I have seen animated discussions going on in classes where none of the students could read well enough to have grasped the content of the book, although they had become adept at spotting keywords and inferring from them the topics of text passages. In these cases the literary work serves the same purpose as the exercises and games that encounter group leaders use to stimulate self-expression. It creates a topical framework within which students can do their own thing.

These are not isolated examples. Reductive pedagogy is the norm—so normal, in fact, that departures from it are likely to be perceived as lapses. The standard reading lesson, as carried out in most North American schools up through the sixth grade, is entirely focused on extracting information from the story or article being read. As observed by Pearson and Gallagher (19xx), the standard method is for the teacher to quiz students on text content, continuing to ask the same question until a correct answer has been obtained. If no one answers the question, the teacher does so (or distorts a student’s response into the desired answer). This is called teaching reading comprehension. But in fact no teaching of reading comprehension skills is going on (Durkin, 19xx). Instead, what we have is reduction to subject matter, the subject matter in this case being whatever the reader selection happens to be about. When the selection is fiction, which it usually is, treating it as subject matter becomes ridiculous. Great attention is lavished on getting every detail of the story right, as if it were something one needed to remember throughout life. Yet if you try to drop this reductive practice in favor of an approach that gets students doing the kinds of things that good readers do in comprehending what they read, there will be complaints that "there's no comprehension."[5] In the traditional teaching of writing, reduction to activities has been the norm. The focus is on producing a good written product rather than on learning to write well. This distinction, I have found, is virtually incomprehensible not only to teachers but to language arts specialists. The result is that any practice that improves the apparent quality of the product—such as assigning a livelier topic—is welcomed, without regard to what effect it may have on learning. Recently there has been a movement away from the standard writing assignment, which sustains this reduction to activity, but the replacement is reduction to self-expression. Students are supposed to find their "voice" and use writing as a medium for rendering their life experiences more meaningful and for expressing their feelings. This is fine as far as it goes, but it reduces writing to only one of its many functions.

Reduction to activities is perhaps unavoidable in schools. Teachers have to see to it that some kind of activity is going on all the time, and it cannot always be well calculated to achieve a learning objective. So every teacher relies on a repertoire of activities to fill up a space of time in a way that conforms in a general way to what is supposed to go on in school. I have never seen an exception, even in the most no-nonsense of school programs. In their proper place, activities, subject matter, and self-expression are all legitimate. What has happened, however, is that reductive practices have become elevated to pedagogical virtues. Reduction to activities becomes hands-on or project-based learning. Reduction to subject matter becomes direct instruction or back-to-basics. Reduction to self-expression becomes 'reader response theory' and 'process writing.' As a result, reductive practices are hidden behind respectable principles, and the principles themselves become degraded.

At the extreme of reductionism, educational thought descends into word magic. Word magic is the belief that things called by the same name are substitutable one for another. In education, word magic appears as the belief that if you teach people a certain thing, the learning will automatically transfer to everything else called by the same name. In one corporate training program, trainees work in teams to climb a high wall. The express purpose is not, of course, to make them into mountain climbers; the wall climbing is supposed to teach them to work cooperatively on arranging business deals. Any close analysis of the two tasks would show that they have almost nothing in common except at the most abstract level. They both involve goals and a need for cooperative effort to achieve the goals. But the kind of cooperation, the obstacles to cooperation, the kind of communication needed in order for cooperation to occur, the emotions, the motives, the clarity of purpose—all of these differ radically.

If the three reductionist pillars of pedagogy were removed, schooling as we know it—in both its traditional and its ‘progressive’ forms—would crash to the ground. But there are two other pillars that could support a nonreductive pedagogy. These are

- enlisting students in the pursuit of learning objectives
- converting learning objectives into knowledge building objectives

The first is the route of intentional learning (Bereiter & Scardamalia, 1989). In terms of the three kinds of goals discussed earlier, it involves shifting students from task-completion goals to learning goals. The second route involves the further shift to knowledge building goals.

These are not alternative routes. Because students engaged in knowledge building are also learning, they can be simultaneously attending to both kinds of goals. In one Ontario classroom students were engaged in a collaborative knowledge building project concerning preColumbian civilizations. Through formulating problems, proposing theories, and revising them in the light of information, the students had compiled a substantial network of notes in a CSILE[6] database. The teacher then inserted into the database a graphic note which itemized the learning objectives that the Ontario Ministry of Education and Training had set out for this unit of subject matter. He had the students link their notes to appropriate objectives and explain how the notes fulfilled the objectives. When this task was completed, the students ended up with notes left over, from which they derived two additional learning objectives that they considered relevant to understanding preColumbian civilizations. Like most curriculum guidelines, those of the Ministry supplemented the objectives with suggested learning activities for achieving them. Normally—in keeping with the practice of reduction to activities—the students would encounter the activities but would have no contact with the objectives. In this case the activities were by-passed in favor of engaging the students directly with the learning objectives and, more importantly, with the problems that made the subject worthy of study in the first place.

Direct versus Indirect Learning Activities

Teachers must continually choose between teaching something directly or leaving it to be acquired incidentally through indirect learning activities . This ought to be a strategic issue, not an ideological one. Strategically, answers will vary. Ideologically, there is a press to have a uniform answer for all occasions. Either everything must be engineered or everything must be acquired naturally. Combined with a failure to distinguish learning from knowledge building, pedagogical ideologies of all sorts tend toward fanaticism and superstition. The teaching of reading provides the most extreme examples. I have heard someone stand up at a meeting of the Reading Reform Foundation and maintain without challenge that no one could become a good reader unless they were taught phonics (thus branding as illiterate everyone who, like me, was brought up on *Dick and Jane*). A leading advocate of whole language, on the other hand, reportedly called for a stake to be driven through the heart of someone who dared to write a book advocating the teaching of phonics. Fortunately, this level of fanaticism has not taken hold in other areas of education, but there are ideological purists everywhere. I remember the dismay of a group of reformers who were promoting an inquiry approach to mathematics, when they discovered that their prize teacher was starting off each lesson with mental arithmetic drill. My response, however, was “Good for her!” She recognized that there was more than one side to competence in mathematics, and that one side might need one approach and another another.

When the issue is approached strategically, it seldom turns out that the decision is a simple either / or. Vocabulary acquisition provides a nice illustration. Children of school age learn thousands of new words each year (Anglin, 1993). The most ambitious of vocabulary instruction programs aim to teach only a few hundred. To think of directly teaching thousands of words a year is unrealistic. It would eat up too much instructional time and might prove impossible: The rate of forgetting would probably overtake the rate of learning. But a purely laissez-faire policy is not adequate either. Children differ greatly in vocabulary knowledge, vocabulary figures prominently in achievement and scholastic aptitude tests, and mastering vocabulary is an important part of learning in every academic discipline. “Encourage wide reading,” is one answer, but only a partial one. Much surely depends on what and how a child reads. And no amount of normal reading is likely to convey the meaning of *parallelogram*, the literal meaning of *osmosis*, or the difference between *heat* and *temperature*. But understanding parallelograms, osmosis, and the relation between heat and temperature is not simply a matter of vocabulary learning. It is a matter of developing important understandings in mathematics and science. And so vocabulary acquisition becomes a subordinate issue in the larger problem outlined in the preceding chapter—the problem of becoming familiar with the abstract artifacts making up various domains of knowledge. In that context, the question is whether special attention needs to be directed to learning key words, or whether that will come about as a natural consequence of knowledge-building activity in the domain. The answer is to be found in data, not pedagogical philosophy.

The Question of “What Is Learned?”

With all learning activities, the big question is what is learned. Where direct teaching is concerned, educators are all familiar with the fact that what is supposedly taught is not necessarily learned. That is why a careful teacher will keep checking to see what is getting across and what is being retained and will keep reteaching or trying different angles until satisfied that the desired learning has taken place. This is the great strength of direct learning activities and is the reason that even teachers ideologically opposed to direct instruction will resort to it when it really counts—for instance, in teaching emergency procedures, computer operations, or classroom routines. But the literature on misconceptions shows how fallible the process is. The instructor teachers Darwin but the students learn Lamark. Case (1985a) has shown that even in arithmetic, that bastion of direct instruction, children will adopt workable procedures that are different from what they have been taught.

When it comes to indirect learning experiences, however, the question of what is learned becomes even less tractable. Word magic and wishful thinking often serve to fill the voids. Hands-on experiences in science and mathematics are supposed to result in deeper understanding, but evidence is lacking and just how the understanding is supposed to come about is seldom explained. In arithmetic, for instance, it is known that children who are directly taught algorithms often learn them as rote procedures and get them wrong, producing senseless answers. A lively industry has grown up around the production of ‘manipulatives,’ clever systems of rods, blocks, or other objects, with which it is possible to enact concretely the operations involved in multidigit addition, subtraction, and multiplication. The hope is that working with these palpable things will make the corresponding symbolic operations (such as regrouping or borrowing) more meaningful. The reality, according to studies by Resnick and her colleagues (Resnick & Neches, 1984; Resnick & Omanson, 1987), is that many students can become proficient in manipulating the concrete objects and in performing the symbolic operations yet never see the connection. In science, the gap between what students are supposed to be getting out of hands-on experiences and what they actually get out of them is often even greater (Roth, Anderson, & Smith, 1987).

With the rising concern about standards and accountability, educators given to indirect learning activities have had to become more definite about what students are supposed to learn from such activities. This has forced word magic to a new level. Whatever words may be used to *describe* an activity are translated into *outcomes* of the activity. Thus, an activity in which children team up to make paper airplanes, measure and record how far they will fly, and criticize each other’s designs may be claimed to teach cooperation, technology, problem solving, force, gravity, measurement, scientific method, subtraction, and critical thinking. ‘Objectives-Based Education’ is a name given to such word play. Although it may also refer to firmer stuff, in its degraded versions ‘OBE’ amounts to a scheme in which the activities *become* the objectives. Fairly expensive software systems are being sold to schools that permit lesson plans and activity descriptions of all kinds to be entered and linked to lists of mandated objectives.

It would be futile, however, to insist that claims about the outcomes of indirect learning activities should be based on evidence. Activities keep being invented and modified at a far greater rate than research could hope to keep up with, and seemingly minor modifications can make a big difference (see, for instance, Scardamalia, Bereiter, Hewitt, & Webb, 1996). Rules of thumb are needed. A rule of thumb implicit in a great deal of computer-based educational fare (and sometimes stated explicitly by software designers) is, "As long as they're having fun, they're learning." Inasmuch as we can be confident that students are also learning when they are not having fun, this is not a very useful rule. It does not help in identifying *what* is likely to be learned.

From the cognitive instructional research of the past quarter-century, two rules of thumb can be derived that do generate inferences about *what* will be learned:

1. People learn what they process.
2. The skills most likely to be learned are the minimal ones necessary to accomplish the range of tasks presented.

The first rule marks the great divide between cognitive and behaviorist conceptions of learning. Its point may be illustrated by a little experiment of many years ago, the source of which is lost to me. Someone got the bright idea of teaching children to recognize the names of colors by printing each color name in its appropriate color ('green' was printed in green, 'yellow' in yellow, and so on) and then gradually fading the colors away until eventually all the words appeared in black. When this was tried in a first-grade class it worked nicely; the children learned to recognize the color words. When it was tried in a kindergarten class, the children correctly named the words so long as there was the faintest trace of identifying color, but when the final transition was made to black, they called every word 'black.' Behaviorally, the children were all doing the same thing except at the last step, accurately naming the color words; cognitively they were evidently doing something quite different. The first-graders were processing the spellings or at least the shapes of the words, whereas the kindergarteners were processing only their color.

The question of what is actually being processed looms large in any serious consideration of so-called 'projects.' A favorite in modern classrooms is having students produce multi-media documents on some topic relevant to the curriculum—on an endangered species, say. With the resources available on the Internet and on CD-ROMs and with the use of scanners and multimedia presentation software that permits the incorporation of video, sound, and graphics, students can produce impressive documents. But what do they learn about polar bears from producing a multimedia document on polar bears? It all depends on what information they process in assembling the document. If the only questions they consider are “Is it about polar bears?” and “Does it look nice?” we may infer that not much polar bear knowledge will be acquired. In a later section I shall argue that knowledge building, as a form of activity, has the advantage over most kinds of project-based activity that it leads to deeper processing of information.

The second rule of thumb encourages us to think small when making claims about the skills to be acquired from learning activities. If the task is sorting buttons, let us not claim that it is teaching ‘classification skills’; let us only claim that it is improving children’s ability to sort buttons. To do well in the paper airplane activity mentioned earlier, students do not really have to master the scientific method (whatever that is) or become all-around critical thinkers (whatever that means). They have to measure and record distances and work out some way to extract conclusions from variable data. They have to learn what some of the significant variables are in paper airplane performance and use these in evaluating designs. Those may not be trivial attainments, but they are not very far-reaching. Some students will get more out of the activity than that, but others will get less. Those who get more will do so because they have pursued objectives of their own, over and above those implicit in the task (cf. Ng & Bereiter, 1991). Those who get less, even though they appear to perform adequately, will have found some way to get along in the actual situation that circumvents the apparent requirements.

This minimalist view of skill learning does not imply that students are shirkers. It simply assumes that students will act like other adaptable organisms and will develop ways of accomplishing tasks that reduce time and effort—including cognitive effort. When I first started doing research on writing, I believed, as many teachers do, that the ideal writing assignment was one that posed a problem that made the writers think. What impressed me in the end was writers’ ingenuity in evading the need to think—except for people who were serious about writing; with them, it didn’t matter how banal the assignment, they would find ways to turn it into a challenge (Scardamalia & Bereiter, 1991). People will go to great lengths to master skills that they want to master. They will seek out challenges and they will also seek out instruction and guidance. But give them a task that calls for a skill they are not eager to develop and their first recourse will be to find a way to accomplish the task without need for that skill. Thus, natural tendencies are naturally at war with attempts to foster high-level skills through indirect learning activities.[7]

The obvious way out of this dilemma is to enlist the students in the effort to learn. This can have striking results. Valerie Anderson developed a program aimed at involving adolescent poor readers in every aspect of the task of becoming a more capable reader, and it has had the effect of turning discouraged and unmotivated students into alert, critical readers (Anderson & Roit, 19xx). Direct learning activities are often criticized for inducing passivity and parrotlike learning, but those are usually secondary consequences of the failure of students to understand or appreciate the value of what they are supposed to be learning.

Not everything can be learned in a direct manner. Sometimes it is because the learners are too young to grasp the purpose. Phonemic awareness—the ability to hear words as combinations of identifiable sounds—has emerged as an important objective for kindergarten and first grade (Adams, 1990); but children of this age could hardly be expected to understand what phonemic awareness lessons are about. Consequently, all the teaching of it is done indirectly, through songs, word play, and exercises. For similar reasons, the most effective approach to teaching number sense to young children is based on games (Griffin, Case, & Siegler, 1984).

Sometimes students can understand an objective but cannot see the value in it until after they have attained it. When I was teaching high school English, I presented my students with a list of possible learning objectives and asked them to rank them according to how interested they were in pursuing them. Up at the top of the list was learning correct grammatical usage. But fourteenth out of fourteen on nearly every student’s list was learning to appreciate good literature. This was a blow to me, because it was near the top of my priorities; but I should not have been surprised. There are adults who feel a need to improve their taste in art, literature, or music, but their need arises out of kinds of social experience unknown to most teenagers. This does not mean it is impossible to develop an appreciation for good literature in high school students, only that it must be done indirectly, through activities that have a different purpose in the eyes of the students.

Some kinds of instructional activities, however, only succeed if the students are trying to learn. I mentioned the Omanson and Resnick research on students using ‘manipulables’ intended to promote understanding of elementary arithmetic algorithms. Some students did grasp the relationship between the concrete operations and the algorithms. Interviews revealed that they were students who realized there was supposed to be a connection and were actively trying to figure out what it was. Evidently, the relation between the concrete activities and numerical operations only comes across to children who are looking for a relationship (Resnick & Omanson, 1987). It is amazing to what extent Dewey’s aversion to the direct approach[8] has permeated educational thought, to the extent that educators often do not even consider the possibility of letting students in on the secret of what they are supposed to be learning. Partly, though, this is a result of conceptual confusion. When students are eagerly engaged in some intellectual activity, such as a science experiment, this is taken to mean that they are eager to learn. But they may not be thinking about learning at all. They may only be thinking about the experiment and its meaning—which is how Dewey would have it. And that may be sufficient, but in some cases it may not. All I am arguing for here is realizing there is a difference.

Learning Through Knowledge Building

Knowledge building, as carried on in the adult world, is not a learning activity. Often it is a form of economic activity, a matter of adding value to knowledge artifacts. Even when it is pursued for its own sake, however, as in pure research and scholarship, it is seldom done for the purpose of learning—that is, for the purpose of mental improvement. It is done to advance knowledge in a more general way. Successful knowledge building may always result in worthwhile learning for the people involved, but that is an incidental outcome, like the gains in physical fitness that may result from manual work.

Knowledge building, as carried on in a school, however, is likely to be viewed and evaluated as a learning activity, whether or not the participants see it that way. According to the distinction developed in a preceding section, knowledge building is an indirect learning activity. It might be interspersed with direct learning activity, with what Hunt and Minstrell (1994, p. 58) call ‘benchmark instruction,’ but that is a recognizable departure from the course of knowledge building work.

Let us revisit the fifth-graders’ inquiry into functions of the brain, only this time let us suppose it is conducted as a knowledge building activity, focused on World 3 rather than on World 2. The students venture tentative theories about how the brain works, and then they seek new information and engage in discussions aimed at improving their theories of brain functioning. Progress is observable. The initial theories tended to be of the homunculus variety: the brain as a little person inside the head who receives information from the senses, thinks about it, and then sends instructions to the muscles. Subsequent work does not entirely eliminate the homunculus (psychologists have always found it hard to get rid of) but a number of automatic functions of the brain are identified and there are the beginnings of a conception of an information processing organ that functions lawfully but without benefit of a wizard behind the curtain.

It is tempting to conclude that these knowledge building accomplishments represent what was learned, but that will not do. We are talking about a collective accomplishment, and we no more know what each individual has learned in the process than we know what has been learned by each individual member of a team that has built a successful museum exhibit or crossed the ocean on a raft. There should be some relation between the collective advance in knowledge and individual learning, but we cannot infer one from the other. The problem of determining what is learned and the even more taxing problem of predicting what is likely to be learned are the same for knowledge building as for any other kind of indirect learning activity. Accordingly, the two rules of thumb presented earlier should apply:

1. People learn what they process.
2. The skills most likely to be learned are the minimal ones necessary to accomplish the range of tasks presented.

To see how these rules of thumb apply, let us consider two different knowledge building activities: planning a trip to Mars and explaining force and motion—what makes things move. The task in each case is to produce a conceptual artifact: a plan, in the first case, and an explanation (or what the students may be encouraged to think of as a theory) in the second case. It will be relevant to issues raised in the next chapter that the first is a sort of pseudo-artifact. No one expects the plan to be carried out or to have any use other than as a pedagogical vehicle; thus the value of the activity lies entirely in the learning that results from it. The explanation, by constrast, is a serious knowledge artifact. It constitutes, for the students who create it, genuine World 3 knowledge. It is something they can use in making sense of the world. For present purposes, however, we shall consider both activities only in terms of the learning that results.

Knowledge building may be considered a variety of problem solving (although not all problem solving is knowledge building). Planning a trip to Mars is a *design problem*; producing a theory of motion is a *problem of explanation*. Applying the first rule of thumb, we may assume in each case that the knowledge that will be processed is knowledge actually used in solving the problem. Similarly, the second rule implies that the skills acquired or refined will be those actually employed in solving the problem. Design problems and problems of explanation are different enough that we should expect rather different kinds of knowledge to be processed and skills to be exercised.

Planning a trip to Mars and similar many-faceted design problems are popular activities among teachers of a constructivist bent partly because, pursuant to the first rule of thumb, they engage students in processing a great deal of knowledge of different kinds. Planning a trip to Mars calls on knowledge about space, the solar system, gravity, properties of matter, nutrition, physiology, and rocketry. Producing an explanation of motion involves a much more limited range of knowledge but presumably knowledge of greater depth. But what does that mean?

Ever since ‘learning by doing’ became a pedagogical watchword, educators have been promoting the belief that practical problems naturally motivate a drive toward understanding. This is true only in a very limited sense. According to Roger Schank (19xx), explanation is ‘failure-driven.’ When our plans or expectations go amiss, we try to find the reason. This leads to understanding, but within the realm of action in which the failure occurred. By solving fuel system problems, auto mechanics develop a good practical understanding of carburetion and fuel injection, but how many of them are ever led to inquire into the chemistry of combustion? And, without an understanding of more basic chemistry, how would they understand an explanation if they sought one out? The kinds of failures that lead to theoretical understanding are failures of explanation. And all explanations fail, in that every explanation posits something which itself needs explaining (Miyake, 1986). But pressing on with explanation upon explanation implies a very different kind of purpose from the practical purposes that motivate ‘learning by doing.’ There is a gap between theory and practice that interferes with exchange in either direction. Theory does not provide solutions to practical problems, but practical activities do not give rise to theoretical problems, either. Yet both are essential to the advance of knowledge.

Students planning a trip to Mars will likely learn about the use of rocket engines to change the course or speed of a space vehicle. But will they be moved to inquire how a rocket engine can work in outer space where there is no air for the jet gases to push against? That would be a digression from their task. That question might well arise, however, in the course of an inquiry into forces and motion. An initial theory about motion might be that things move because they are pushed. But examples would arise of things that provide their own motion, such as animals, automobiles, and rockets. A theory that could account for these cases would be that the motion is caused by pushing against something: The animal’s feet push backward against the earth, the canoe paddle pushes backward against the water, and so on. By this reasoning, jet propulsion would be explained by the explanding gases pushing against the surrounding air—a plausible explanation, but one that fails to explain jet propulsion in outer space. To think of this counterexample, however, students would need to be well aware of jet propulsion working in outer space; and a fact like that might only be thoroughly enough absorbed through knowledge building activities like planning the trip to Mars. Thus both kinds of activities have their place, the one leading to the processing of a lot of useful world knowledge but skirting basic principles, the other going more deeply into problems of understanding but being hampered by students’ limited practical and factual knowledge.

With respect to skills, knowledge-building activities of both kinds encourage high-flown claims: planning skills, reasoning skills, problem solving, creativity, scientific method, and all the higher levels of Bloom’s *Taxonomy* (1956). Such claims, as I have been suggesting, are really advertizing copy. There might be some truth behind them, but that is more a fancy than a warranted belief. The second principle enjoins us to keep our expectations closer to what the students actually do in carrying out the activities. On this basis, knowledge building activities have an outstanding advantage over most other kinds of direct and indirect learning activities that make up formal education. The activities themselves represent worthwhile things to be able to do. Being able to sit down with a group of people and work out a complex design or plan is something real people do in real life. It is not the same as bandying opinions or exchanging anecdotes. It is a demanding sort of discourse, which presents problems in keeping things moving forward while not shutting out objections and divergent ideas and in taking account of relevant facts without getting overwhelmed by complications. It seems reasonable to expect that through taking part in the solution of design problems of various kinds throughout their school careers, students would emerge better prepared to take part in similar activity out in the world than if they had not done so. As for the high-sounding general skills, many are involved in the design process, but it is not necessary to claim the skills developed through design work apply outside it. Being able to participate skillfully in design work itself is enough to justify the activity.

The same argument may be made with respect to the second kind of knowledge building activity, the kind that involves problems of explanation. What is being learned, essentially, is how to participate in theoretical, explanation-seeking inquiry. This may not be a highly marketable skill in its own right, but it is something expected of people in the higher levels of every profession. Furthermore—and this is not a minor consideration in school curricula—it takes on great importance in higher education. High grades and SAT scores may be important for getting into a top-flight university, but being able to participate in intellectual inquiry is essential to doing well there and proceeding into a graduate or professional school. Beyond such practical considerations, there is also the value of explanation-seeking inquiry in any kind of active mental life.

Knowledge Building as Productive Work

If students built the computers they were to use in school, this would have a marked educational advantage compared to doing electronics work of no practical consequence.[9] Part of the advantage would come from what is learned during the building. The students would likely pay closer attention to what they were doing, try harder to detect and solve problems, and thus learn more from it. The other part of the advantage would come from continued learning as they did things with the computers. Because of their knowledge of the computer's innards, they would be able to make more connections between structure and function and would be less bound to rote procedures. Besides that, there would be the enduring satisfaction of having made something they could use.

All these same advantages can be claimed for building knowledge that is subsequently put to use in school activities. I am not speaking metaphorically. As I have been arguing all along, conceptual artifacts should be regarded as real things, albeit immaterial ones. The difference between building a conceptual artifact and building a computer is that with the computer you have the alternative of going out and buying one, whereas with a conceptual artifact there is no alternative to building your own. What that means will warrant closer examination, but let us for the moment take it as a constructivist article of faith.

The comparison to building computers reveals an important difference between the two knowledge building activities considered previously. Planning a trip to Mars, no matter how engaging and instructive it might be, does not yield a product that the students will put to use. Therefore it lacks something in immediate motivation—the plan doesn't really have to work, it only has to meet with approval—and it does not make for continuing learning. By constrast, working on an explanation of force and motion does yield a potentially usable artifact: Call it a theory of motion. If students can see that it is going to be their theory and that they will be using it in subsequent work, there should be motivation to do a good and thoughtful job. And they should learn more in subsequent uses of the theory, having constructed it themselves and being sensitive to its strengths and weaknesses.

Such claims might appear to push the notion of constructivism too far, however. The theory of motion that the students allegedly construct should, if it is to be useful, bear a very close resemblance to Newton's and will in fact owe a great deal to reference book accounts of that theory. So is it really *their* theory? Are their constructive efforts really comparable to those that would go into building a computer? Interestingly, they are. Students would not build a computer from scratch. They would assemble it from available components and according to directions. The result, if successful, would closely resemble something already on the market. Yet they would learn something from building it and would have a sense of ownership. How much learning and how much sense of ownership would at bottom depend on how much problem solving the students had to do in building the computer. This could vary, depending on how the task was structured, but in general the complexity of the assembly and of the testing that has to be carried out ensures that a lot of problem solving would go on, even if there is not the slightest effort to innovate.

The parallels to building a theory of motion are very close. Students do not build the theory from scratch (not even Newton did that), but make use of available resources such as texts, computer simulations, and packaged demonstrations. The learning and sense of ownership that result do not depend on the originality of the theory but on the problem solving and other constructive thought that go into producing the final product. This can vary depending on how the activity is structured. In one approach, where students begin by formulating their own theories and then improve them through research, criticism, and comparison to other theories, a great deal of original work goes into the activity (Hewitt & Scardamalia, in press). But even in the most didactic approach, which starts with a lecture or textbook explanation, students have to do substantial constructive work and problem solving before they can be said to *have* a Newtonian theory of motion. Sir Karl Popper put the point very well:

> What I suggest is that we can grasp a theory only by trying to reinvent it or to reconstruct it, and by trying out, with the help of our imagination, all the consequences of the theory which seem to us to be interesting and important. . . . One could say that the process of understanding and the process of the actual production or discovery of... [theories, etc.] are very much alike. Both are making and matching processes. (Popper & Eccles, 1977), p. 461.

Of course, didactic teaching often fails to mobilize such constructive effort, but insofar as that is true it fails to do its job and is teaching mere verbalisms.

Speaking more broadly about understanding existing World 3 objects, Popper said:

> According to my view, we may understand the grasping of a World 3 object as an active process. We have to explain it as the making, the re-creation, of that object. In order to understand a difficult Latin sentence, we have to construe it: to see how it is made, and to re-construct it, to re-make it. In order to understand a problem, we have to try at least some of the more obvious solutions, and to discover that they fail; thus we rediscover that there is a difficulty--a problem. In order to understand a theory, we have first to understand the problem which the theory was designed to solve, and to see whether the theory does better than do any of the more obvious solutions. In order to understand a somewhat difficult argument like Euclid's proof of the theorem of Pythagoras (there are simpler proofs of this theorem), we have to do the work ourselves, taking full note of what is assumed without proof. In all these cases the understanding becomes "intuitive" when we have acquired the feeling that we can do the work of reconstruction at will, at any time." p. 44
> [Popper, 1977 #1581]

Should the student be asking, "Have I learned this content? Do I understand this concept?" The first question is generally one that the student should ask, as an intentional learner. The second can generally be restated as a World 3 question. Suppose the subject matter is the formula for solving quadratic equations. Instead of asking, "Do I understand this formula?" the student may better ask, "Does this formula make sense?" That then becomes a question to be discussed, explanations or proofs constructed, and so on. In short, restating it as a World 3 issue has more indications for positive action. On the other hand, "Do I remember this formula? Can I apply it?"--these are questions about one's own competence. They are questions to guide learning rather than to guide World 3 discourse.

We need students who are *both* intentional learners and knowledge builders. We need teachers who are good at *both* promoting learning and promoting knowledge building. The hardest idea I have ever tried to get across to teachers is that there are two jobs here, not one. No amount of enriching learning activities and turning responsibility for them over to the students will ensure knowledge building. No amount of knowledge building will produce all the learning that students need to acquire in school.

The Place of Teaching

> "Teach us science! We want to learn science!"
> "What kind?"
> "Any kind. We'll stay in at lunch. We'll be your assistants. We want to learn it all!"
>
> Fourth-grade children in an inner-city school, after what may have been the first time they ever read and understood a scientific article.

One of the more tedious exhortations to which teachers are subjected (often by other teachers) is to pay more attention to learning and less to teaching. This exhortation is often lightened by the witticism, "I taught it but they didn't learn it." On any reasonable construal of the word 'teach,' this is nonsense. If they haven't learned it you haven't taught it. But the fact that the witticism always draws an appreciative chuckle, overworked as it is, suggests that it resonates with the experience of teachers.

The three kinds of reduction discussed earlier all tend to separate teaching from learning. Reduction to subject matter converts complex objectives into subject matter to be covered. There is a focus on learning, but it is learning dissociated from the original objectives. When, for instance, the teaching of problem solving strategies is reduced to subject matter, the learning consists of little more than memorizing a short list of slogans. The trouble in such cases is not that students fail to learn what they are taught but that what they are taught is not what they are intended to learn. Cognitive strategies are ways of managing your mind so as to achieve certain results, and nothing like that is being taught. Really to teach cognitive strategies rather than mere slogans or rituals requires a much more complex and extended kind of teaching (cf. Bereiter & Bird, 1985; Pressley, Harris, & Marks, 1992).

Reduction to activities—by far the most common type of reduction in elementary schooling—pushes learning out of the picture. What the teacher does is still, for historical reasons, called teaching, but it might more accurately be called management (Berliner, 1983). It is concerned with keeping the students engaged and seeing that the activities move ahead. The third type, reduction to self-expression, puts the students' interests at the center, leaving the teacher in the role of facilitator and responder. This is often what is meant by focusing on the learning rather than the teaching, although both teaching and learning are rendered ambiguous by this kind of reduction.

Dictionaries offer a range of meanings of the word 'teach.' One meaning is to give lessons in or hold courses in a subject. By that definition, all these reductive practices count as teaching, and there is merit in calling teachers' attention back to learning. But in the more generic sense of the word, to teach is to cause learning. The teacher can no more ignore learning than the comedian can ignore laughter. It is what defines the activity. Teaching, in this more demanding sense, is inherently problematic. It requires what Newell and Simon (1972) called "means-end analysis," and the end—whether specified in advance or allowed to emerge opportunistically—is learning. There are always things that need monitoring, that may deflect progress toward the hoped-for learning and that require strategic teaching moves.

The effect of reductive practices is to remove this problem-solving element, reducing teaching to something that can just be carried out or that presents problems of a more manageable sort. Eliminating problems in favor of routines is normal adaptive behavior. The crucial question always is whether you eliminate peripheral problems in order to devote more mental resources to the central ones or whether you routinize the central problems and thereby become deadwood (Bereiter & Scardamalia, 1993, Ch. 4). In all the teachers' meetings I have attended, the most resounding applause is for those whose message is that you don't really have to teach. The message is never put in those bald terms, of course. Most often the message takes the form of an attack on the empty bucket theory of learning, a theory that invites attack because of its complete absence of defenders. Teachers are told that adherence to this discredited theory results in subverting young people's natural readiness to learn, drilling them on meaningless subskills and boring them with lectures and busywork. But the underlying message, translated into normal English, is that teaching is the enemy of learning. It is in this context that it becomes pernicious to exhort teachers to pay attention to learning rather than teaching. It is shameless exploitation of teachers' uncertainties and disappointments, and yet the audience for it seems inexhaustible.

This probably reads like taking time out for a rant. The reason I think it is important in a chapter on the topic of learning is that this spurious antagonism between teaching and learning has come to be associated with ideas like constructivism and knowledge building that I have been advocating here. Therefore it is important to make clear that this isn't so. Sometimes the most important constructivist move you can make, the most vital way of promoting knowledge building, is to sit students down and teach them something.

The epigraph to this section comes from my own experience. We had a project going in a school where literacy levels were low, much lower than the brightness and eagerness of the children would have led one to expect. The particular kind of reductionism that flourished in this school was reduction to self-expression. This was used to good effect in creating what could properly be called a therapeutic environment for children who lived where murder was a common event. But where reading was concerned it meant chat sessions that required little attention to the book in hand. We were trying to do some science and the teacher brilliantly hit on the topic of sleep, for lack of sleep and troubled sleep were common in the experience of these fourth-graders. So they talked about sleep and raised interesting quesions about sleep, but how were they going to learn anything?

The don't-teach philosophy might say that kids talking about sleep, sharing their experiences and feelings, is enough, that that *is* learning. But these kids had developed some real curiosity about the topic. They had begun to ask the question that has motivated much of the research on sleep and dreaming: What is it, anyway? Uninformed discussion cannot get you very far in seeking to answer that question. But what do you do with kids who would be helpless trying to read a children's encyclopedia article? Give a little lecture perhaps or scout up an appropriate video? These might have served the immediate purpose, but we were hoping to make some advance on the more fundamental problem of the children's inability to learn from reading.

Reciprocal teaching (Palincsar & Brown, 1984) is an approach that has been successful in enabling poor readers to learn from reading. The students are taught to ask each other questions answerable from the text, raise points in need of clarification, predict what kind of information is coming next, and sum up what they have learned so far. I have had the hunch that a considerable part of the effectiveness of this approach comes from getting kids to realize that a text can tell them things about the real world that they do not already know.[10] This seems so obvious that it does not occur to educators that it might be anything a person needs to learn. Reciprocal teaching is designed to be carried out in small groups and must be worked up to through coaching. It is not something a teacher can pull out of a hat. I wanted to see whether some of the same effect could be achieved by leading the whole class through reading an informative article, using a rather free translation of the reciprocal teaching procedures. I handed out copies of the simplest article we could find that explained the use of electroencephalographic recordings (EEGs) to study the various phases of sleep. It was, however, a good four years beyond the prevalent reading grade level in the class.

My role did not consist of much more than managing the process. I would call on someone to read a few lines, then he or she would call on someone to ask a question, and so on, with plenty of freedom for discussion. The kids read ‘EEG’ as ‘egg,’ but I let that pass. Their interest mounted as they began to get on to the idea of electrical events in the brain revealing what the brain was doing in during sleep, and before long the formalities of reciprocal teaching gave way to excited discussion of what this all meant. When we got to the end of the article and several children had taken their shots at summing up what they now knew, I raised the question, “How do you know all this?” The question produced a stunned silence. “Do the kids in the next room know this?” I continued. No, they agreed, they probably did not. “Well, then how is it that you know all this and they don’t? Did I tell it to you?” Some of them wanted to say yes, but they quickly worked it out that I had not in fact told them anything. Then you could see the light starting to dawn in face after face as they realized that they had learned it themselves by reading. And it was after that, as I was preparing to leave, that they crowded around me and the “Teach us science!” dialogue ensued.

This was one of the highest points in my life as a teacher; but did I really do any teaching? I had not conveyed information. No knowledge had poured from my head into theirs. I had merely managed an activity, as a result of which the children learned something; but was this any different from the ‘reduction to activities’ that I have been maligning? The crucial difference, in my view, is that I was engaged in a sustained goal-directed effort to achieve a learning objective—in fact, two learning objectives. I wanted the kids to learn something of what scientists know about sleep and I wanted them to learn that they had the capacity to expand their knowledge through reading. It was 40 minutes or so of very hard work during which I was continually problem solving, choosing various strategic moves within the bounds I had set for myself, monitoring progress, and adjusting strategies to move closer to the two goals. Of course, neither goal was fully achieved, nor did I expect that; all I was hoping for was what Marion Blank called a “glimmer.” That seemed enough for 40 minutes work. So, yes, I would say, that is teaching. Preparing a lecture and rattling it off or showing a videotape would not be teaching. But embedding either of those in a discussion and working at it until there was evidence that learning had been achieved would be teaching. In short, teaching is not defined by the kinds of actions the teacher engages in but by the ends towards which those actions are adaptively employed.

In simple terms, teaching means taking responsibility for someone else’s learning and carrying through the actual problem solving required to bring that learning about. Parents accept responsibility for their children’s learning, but in sending their children to school they delegate the problem solving to someone else. What is becoming increasingly the practice in schools is that the teachers further delegate the problem solving to specialists. The classroom becomes a center of activities that result in most of the children learning enough to satisfy the parents and other monitors. The children who come up short are then passed on to remedial instructors and therapists, who do the actual problem solving—if it is to be done at all. The unfortunate fact is that many remedial instructors are just as committed to reductive practices as the classroom teachers. Failures may then be passed on to still another specialist, and so on to the extent of the school system’s or the parents’ resources. This can result in the anomaly I have encountered more than once, where a youngster, by now badly damaged from repeated experiences of failure, ends up with a psychiatrist who becomes the first person in the chain to actually sit down and try to teach the youngster to read.

There is, moreover, something important to be said for teaching even in those cases where reductive practices yield satisfactory learning of the curriculum content. When the educational process is reduced to activities and self-expression or to the teaching of incidental subject matter, the personal rewards of teaching become attenuated. The teacher can take pleasure in the happy buzz of well-occupied students and can share students' pride in jobs well done, but this is pretty thin nourishment to the teacher's ego, especially as the years roll on. When I say that the 'EEG' episode was a high point in my life as a teacher, it is clearly because I feel that I made something important happen. Such feelings, whether well-founded or not, make teaching a rewarding experience. I suspect that the absence of such experiences has much to do with the reported decline in morale of school teachers. Such feelings only come from problem solving, from trying to make important things happen and occasionally succeeding. Reductive practices, by eliminating the problem solving, eliminate the experience of having successfully taught.

More important, however, is what teaching episodes mean to the students. As everyone recognizes, successful teaching requires involvement on the part of the learners. This is not easily gained. It is often much easier to get students involved in playing a game, building something, having an argument, doing an experiment, or even just doing pages in a workbook. Students can learn from all of these, and the learning may be of great significance. But the students have no experience of learning. They may experience the excitement of the activity, the joy of winning a game or argument, the satisfaction of having produced something good or of having performed well, but the learning itself is unconscious. The students who crowded around me and begged to be taught science had, I think it is safe to say, experienced something important having happened to themselves and they wanted more of it. I have seen the same thing happen with children as young as five or six. The experience of learning is deeply satisfying, and when it happens for the first time, after years of unconscious learning, the effect can be explosive.

There is much talk these days about how people must become lifelong learners. I am not sure how you produce lifelong learners or even whether that is quite the right objective, but I do think it is clear how not to produce lifelong learners, and that seems to be the way education has been heading. The way not to produce lifelong learners is to divert students from the experience of learning throughout most of their school lives, confining the experience to unpleasant circumstances such as preparing for tests or overcoming a deficit.

Another way not to produce lifelong learners, however, is to make students' experiences of learning always dependent on a teacher. In the example I have been using, it is clear from the students' urgings at the end that they saw the learning they had experienced as dependent on me as teacher. One of my goals had been for them to realize that they could learn by themselves from reading, but they were no doubt correct in believing that what had happened that day could not be repeated without a teacher's help. Being able to proceed independently would lie at the end of a process during which more and more of the direction was turned over to the students. So there is a bit of a dilemma here. If the teacher stays in the background and just oversees indirect learning activities, the students get no experience of learning. If the teacher takes over and teaches, the students experience learning as something that requires a teacher.

I say 'a bit of a dilemma,' because it is easily resolved by adding a time dimension. In earlier work, Scardamalia and I framed the issue in terms of three kinds of teachers (Bereiter & Scardamalia, 1987; Scardamalia & Bereiter, in press[xxCIAR]). Teacher A reduces teaching to activities, which could be anything from old-style workbook activities to the trendiest of space-age projects. Teacher B assumes responsibility for students' learning, doing all the problem-solving that that entails. Teacher C does that as well, but with the added objective of continually turning more of the learning process over to the students. Walking into a classroom, you cannot immediately tell these three kinds of teachers apart. One of the things you might see going on these days is the students working in groups to produce videos or multimedia presentations. The teacher is likely to be found going from group to group, checking how things are going and responding to requests. Over the course of a few days, however, differences between a Teacher A and a Teacher B would become evident. Teacher A's focus is entirely on the production process and its products—whether the students are engaged, whether everyone is getting fair treatment, and whether they are turning out good pieces of work. Teacher B, of course, attends to all of this as well, but Teacher B would also be attending to what the students were learning from the experience, and would be taking steps to ensure that the students are processing content and not just dealing with show. To see a difference between Teachers B and C, however, you might need to go back into the history of the media production project. What brought it about in the first place? Was it conceived from the start as a learning activity, or did it emerge from the students' own knowledge building efforts? In one striking example of a Teacher C classroom, the students had been studying cockroachs and had learned so much from their reading and observation that they wanted to share it with the rest of the school, and production of a video came about to achieve that purpose (Lamon, Caswell, Scardamalia, & Chandra, 1997).

The differences in what might seem to be the same learning activity are thus quite profound. In the Teacher A classroom the students are learning something of media production, but the media production is likely getting in the way of learning anything else. This is much like the evaporation-condensation project reported by Anderson, Holland, and Palincsar (1997) and discussed earlier—a pure example of Teacher A pedagogy. In the Teacher B classroom, the teacher will be working to ensure that the original educational purposes of the activity are met and that it does not deteriorate into a mere media production exercise. In the Teacher C classroom, the media production is continuous with and a direct outgrowth of the learning that is coming to be embodied in the media production. The greater part of Teacher C's work has already been done before the idea of a media production even comes up, and it remains only to help the students keep sight of their purposes as they carry out the project.

In a typical classroom activity, whatever model of teaching is supposed to be in force, the teacher will manage the intellectually higher-level parts of the process, leaving lower-level parts to the students. The teacher asks questions, the students answer; the teacher poses problems or issues, the students pursue solutions or offer opinions; the teacher sets out an objective, the students work out ways to achieve it. In taking fourth-graders through reading the article about sleep, I might have seemed to be turning most of the process over to the students, but in fact it was I who decided as they went along whether they were understanding it or whether other questions needed to be raised or other answers or summaries solicited. And if they had become bogged down and were not making progress in understanding the article I would have done more. That is normal Teacher B behavior. Only over a course of weeks might it become evident whether I was helping the students to take over more of these higher-level parts of the process, as the Teacher C model stipulates, or whether I was holding on to the reins, perhaps because only in doing so could I continue to see myself as an accomplished teacher.

Teachers A, B, and C are abstract models that of course fit real teachers only to a degree—and more on some days than others. The three models correlate with the three levels of goals discussed earlier, but not perfectly. Teacher A's focus is task completion, and any knowledge building that takes place will likely encounter unwitting interference from the teacher—as in the Yarnal and Kafai (1995) study, where the teacher discouraged subject-matter questions because they deviated from the rules of the game-design project. Teacher B's focus is learning goals—that is, goals held by the teacher for what the students are to learn—but this does not preclude knowledge building as a way of achieving those goals. A climate that genuinely fosters knowledge building is likely to be found only in a Teacher C classroom, but that does not mean Teacher C will necessarily favor this approach. It is quite within the bounds of the Teacher C model to focus on learning but not knowledge building, getting students to take the initiative in pursuit of established learning goals and to help other students in the same pursuit. Probably most real-life approximations to Teacher C would see their job that way.

Conclusion

I see an important line of pedagogical evolution leading toward knowledge building. Its early origins were in Dewey, Froebel, and Montessori. It took definite shape in the English infant school movement, with its emphasis on 'the child's own question' (Isaacs, 1930; Weber, 1971), and it found a theoretical basis in Piaget's 'constructivism' (Piaget, 1929). In the 1950s and 60s, it began to be modeled on scientific research, and 'learning by discovery' and 'guided discovery' were born (Shulman & Keisler, 1966). In the 1980s and 90s, aided by a growing appreciation of the social character of scientific research, the idea of communities of inquiry replaced what had up till then been a strangely individualistic conception of knowledge construction (Brown & Campione, 1990). It received a technological boost and a conceptual setback as computerists took up the cause of ‘project-based learning.’

Throughout this century-long evolution, however, the ideas of knowledge construction and of learning remained not intertwined but hopelessly muddled together. This confusion has not diminished but has perhaps reached its zenith in project-based learning. In its more exalted forms project-based learning amounts to knowledge-building:

> Project-based science focuses on student-designed inquiry that is organized by investigations to answer driving questions, includes collaboration among learners and others, the use of new technology, and the creation of authentic artifacts that represent student understanding. (Marx, Blumenfeld, Krajcik, & Solloway, 1992, p. 341).

The last phrase is a killer, however. We have seen how producing “authentic artifacts” can take over as the purpose of the enterprise, leaving the driving questions and the pursuit of understanding out in the cold. This is reduction to activities. Question-driven inquiry, an intellectual pursuit of a high order, becomes reduced to the production of posters, skits, and movies.

So common is this reductionism that I predict “project-based learning” will soon become a term of ridicule, the way “learning by doing” and “basket weaving” did in reaction against the excesses of progressive education. Indeed, what brought about that reaction against progressive education was the same kind of reduction to activities—activities that become silly as soon as they are detached from their original purposes.

To ward off such reductionism, it is essential to distinguish between learning and knowledge building. Such a distinction enables us to see that the above description of project-based learning muddles together two purposes that need separate attention. One is the knowledge-building purpose. This is inquiry driven by real, not contrived questions. The intended product of such inquiry is an artifact, all right, but it is a *conceptual* artifact, not a material one like a poster or a movie. Scientists do not set about their inquiries with the objective of producing posters or skits or even books and scholarly articles—or to the extent they do, they should not be serving as examples to the young. Their objective is to produce a piece of knowledge, which is likely first to find embodiment in messages exchanged with peers, before it ever finds its way into a public display. And that is a natural way for things to proceed with students' knowledge-building efforts. The resulting conceptual artifacts will be “authentic” to the extent that they are things the students can actually use—primarily for purposes of understanding real-world phenomena and texts that refer to them. That is why, in the classroom examples I discussed earlier, an explanation of jet propulsion is a more authentic artifact than a plan for a trip to Mars.

The other purpose implicit in the above description of project-based science is learning. As a result of their question-driven inquiry the students are supposed to acquire some scientific understanding. That seems to be the intent behind the reference to “artifacts that reflect student understanding.” But the appropriate concern here is not with what the class collectively has found out. That concern belongs to knowledge building. The concern should be with what individual students get out of the effort. A collaboratively produced demonstration is not a very accurate means of assessment for this purpose. It is likely to reflect what the most successful learner understood, or possibly something beyond what any individual has grasped.

The likely result of this muddling of purposes is that knowledge building gets diminished if not obliterated by the intrusion of activities devised to promote learning, while learning comes to be neglected as a result of its getting confused with collective accomplishments. For the teachers, the only things they can get a firm hold on are the inquiry procedures—the observations, records, discussions, and the like—plus the “authentic artifact,” the tangible, visible thing that is produced in the end. Hence reduction to activities becomes complete.

The best hope I can see for straightening this out is to scrap project-based learning, with its built-in connotations of reduction to activities, and to substitute the idea of knowledge building. But this is not merely a matter of labels. To grasp the idea of knowledge building, educators have to understand the following:

- Knowledge building is not just a process; it is aimed at creating a product.
- That product is some kind of conceptual artifact—for instance, an explanation or a historical account or an interpretation of a literary work.

- A conceptual artifact is not something in the minds of the students.
- It is not something material or visible, either.
- It is nevertheless real and preferably something students can use.

Once these ideas are assimilated, it becomes obvious that what students have learned is a different issue from whether they have created a worthwhile piece of knowledge. The conceptual artifacts students have produced through knowledge building provide evidence that some learning has occurred, but who has learned what and how well are questions still to be answered—as they are with any other activity carried out for educational purposes.

It is always hard to determinine what students are learning and even harder to know what to do in case their learning is inadequate, especially because these judgments must usually be based on grossly inadequate information. But they are what makes teaching teaching. They are what it means to take responsibility for someone else's learning. To avoid them is to become a mere classroom manager. Pretty much everyone recognizes this, but there is a strong temptation to believe that some new technology or some new leap forward in educational thought has rendered those onerous judgments no longer necessary. Constructivism, learning by discovery, learner-centered education, computer-assisted instruction, whole language, hands-on learning, computer conferencing, knowledge building: each one comes accompanied by a little demon who whispers in the ear, "You won't have to worry about learning anymore. It will take care of itself. Stop teaching. Just let it happen."

Sometimes, as occurred in the early 1990s in California, the demon's tempting words become official policy, and then you have a disaster. More often, teachers only half believe the message, but going halfway can mean both weaker teaching and missing out on what is most valuable in the new approaches. Instead of such halfway accommodation of inquiry approaches, there ought to be whole-hearted appreciation of both the values and the limitations of three different aspects of knowledge building:

1. Knowledge building as productive work: This is the same kind of work in the classroom as it is in the research laboratory. It is working collectively to produce conceptual artifacts that are of some

use. For students, it is mainly a matter of producing conceptual artifacts that help them understand the world.

2. Learning through knowledge building. This is the learning of scientific, historical, literary, or other kinds of content through knowledge building—through solving problems of understanding in these domains. This is indirect learning, learning that occurs as a byproduct of activity carried out for another purpose, and it cannot be taken for granted.

3. Learning to be a knowledge builder. This is the unique added advantage of knowledge building as an educational approach. It has great potential value for living and working in a knowledge society. But like any other kind of learning, it cannot be taken for granted just because students appear to be engaged in the relevant activity. Teachers have to ascertain whether the learning is actually happening and marshall their best pedagogical resources when it is. not.

---

[1]. Our first public use of the term was in a paper titled "Schools as Knowledge-Building Communities," presented by Scardamalia in 1989 at a workshop on Development and Learning Environments, University of Tel Aviv, Tel Aviv, Israel.

[2] . In Ng & Bereiter (1991), three kinds of goal orientations are identified: task goals, instructional goals, and knowledge-building goals. These correspond approximately to the three kinds of students described below, except that knowledge-building was not so clearly distinguished from learning in the earlier work.

[3] There is also a very large research literature (see xx for a compilation) on students' orientations to learning, all of which makes some distinction similar to the one we make between task completion and learning goals, but none of it distinguishes what we call knowledge building goals from learning goals.

[4] In all the quotations from the student database, punctuation and spelling have been corrected and real names have been replaced with pseudonyms. Otherwise, the quotations are verbatim.

[5] I write this from the experience of having worked for many years, as an author of Open Court reading programs, to get away from reductive practices and to incorporate what research showed were effective ways to improve reading comprehension abilities. Twenty-five years of well-publicized research on reading comprehension have yet to quiet complaints that any departure from the quiz mode means "there's no comprehension."

[6]. Computer Supported Intentional Learning Environment (Scardamalia & Bereiter, 1994; 1996).

[7]. The principle of "they only learn what the tasks require" applies even when motivation is high. Students who join a chess club may be highly motivated to learn to do well in chess and they may study hard and think hard in pursuit of this objective. But the result will be that they learn skills applicable to chess; they are unlikely to learn much in the way of skills, concepts, or habits that are of use anywhere else. To claim that they will learn reasoning, planning, or pattern recognition skills is to indulge in the word magic I have been criticizing. They will learn chess reasoning, chess planning, and chess patterns. Beyond that, we should not be venturing claims.

---

[8] Dewey called making learning a “direct and conscious end in itself” one of the “evils in education” (Dewey, 1919, pp. 168-169).

[9]. The example is somewhat passé. As computers have gotten cheaper and to involve less use of the soldering gun, making one’s own computer has become less instructive and less economical. I have not been able to turn up any recent examples of its being done in schools except at the postsecondary level. Students’ building their own computers has given way to building their own web sites, which is not the same kind of activity at all. So let’s imagine the year is 1980 and that building a machine with 8 kilobytes of RAM would result in a useful tool.

[10]. Stories, even realistic ones, don’t count for this purpose, although they have other estimable values. Indeed, I suspect that an exclusive concentration on fiction in the beginning years of reading instruction contributes to children’s seeing reading as having no connection with their efforts to learn their way around in the world.