Bereiter, C. (in press). Artifacts, canons, and the progress of pedagogy: A response. In B. Smith (Eds.), Liberal education in a knowledge society. Chicago: Open Court.

Artifacts, Canons, and the Progress of Pedagogy: A Response to Contributors

How, other than through the occasional updating of content, ought liberal education to change to meet changing conditions? That was the question addressed in my target paper, "Liberal Education in a Knowledge Society." The answer I offered is a radical one, radical in that it entails a change in goals and conception rather than only a change in method or content. Yet I see it as preserving rather than abandoning the essence of liberal education. Quite appropriately, most of the responses dealt with conceptual issues and underlying assumptions. Disappointingly, there was little direct questioning of my conclusion: that schools should change to become workshops for the production of knowledge. I fear that the reason this proposal escaped criticism is that I have not yet succeeded in making it intelligible.

Perhaps I can make the proposal clearer by moving away from generalities to focus on one topic, the teaching of evolution. This may be an unwise choice, because the topic has given rise to much impassioned debate that is irrelvant to the points I want to make. Most of the contemporary debate, however, including the notorious dispute over the demotion of evolution in Kansas, has concerned whether evolution is "just a theory" or whether it is established fact. In the Popperian terms that I adopted in my target paper, this is an argument about whether the teaching of evolution is teaching about a World 3 object, Darwinian theory, or whether it is teaching about World 1, what really happened over the course of life on earth. Obviously, it should be both, but the sad fact is that school teaching is likely to focus exclusively on Worlds 1 and 2, ignoring World 3. The textbook or some related kind of document will be the basis of study. The textbook treatment of evolution will typically be treated as a surrogate for World 1—as a stand-in for the flesh-and-blood reality of evolutionary history.

Alternatively, it may be treated as a statement of beliefs, a window on the scientists' World 2. Or it may, with little regard for consistency, be treated as both. Uncertainties may be acknowledged, but they are uncertainties about what is *really* the case, not uncertainties about the status or qualities of a World 3 object. Although the term "theory" may appear in the discussions, it is likely either to refer to personal opinions or else simply be used to indicate uncertainty. But Darwin's theory and competing or revisionist theories will not be dealt with *as such*, as objects of inquiry in their own right.

This focus on Worlds 1 and 2 and ignoring of World 3 leaves several serious gaps. There is a cultural gap: Students acquire no sense of what a major intellectual leap Darwin's theory was over its predecessors and what a profound impact it had on 19th century thought. There is an epistemological gap. Thomas Kuhn (1970, p. 77 ff) has emphasized that scientists never evaluate theories solely by comparing them to the world but always also by comparing them to each other. Without this kind of World 3 analysis, students are unlikely to escape from naive empiricism in their scientific thinking. There is also a more specific comprehension gap: Large numbers of students fail to grasp the idea of natural selection. Failure to grasp the idea leads to an intellectual gap: The idea of natural selection has applications far beyond explaining the origin of species (Ohlsson, 1993). It applies to such diverse mechanisms as the immune system and operant conditioning. It explains why misuse of antibiotics produces resistant strains of disease germs and why cancers develop resistance over the course of chemotherapy. It is essential to understanding the difference between using knowledge of the genome to improve classical breeding practices and using it to produce genetic alterations directly (a difference with significant implications for what appears on our grocery shelves). The idea has been used and misused in social theories, is basic to sociobiology and psychobiology, and is the basis for one of the more plausible explanations of creativity and genius (Simonton, 1999). In short, the idea of natural selection—the World 3 object—should be part of the intellectual equipment of every educated modern, including those who for whatever reason reject some or all of the empirical claims of evolutionists.

According to Stellan Ohlsson (1991), who has studied evolution beliefs among university students, it is not that students acquire a wrong theory—a Lamarkian theory, as is commonly said. Ohlsson concluded that most students have no theory of evolution at all, that is, no explanation of how evolution works. They simply take species adaptation as a fact. This suggests that they have not visited World 3 and got things wrong; they have not visited World 3 at all. Their study has been confined to World 1 (albeit World 1 as represented in textbooks).

If you look at scientifically respectable biology textbooks, it is not apparent why the teaching of evolution should fail on so many counts. The necessary information and explanations are there. With normal instruction by teachers who understand natural selection themselves, half or more of the students will indeed grasp the idea, but a sizable minority will not. But of those who do grasp it, how many will be able to use it, and use it with proper care, as an intellectual tool for purposes other than explaining biological adaptation? I don't know, but I suspect the number is small, probably as small as the number who master algebra as an intellectual tool for purposes other than solving textbook algebra problems.

Something more seems to be required than clear explanations and careful instruction. My paper addresses what that something is. It is immersion in World 3, not just as a visitor but as an active participant—a knowledge worker, a producer of knowledge. Popper said that in order to grasp a theory you must first understand the problem that the theory was intended to solve and try out the obvious solutions to discover that they don't work (Popper & Eccles, 1977, p. 44). Doing that means getting into the theory-building business yourself. According to this proposal, the biology class needs to become a workshop for developing, testing, and improving explanations of biological phenomena. Trying to develop an evolutionary theory of their own and testing it against facts—does it explain this? does it explain that? does it imply that apes should be extinct, which they obviously are not?—leads students to recognize the problems that a theory of evolution must solve. As they begin to see how Darwinian theory solves these problems they will not only grasp its core idea

but they will appreciate keenly what a powerful idea it is, how it changes your whole outlook on nature. And, with encouragement, they may begin to apply the idea to other problems of understanding. Of course, they are likely to apply it simplistically at first (as many respectable 19th century thinkers did), but that is liable to happen with any new tool. What is important is to sustain the effort at idea improvement.

The proposal, obviously, is a constructivist one. The phrase, "production of knowledge," gives that much away. There are a lot education proposals in the air these days that carry the label 'constructivism,' and there is also a mounting counter-reaction. In its most degraded forms, constructivism is distinguished by the absence of instruction. A slightly more coherent form goes by the name of "project-based learning." It is the almost universally favored way of putting information technology to work in schools (Moursund, 1999). Although my proposal is compatible with "project-based learning," it insists on a distinction that is not made by advocates of that approach—a distinction between projects aimed at creating a piece of knowledge and those aimed at producing a tangible object. The trouble with constructivism as generally practiced in schools is that, compared to traditional subject-matter instruction, it takes students even farther away from World 3, into an almost exclusive occupation with the palpable and observable.

My proposal is not, as Miller seems to suspect, a plot to pervert liberal education into "the production of knowledge that is useful in the global marketplace." Usefulness enters the picture only in the way that Whitehead intended when he said (1929, p. 14),"Of course, education should be useful, whatever your aim in life. It was useful to Saint Augustine and it was useful to Napoleon. It is useful, because understanding is useful." The most likely products of youthful knowledge building are explanations of physical and cultural phenomena, interpretations, historical accounts, and other more-or-less scholarly efforts. Thus the problem with my proposal is not to see what it preserves of liberal education but what it changes. Seeing what is different about it requires a distinction between learning and knowledge building; this distinction in turn rests on a distinction between knowledge in the head and knowledge that is in

some reasonable sense independent of individual minds. These distinctions were the object of much of the criticism in response to my target paper.

The Constructedness of Knowledge

Everyone seems willing to recognize a kind of knowledge that in some sense exists independently of its authors or of individual knowers. Conceptions differ widely, however. There is the commonsense view, probably still dominant in schooling, of a body of truths of which personal knowledge constitutes a small and probably corrupt sample. Then there is the opposite view, enunciated by Nonaka and Takeuchi (1995) and seemingly taken up as the final word by business pundits. It sees 'explicit knowledge' as simply personal knowledge made public. Then there is a sociological view according to which certain beliefs and categorizations become certified as knowledge through the actions of institutions such as scientific societies and law courts (Bloor, 1998). Philosophers over the centuries have produced a number of conceptions, from Plato's pure ideas to Reck's 'conceptual possibilities.' Semioticians and socio-cultural theorists have on the one hand rejected the idea of knowledge as an "immaterial object that... exists, independent of the linguistic formulation and argumentation through which it was constructed" while at the same time recognizing something like what Wells calls "semiotic artifacts," which "play a central role in the knowing of those involved." (Wells, this volume, p. xx)

It is tempting to leave the sorting out of differences to philosophers and to adopt some suitably vague notion like 'explicit knowledge' in order to get on with educational argument. But that will not work if the argument you wish to make is about changing the way knowledge is treated in schools. It then makes a difference whether you are talking about knowledge as a body of truths, as externalized personal beliefs, as institutionally certified beliefs, as hypothetical possibilities, or as semiotic tools. Rigor, though always desirable, is not of the first importance. Serviceability is. The most serviceable conception of knowledge-outside-the-head that I have been able to find is Popper's idea of World 3. Despite the criticisms, of which those by Reck and Wells are among the more charitable, and despite the general disdain that falls from some quarters on

anyone who says a good word for Popper, I have not found an alternative that serves the purposes of educational modernization as well.

To be serviceable in the contemporary context, a conception of knowledge has to represent it as something that can be created and worked with, criticized and improved. Also—and this is essential for the proposal I presented—it needs to allow the work of scientific research laboratories and leading-edge scholars to share the same conceptual space as the efforts of school children to formulate problems and explanations, to conjecture and refute, and to interpret the products of those scientists and scholars. Popper's World 3 does all this. The immaterial objects composing World 3 are unequivocally human constructions, yet they have a certain autonomy that allows them to be studied, modified, adapted to new uses, and so on. Popper was also clear that creating a World 3 object and understanding one that already exists are essentially similar processes (Popper & Eccles, 1977, p. 461). The great virtue of Popper's three worlds scheme from an educational point of view is that it provides the basis for a reasonable and practical constructivism.

There is another meaning of constructivism only tangentially related to the one I have been advancing. This is a meaning that has come to prominence in the sociology of knowledge. Much more controversial and inflammatory, it tends to drag the educational version of constructivism along with it. So it is important to separate them. The inflammatory version, commonly referred to as "social constructivism" or "constructionism," asserts that *truth* is socially constructed (Berger & Luckman, 1967). According to this view, the fact that most scientific theories are the work of European males is relevant to judging them. This is in contrast to the more traditional view (held by Popper, among many others) that the truth of a proposition is independent of its origins and of what people happen to think.

The confusion about constructivism is aggravated by the ambiguity of the term knowledge. If you take the traditional epistemological definition of knowledge as "true or warranted belief," then to say that knowledge is constructed is the same as saying that truth is constructed. But if you hold the more workaday conception of knowledge as something like "intellectual

property," then saying that it is constructed is almost self-evident and it does not commit you to any position one way or another concerning the nature of truth. This latter is the position I have been trying to maintain.

When Reck (Chapter xx) criticizes constructivism he is apparently doing so on the basis of a "true or warranted belief" conception of knowledge. The *truth* of the Pythagorian theorem, he seems to be saying, is timeless and therefore it should not be regarded as a human construction, with a history and all the other attributes that a constructed artifact may have. The credit due to Pythagorus (or whoever was actually responsible) should be credit for formulating the theorem, not for creating it. Referring as it does to ideal entities (sides of a triangle), the theorem does not carry a date stamp. There is not a time before which it was not true and after which it was true or at least arguably true. That sounds right to me, but if some postmodern theorist wants to dispute it, I don't (for present purposes) care; it has no relevance to the kind of constructivism I am advocating.

There is one point, however, at which the issue of timeless truth impinges on educational constructivism. If the present value of a World 3 object (it's truth, applicability, or whatever) is independent of its origins and history, then why trot in all this baggage about knowledge being humanly constructed at all? Certainly if you are going to teach children that two plus two is four, it does not help the project along to explain that arithmetic is a human invention. Where the constructedness of human knowledge becomes important is in achieving what I called "enculturation into World 3" and which I have defined as "joining the ranks of those who are familiar with, understand, create, and work with the conceptual artifacts of their culture" (Bereiter, in press, Ch. 7). If students are to become enculturated into this world, it is essential but by no means sufficient that they become acquainted with the major intellectual products of past generations. "Joining the ranks" means becoming part of the knowledge-producing subculture, identifying with its members, learning their skills, sharing their aims and their norms. Except for adding some well-motivated qualifications, I don't think many people would disagree with this as a prescription for graduate studies, nor would they doubt the wisdom of earlier preparation that would ease the shock and panic some graduate students

experience when informed that getting good marks is no longer sufficient, that they are expected to "make a contribution to knowledge." The question is how early enculturation into World 3 should start. As Scardamalia's contribution (this volume) suggests, the first grade is not too soon. Others have found this also to be the case. Even at the stage of learning two plus two equals four, children can begin producing generalizable mathematical ideas and learning the norms that apply to such knowledge production (Cobb et al, 1997).

Conceptual Artifacts

Popper's term for what populates World 3, "objective knowledge," invites the sorts of criticisms Reck has made, and also criticisms by sociologists of knowledge and others who detect yet another European male declaring that his beliefs are "objectively" true. That is why I have never used Popper's term except to criticize it. A preferable term, I believe, is 'conceptual artifact.' This captures the notion of human constructions that are created to some purpose, but it avoids the ambiguity of the term 'knowledge' and leaves it quite open to what purpose such artifacts are created, how successful they are in serving that purpose, and what other unplanned purposes they may come to serve.

The qualifier 'conceptual' is meant to distinguish these kinds of artifacts from (a) material artifacts, such as ATM machines and ironing boards, and (b) immaterial but nonconceptual artifacts, such as sonatas and jokes. As with most artifacts, rigorous definition is impossible and boundary problems are to be expected.[1] The distinction between cars and trucks used to be clear but now it is fuzzy—not because analysts have discovered ambiguities but because manufacturers have started producing ambiguous vehicles. Despite its growing uncertainty, we retain the distinction between cars and trucks because it still

[1] Exceptions, of course, are theoretical and mathematical concepts. Note, however, that while the conceptual artifacts such as line, ray, angle, and triangle, as defined in Euclidean geometry, are (for instructional purposes, at least) precise, the conceptual artifact called 'Euclidean geometry' is not. It has boundary problems, problems of when it came into existence, problems of whether variations are still 'Euclidean geometry' or are something different, and so on.

serves to sort out cases warranting different treatment.[2] We should similarly expect a category like *conceptual artifacts* to be useful only insofar as it helps us sort out cases needing a certain kind of treatment. The point of my paper was that these objects making up World 3, which I now propose to call conceptual artifacts, do indeed warrant special attention in education, attention they have failed to receive partly because the two-worlds conceptual framework of educational thought has provided no place for them.

Culture and World 3

Students need to become knowledgeable about the important conceptual artifacts of their culture as well as other important cultural artifacts. I had no intention of playing down the arts and humanities nor other educational missions such as the development of values and virtues. The issue before us was what, if anything, needs to be changed to adapt liberal education to present and foreseeable circumstances. What needs changing, I have been arguing, is just that part of liberal education that deals with conceptual artifacts.

Cultural artifacts may be defined as artifacts of any kind that are preserved because of their meaning. An old household appliance may be preserved because of its historical interest or because of its artistic interest as an example of, say, art deco design. Statues and cave drawings are more obvious examples. Some cultural arftifacts are immaterial. Poems, myths, and musical compositions are examples of these. Copies of the first edition of an important literary work may be preserved as material artifacts of cultural value, but the literary work as an immaterial artifact persists through many changes in physical embodiment. When we discuss Milton's *Paradise Lost*, we are discussing the immaterial artifact rather than any material one. None of this should present any particular conceptual problems. It is therefore uncontroversial to say that a major concern

---

[2] Distinctions among artifacts not only lack the rigor some intellectuals demand, but they also shift depending on context and purpose. For determining license fees, the relevant distinctions between cars and trucks have to do with size, horsepower, and use. From the standpoint of highway safety, however, a truck is best defined as any vehicle built on a truck frame, for it is the frame that in a collision causes a truck to ride up over a sedan and kill the occupants.

of liberal education has been to develop not only a knowledge of but also a deep intimacy with a range of cultural artifacts.

What I am calling conceptual artifacts would fall within that range. Students are expected to become familiar with Newton's laws and Plato's scheme for a republic, just as they are expected to become familiar with *Paradise Lost* and the Acropolis. So why make a special case of *conceptual* artifacts? One reason, advanced by Carus (this volume) is that they include scientific theories and principles, and these are of exceptional importance in contemporary society. Although I agree, that is an argument that picks out only a particular subset of conceptual artifacts—those that assert something to be the case and, among those, assertions growing out of scientific methods (however those may be characterized). Conceptual artifacts cover much more ground than that. They are not limited as to subject matter and they include plans, problem formulations, proposals (like the one I have been advancing here), interpretations, and criticisms. Given this loose and generous inclusiveness, the question arises whether conceptual artifacts can be distinguished at all from the full range of cultural artifacts, leaving aside the question of whether there is any point in doing so.

Let me respond, however, by emphasizing that there is a point in making a distinction. With respect to most kinds of cultural artifacts, students necessarily take a rather passive, receptive role. Although, as part of their liberal education, they may engage in arts and crafts, do creative writing, put on plays and musical performances, and other such active pursuits, these are usually at some remove from the cultural objects they are studying. Their study of major works of art and literature, of historical and natural wonders, and of the lives of culturally significant people is focused on understanding and appreciation. The students are not prospective builders of new pyramids, and making models of the old ones is a questionable activity, often reserved for the less literate. With conceptual artifacts, students can and indeed should take a much more active role. The best way to understand a scientific theory, as Popper declared, is to reconstruct it, starting with trying to solve the problem that the theory was created to solve. Students can then set about actively using the theory, sometimes for practical

purposes but more often as a device for explaining phenomena. The explanations they produce are also conceptual artifacts, and they can work collaboratively to improve them. Moreover, their interaction with other kinds of cultural artifacts can be enriched and enlivened by creating conceptual artifacts related to them—interpretations, criticisms, close readings, historical accounts, psychological or sociological analyses. Although there is a danger here of overintellectualizing things that should be experienced more viscerally, a good teacher should be able to ensure a healthy balance (and help students to experience Euclid's geometry and Newton's laws viscerally as well).

What, then, distinguishes conceptual artifacts from other cultural artifacts? Keeping in mind that, as with all artifacts, strict categorization is impossible because of the possibility of creating intermediate instances, I suggest the following rough distinction: Conceptual artifacts are distinguishable from other cultural artifacts by virtue of the logical relations that may obtain among them. One conceptual artifact may be derivable from or constitute a limiting case of another; one may contradict or support another; whether two conceptual artifacts are the same or different or whether one is a modification of or improvement on the other are substantive questions. Material artifacts cannot relate to one another in these ways, and neither can such nonmaterial cultural artifacts as songs and poems. An important part of mastering any intellectual discipline is grasping the relations among ideas; the more theoretically developed the discipline, the more interdependent the ideas tend to be and the more essential it becomes to grasp them as a system. This systematicity, based on logical relations among elements, is what sets World 3 apart from the larger world of cultural artifacts, which may relate to each other in many ways but not in the particularly powerful and difficult-to-grasp way of conceptual artifacts.

The educational value of work with ideas has long been recognized. We can see it in the medieval *disputatio*, in the recurrent popularity of debate and Socratic dialogue, and even in that contemporary travesty, the critical thinking skills program. But all of these are focused on but one aspect of intellectual work, the arrival at truths or the resolution of competing truth claims. Important as this is, it makes for a static and hide-bound mental life if it is detached from the

formulation and elaboration of problems, the creation of ideas, the shaping of intuitions into discussible propositions, the exploration of what particular ideas and systems of ideas can and cannot do, and what is absolutely central to all the sciences and progressive disciplines: idea improvement. The narrow focus on issues of truth or validity also makes for an unhealthy separation of intellect from feeling. It denies to conceptual artifacts the opportunity that is allowed to all other kinds of artifacts: the opportunity to become objects of passion, to become part of the inner core of things that people really care about.

The preceding paragraph does not say anything glaringly controversial. Indeed, it is vulnerable to the death sentence that practitioners so often pass on novel educational ideas: "We already do that." Without a different epistemology, it is difficult to appreciate the gap between what actually goes on in even the most intellectually alive classrooms and what would ideally go on in what I am proposing the classroom of tomorrow should become: a workshop for the production of conceptual artifacts. Although commonsense epistemology can undergird quite a lively program of inquiry and analysis, it leaves two questions unaddressed, virtually unaddressable:

What is this idea (concept, theory, principle) good for?

How could it be improved?

As far as I can recall, I never in my years of schooling, including university and graduate school, heard these questions raised. Only in the knowledge-building classrooms Scardamalia refers to have I ever encountered them since. Unless you hold some notion of ideas as real things, and recognize work with ideas as real productive work, these questions will not arise. If there is anything to the claim that we are entering a new phase of civilization, a knowledge age, then these questions are ones that students must learn to ask and ponder.

Is The Knowledge Society a Fiction?

But is there anything to the claim that we are entering a new age? Edwards and Ogilvie cite data that raise doubts whether there is any general rise in demand for knowledge workers or that society is undergoing a massive shift comparable to the shift from an agrarian to an industrial economy. Instead, they say, what is going on is a shift from manufacturing to service jobs and a general rise in the

level of schooling required for good jobs. Their educational message is to concentrate on raising the educational floor rather than striving to extend the high end of educational attainment.

One could dismiss their arguments by saying that the present discussion is not about increasing employability but about increasing the capacity of society to take advantage of the new opportunities offered by the explosion of information and communications technology. There is a factual basis for the futuristic business literature, but it is not in labor statistics. Most of its authors are management consultants or business magazine journalists, and the material they draw on consists mostly of cases in which companies, large and small, are reorganizing in ways that make knowledge creation and innovation more central to their business. These companies are not claimed to be typical but to be front-running and therefore likely to force other companies to follow suit in order to remain competitive.

Nevertheless, the challenge Edwards and Ogilvie raise does strike close to the core of the issue with which this conference dealt. For if there is not a major economic transformation taking place, then perhaps there is no need for liberal education—or any other variety of education, for that matter—to change. See to it that kids learn to do something with computers beyond playing video games and otherwise go ahead with normal efforts to do a better job of what schools already do. That is probably a fair description of the present state of educational reform in North America and the U.K., and so it can stand as the default position on the question of educational change.

It is not evident to me, however, how the kinds of economic indicators Edwards and Ogilvie consider could confirm or disconfirm the millennialist claims. Suppose the most radical claim is right, that there is a shift underway toward a society organized around the production of knowledge, and that this is as profound as the earlier shift from a society organized around agriculture to one organized around manufacture. The earlier shift was marked by massive movement of people from the countryside into the cities and by an accompanying change in occupation. Nothing of this scale is in the offing. The supposed changes are internal to firms and occur within rather than between

large occupational categories. One of the skeptical arguments is that if knowledge workers are so much in demand their incomes ought to have risen more rapidly than they have. But this argument does not reckon with the strategies large companies now have available for holding down labor costs, which include hiring younger workers, offering nonsalary incentives, and outsourcing knowledge work to cheaper labor markets.

In short, something quite complex is going on. In the face of uncertainties, we need education policies that make sense under a variety of possible scenarios. Although radical predictions, such as Rifkin's "end of work" (1995), may prove mistaken, so may advice such as "Stay in school and you will earn a higher income." Perhaps the sounder advice would be, "Learn to do work that cannot be outsourced." The kinds of work that cannot readily be outsourced to distant places include on one hand personal services, most of which are poorly paid, and on the other hand ill-structured tasks in knowledge-rich domains (Bereiter, in press, Ch. 7). Serious reform in education systems ought to be aimed at qualifying more people for the latter. That is what they are trying to do in Singapore and Hong Kong, for instance, while the U.S. is preoccupied with "raising the floor." The question to which educators the world over lack sound answers, however, is *how* to equip students to handle ill-structured tasks in knowledge-rich domains. They may opt for well-structured tasks in knowledge-rich domains (e.g., memorizing facts) or ill-structured tasks in knowledge-poor domains (the stuff of thinking skills programs) or, the favorite these days, "project-based learning," which gives the impression of creative knowledge work without actually requiring that students do any. Knowledge building, the creation and improvement of useful conceptual artifacts, is intended as a fourth option that more directly addresses the need.

Canons to the Right of Us, Canons to the Left of Us

Demonstrating that it often takes a radical to be a true conservative, Miller eloquently argues two points:

1. That what students need to learn does not necessarily coincide with what they want to learn—or with what they want to do, which may be only incidentally productive of learning.

2. That teachers ought to know better than the students what students need to learn.

These, the pillars of 'basic education,' have often been challenged by child-centered educators, but I believe they should be allowed to stand, provided they are augmented by two additional principles:

3. That understanding, in particular, cannot be received but must be constructed or reconstructed by each person who would possess it.
4. That, as Dewey and Whitehead both insisted, subject matter should have value to students at the time they learn it; it should not merely be banked against future needs.

In my target article, I assumed the first two principles—acknowledging the first and taking the second for granted—in order to concentrate on the latter two, toward which I thought I had something to offer. It seems I miscalculated. If you do not emphatically proclaim the first two principles, then anything you say in favor of the latter two is liable to be taken as a rejection of the first two. More seriously, however, I did not deal with the tension that exists between them, a tension that any educator must live with who tries to do justice to all four principles.

Miller discusses, with a good deal of wisdom and charity, the tension between what students ought to learn, for their own and society's good, and what their peer culture disposes them toward. His optimistic conclusion, based on experience, is that if you offer what students need, some at least will come. Carus's discussion represents, among other things, a creative response to the tension between the need for a shared body of knowledge and the need for students to construct their own understanding of the world. An interesting way to think about this tension is introduced by Edwards and Ogilvie and taken further by Carus. It is to see it as a tension between education for consumption and education for production. There is no doubt that the average person has more impact on society as a consumer than as a producer. Schools invest considerable effort in trying to discourage bad consumption, with their anti-drug and other "just say no" campaigns. But where is the effort to encourage good consumption—that is, socially beneficial consumption? Isolated in those few

remaining bastions of traditional liberal education, it would seem. But Carus's point is not about the social value of education for consumption, it is on its value in helping students to develop a personal mental life.

Carus's chapter, as I read it now, addresses a question that I raised at the beginning of my target article but that neither I nor any of the other participants addressed directly: "What should it mean to be an educated person in the 21st century?" That is, what should it mean *on the whole* to be an educated person under the rapidly changing conditions of 21st century civilization? This is not a question that can be answered by citing specifics of what is to be learned or indicating changes that need to be made. It is a question about what holds all those separate learnings together.

We do not expect students on their own to discover universal gravitation, the circulation of the blood, or the rise and decline of feudalism; but it seems that we expect them to discover Western thought (or whatever is the leading thought of their culture) on their own. Thus the implicit belief, shared by traditionalists and constructivists alike, is that students need help with the pieces but the whole will take care of itself. Carus challenges this belief. There has to be a canon, he argues, but it cannot in this century be the canon of an all-encompassing religion or of an institutionalized and static world view or body of wisdom. It needs to be a canon that itself embraces the continual advancement of knowledge, the continual breakdown of established beliefs, and the unending clashes between world views. Science, he somewhat reluctantly concludes, provides the closest things we have to such a canon, and so fostering scientific thinking should be central to modern liberal education.

I pretty much agree with all of this, with one caveat: The whole does, necessarily, take care of itself. It is the product of self-organization among pieces that do not embody all the essential characteristics of the whole. But, then, education of every sort can be viewed as only an effort to help and influence self-organization. The question is whether education needs to do something about the whole over and above what is done about the pieces. I argued the negative years ago in an article whose title summarized the argument: "The Soul Builds Its Mansion by Nailing Pieces Together" (Bereiter, 1970). Where I have moved since

is to a position shared by many contemporary educators, including Wells and Scardamalia, whose contributions to this volume reflect it: Schools need to create a social environment in which the active construction of knowledge can flourish. It needs to be more than the nurturant environment beloved of child-centered educators (e.g., Weber, 1971). It needs to embody the modern canon, which is one of sustained effort to advance the frontiers of knowledge.

Despite considerable differences, I think there is a common impulse in the contributions of Carus, Wells, Scardamalia, and myself, toward a kind of education attuned to the modern dynamism. We have provided different perspectives—philosophical, semiotic, psychological, anthropological—on what may or may not prove to be the same object. We cannot tell for sure from the accounts given, much less synthesize a coherent description of the whole object. It is early days, but it does look as if an important shift is taking place in what it means to be an educated person. First there was the educated person as the embodiment of an inherited high culture. Along about the time of Sputnik came the realization that education must adapt to the continual advancement of knowledge. "Lifelong learning," which earlier meant a lifetime devoted to absorbing the riches of the past, now took on the added (or alternative) meaning of keeping up with progress. With the coming, whether imaginary or not, of the "new economy" driven by innovation, a whole new level of expectation is introduced: The educated person of the 21st century is one who can help create progress, not merely keep up with it.

Making It Happen

My target article had little to say about how to bring about the educational change it advocated or indeed what a knowledge-building approach to education would be like. The chapters by Wells and Scardamalia partly remedy this lack, but they too stop well short of how-to-do-it advice. What they do convey, however, is how formidable a job of change is required.

"Systemic change" is the well-chosen label for what is required (Smith & O'Day, 1991). There is an interlocking, self-perpetuating system comprising regulations, curriculum guidelines, textbook adoption procedures, achievement tests, credentialing, teacher education, and teacher development that constitute

institutionalized stupidity, or what Harriet Tyson-Bernstein (1988) more generously termed "a conspiracy of good intentions." In order to succeed, systemic change has to reach down into the classroom, and in order for systemic change to occur there, a radical shift has to occur in how teachers conceive of their jobs. From their earliest preservice training to the workshops and teacher's guides that are supposed to advance their professional development, teachers are encouraged to believe that their job is to enact procedures that result in learning. Procedures become the focus. Ideas serve to justify procedures rather than procedures serving to implement ideas.

Project-based learning, hands-on learning, process-writing, strategy instruction, whole language, cross-curriculum integration, and teaching for multiple intelligences are all currently degraded practices that came about when an idea of some merit was reduced to procedures. Once the idea is reduced to procedures, two things happen. The procedures take on a life of their own, come to be valued in their own right, and evolve in ways that drift away from the original idea. The more serious consequence, however, is that criticism and efforts at improvement focus on the procedures and the original idea fails to undergo criticism and improvement. That can happen with knowledge building as much as with any other idea, unless it can somehow be saved from the degenerative effects of proceduralization.

Wells's quotations from teachers show how much has to be overcome in order to break loose from procedures, including supposedly progressive ones, and to confront the realities of the intellectual life of a classroom community. The fundamental shift, as I see it, is from managing the activities of the classroom to transforming its intellectual life. Of course, teachers must continue to manage classroom activities. Much of what has gone by names such as progressive education and child-centered education consists of changing the style of management so as to provide more scope for student initiatives. Such loosening of authority has proved troublesome enough, but it is nothing compared to turning the classroom into a community in which the students themselves begin to assume responsibility for its intellectual life—for raising its collective

intelligence, as Scardamalia says; for advancing the "state of knowledge," as that term is understood in the learned disciplines.

The teachers in DICEP perceptively pointed out to Wells the difficulty they were having in creating knowledge building communities because they had never experienced such a community themselves. This is the basic "can't get there from here" problem of all utopian schemes, and we must face up to the fact that the idea of schools as knowledge building communities is utopian. In order to "get there," people need, first of all, concepts that make the utopian ideas discussible in realistic terms. That is what much of the discourse in this book is about. Secondly, they need images—images of the possible that are strong enough to hold firm against the degenerative forces of reduction to procedures. Wells and Scardamalia provide glimpses; but compelling images would require a different kind of book from this—or, better perhaps, a multimedia document devoted to making such images clear and memorable.

Finally, teachers need something concrete to take the place of recipes. Scardamalia's chapter offers reason to believe that the right kind of technology may provide this. That is, the technology may provide tools, supports, channels of communication, a kind of environment in which teachers and students can bootstrap a classroom culture different from anything they have previously experienced. An important question is the extent to which software to support a new educational vision needs to embody that vision. There is no logical requirement that it do so. The sculptor's chisel does not have to embody a vision of the statue. However, it seems to me that a if a technology is to stand a chance against the "conspiracy of good intentions" it has to consist of more than neutral tools. Like Knowledge Forum, it needs to have a very strong pedagogical bias—yet it must not proceduralize or micromanage the learning process, lest it become itself part of the "conspiracy."

The idea of "collective cognitive responsibility" advanced by Scardamalia can help to clarify some of the ideas at issue in this conference. When I suggest that classrooms should become "workshops for the production of knowledge," this may bring to mind the image of children at their desks stitching away like so many little tailors in a garment factory. That is an altogether wrong image, and it

is not much improved by introducing notions of cooperation or collaboration. What is missing is a common mission, shared by teacher and students alike, the success of which counts for something beyond the achievements of individual students. Although it may be expressed differently for different audiences, what the knowledge-building mission comes down to is making the world more intelligible. Improving one's personal knowledge and competence is a different, though complementary task, and one that can benefit from participation in collaborative knowledge building. It is work in World 2, whereas collective cognitive responsibility, as I construe it, has to do with work in World 3. Both are important in education, but traditionally no distinction is made and the result is an often dreary amalgam in which the student's responsibility is limited to carrying out tasks that the teacher believes will result in the desired learning. Scardamalia and I discuss the World 2 aspect of cognitive responsibility in an article on "intentional cognition," attached as an appendix to this volume. Its concern is with the student's development of a mental life. The work Scardamalia reports on collective cognitive responsibility is concerned with students' contributions to advancing the state of knowledge in their community. Both of these figure in Carus's treatment of moral expertise and the canon. You can, as you prefer, regard them as two sides of the same coin, as the same task carried on at two different levels, or as complementary work in two worlds; what matters is that the two not be allowed to settle into the conceptual sludge of turn-of-the-century pedagogy.

Extending the Limits

One legitimate concern about liberal education is that it seems to represent a timeless pedagogy, which is from time to time applied to new content. The tacit assumption, accordingly, is that pedagogy is not one of those disciplines, like medicine and engineering, that has a leading edge, that keeps advancing the limits of what it is able to do. Rather, education is like politics and law, professions that do their work well or badly and that are sometimes in need of reform, that occasionally sport innovations, but that have no internal dynamic that generates progress.

The past century has seen wave after wave of educational reforms that gathered energy for a while and then lost it. If we ignore the counter-reforms and look at those that struck people as new and exciting, we find that almost always the perceived novelty lay in a more enlightened and humane way of carrying on the process of education. It was not perceived as extending the limits of the possible. Yet what we find and indeed have come to demand in most other aspects of modern life is a continual expansion of the possible. We expect cures for the previously incurable, travel to the previously unreachable, explanations of the previously incomprehensible, and low-cost versions of the previously unaffordable. To a large extent these expectations are being met. But not only has this failed to occur in education, it has not even occurred to people to look for it. People cannot imagine what an extension of the limits of the possible in education would mean. I believe the work that casts students as legitimate creators of knowledge is the first in a very long time to be seriously trying to extend the limits. If it succeeds—and the indications so far are positive—this may finally herald a reform that can survive.

References

Bereiter, C. (in press). Education and mind in the knowledge age. Mahwah, NJ: Lawrence Erlbaum Associates.

Bereiter, C. (1970). A last word: The soul builds its mansions by nailing boards together. *Interchange*, *1*(1), 49-51.

Bereiter, C., & Scardamalia, M. (1993). Surpassing ourselves: An inquiry into the nature and implications of expertise. La Salle, IL: Open Court.

Berger, P. L., & Luckmann, T. (1967). The social construction of reality. Garden City, NY: Doubleday Anchor Books.

Bloor, D. (1998). The strong program in the sociology of science. In R. Klee (Ed.), Scientific inquiry: Readings in the philosophy of science (pp. 241-250). New York: Oxford University Press.

Cobb, P., Gravmeijer, K., Yackel, E., McClain, K., & Whitenack, J. (1997). Mathematizing and symbolizing: The emergence of chains of significance in one first-grade classroom. In D. Kirshner & J. A. Whitson (Eds.), Situated

cognition: Social, semiotic, and psychological perspectives (pp. 151-233). Mahwah, NJ: Erlbaum.

Kuhn, T. (1970). The structure of scientific revolutions. Chicago: University of Chicago Press.

Moursund, D. (1999). Project-based learning using information technology. Eugene, OR: International Society for Technology in Education.

Nonaka, I., & Takeuchi, H. (1995). The knowledge creating company. New York: Oxford University Press.

Ohlsson, S. (1991). Young adults' understanding of evolutionary explanations: Preliminary observations (Tech. Rep. to OERI) University of Pittsburgh, Learning Research and Development Laboratory.

Ohlsson, S. (1993). Abstract schemas. Educational Psychologist, 28(1), 51-61.

Popper, K. R., & Eccles, J. C. (1977). The self and its brain. Berlin: Springer-Verlag.

Rifkin, J. (1995). The end of work: The decline of the global labor force and the dawn of the post-market era. New York: Tarcher/Putnam.

Simonton, D. K. (1999). Origins of genius: Darwinian perspectives on creativity. New York: Oxford University Press.

Smith, M. B., & O'Day, J. (1991). Systemic school reform. The politics of curriculum and testing, 233-267.

Tyson-Bernstein, H. (1988). A conspiracy of good intentions. Washington, DC: Council for Basic Education.

Whitehead, A. N. (1929). The aims of education. New York: Macmillan.