

Situated Cognition and the Culture of Learning

by John Seely Brown, Allan Collins and Paul Duguid
Educational Researcher; v18 n1, pp. 32-42, Jan-Feb 1989.

Abstract: Many teaching practices implicitly assume that conceptual knowledge can be abstracted from the situations in which it is learned and used. This article argues that this assumption inevitably limits the effectiveness of such practices. Drawing on recent research into cognition as it is manifest in everyday activity, the authors argue that knowledge is situated, being in part a product of the activity, context, and culture in which it is developed and used. They discuss how this view of knowledge affects our understanding of learning, and they note that conventional schooling too often ignores the influence of school culture on what is learned in school. As an alternative to conventional practices, they propose cognitive apprenticeship (Collins, Brown, Newman, in press), which honors the situated nature of knowledge. They examine two examples of mathematics instruction that exhibit certain key features of this approach to teaching.

The breach between learning and use, which is captured by the folk categories "know what" and "know how," may well be a product of the structure and practices of our education system. Many methods of didactic education assume a separation between knowing and doing, treating knowledge as an integral, self-sufficient substance, theoretically independent of the situations in which it is learned and used. The primary concern of schools often seems to be the transfer of this substance, which comprises abstract, decontextualized formal concepts. The activity and context in which learning takes place are thus regarded as merely ancillary to learning---pedagogically useful, of course, but fundamentally distinct and even neutral with respect to what is learned.

Recent investigations of learning, however, challenge this separating of what is learned from how it is learned and used.¹ The activity in which knowledge is developed and deployed, it is now argued, is not separable from or ancillary to learning and cognition. Nor is it neutral. Rather, it is an integral part of what is learned. Situations might be

said to co-produce knowledge through activity. Learning and cognition, it is now possible to argue, are fundamentally situated.

In this paper, we try to explain in a deliberately speculative way, why activity and situations are integral to cognition and learning, and how different ideas of what is appropriate learning activity produce very different results. We suggest that, by ignoring the situated nature of cognition, education defeats its own goal of providing useable, robust knowledge. And conversely, we argue that approaches such as cognitive apprenticeship (Collins, Brown, & Newman, in press) that embed learning in activity and make deliberate use of the social and physical context are more in line with the understanding of learning and cognition that is emerging from research.

Situated Knowledge and Learning

Miller and Gildea's (1987) work on vocabulary teaching has shown how the assumption that knowing and doing can be separated leads to a teaching method that ignores the way situations structure cognition. Their work has described how children are taught words from dictionary definitions and a few exemplary sentences, and they have compared this method with the way vocabulary is normally learned outside school.

People generally learn words in the context of ordinary communication. This process is startlingly fast and successful. Miller and Gildea note that by listening, talking, and reading, the average 17-year-old has learned vocabulary at a rate of 5,000 words per year (13 per day) for over 16 years. By contrast, learning words from abstract definitions and sentences taken out of the context of normal use, the way vocabulary has often been taught, is slow and generally unsuccessful. There is barely enough classroom time to teach more than 100 to 200 words per year. Moreover, much of what is taught turns out to be almost useless in practice. They give the following examples of students' uses of vocabulary acquired this way:

"Me and my parents correlate, because without them I wouldn't be here."

"I was meticulous about falling off the cliff."

"Mrs. Morrow stimulated the soup."²

Given the method, such mistakes seem unavoidable. Teaching from dictionaries assumes that definitions and exemplary sentences are self-contained "pieces" of knowledge. But words and sentences are not islands, entire unto themselves. Language use would involve an unremitting confrontation with ambiguity, polysemy, nuance, metaphor, and so forth were these not resolved with the extralinguistic help that the context of an utterance provides (Nunberg, 1978).

Prominent among the intricacies of language that depend on extralinguistic help are indexical words --- words like I, here, now, next, tomorrow, afterwards, this. Indexical terms are those that "index" or more plainly point to a part of the situation in which communication is being conducted.³ They are not merely context-sensitive; they are completely context-dependent. Words like I or now, for instance, can only be interpreted in the 'context of their use. Surprisingly, all words can be seen as at least partially indexical (Barwise & Perry, 1983).

Experienced readers implicitly understand that words are situated. They, therefore, ask for the rest of the sentence or the context before committing themselves to an interpretation of a word. And they go to dictionaries with situated examples of usage in mind. The situation as well as the dictionary supports the interpretation. But the students who produced the sentences listed had no support from a normal communicative situation. In tasks like theirs, dictionary definitions are assumed to be self-sufficient. The extralinguistic props that would structure, constrain, and ultimately allow interpretation in normal communication are ignored.

Learning from dictionaries, like any method that tries to teach abstract concepts independently of authentic situations, overlooks the way understanding is developed through continued, situated use. This development, which involves complex social negotiations, does not crystallize into a categorical definition. Because it is dependent on situations and negotiations, the meaning of a word cannot, in principle, be captured by a definition, even when the definition is supported by a couple of exemplary sentences.

All knowledge is, we believe, like language. Its constituent parts index the world and so are inextricably a product of the activity and situations in which they are produced. A concept, for example, will continually evolve with each new occasion of use, because new situations, negotiations, and activities inevitably recast it in a new, more densely textured form. So a concept, like the meaning of a word,

is always under construction. This would also appear to be true of apparently well-defined, abstract technical concepts. Even these are not wholly definable and defy categorical description; part of their meaning is always inherited from the context of use.

Learning and tools. To explore the idea that concepts are both situated and progressively developed through activity, we should abandon any notion that they are abstract, self-contained entities. Instead, it may be more useful to consider conceptual knowledge as, in some ways, similar to a set of tools.⁴ Tools share several significant features with knowledge: They can only be fully understood through use, and using them entails both changing the user's view of the world and adopting the belief system of the culture in which they are used.

First, if knowledge is thought of as tools, we can illustrate Whitehead's (1929) distinction between the mere acquisition of inert concepts and the development of useful, robust knowledge. It is quite possible to acquire a tool but to be unable to use it. Similarly, it is common for students to acquire algorithms, routines, and decontextualized definitions that they cannot use and that, therefore, lie inert. Unfortunately, this problem is not always apparent. Old-fashioned pocket knives, for example, have a device for removing stones from horses' hooves. People with this device may know its use and be able to talk wisely about horses, hooves, and stones. But they may never betray --- or even recognize --- that they would not begin to know how to use this implement on a horse. Similarly, students can often manipulate algorithms, routines, and definitions they have acquired with apparent competence and yet not reveal, to their teachers or themselves, that they would have no idea what to do if they came upon the domain equivalent of a limping horse.

People who use tools actively rather than just acquire them, by contrast, build an increasingly rich implicit understanding of the world in which they use the tools and of the tools themselves. The understanding, both of the world and of the tool, continually changes as a result of their interaction. Learning and acting are interestingly indistinct, learning being a continuous, life-long process resulting from acting in situations.

Learning how to use a tool involves far more than can be accounted for in any set of explicit rules. The occasions and conditions for use arise directly out of the context of activities of each community that uses the tool, framed by the way members of that community see the world. The community and its viewpoint, quite as much as the tool

itself, determine how a tool is used. Thus, carpenters and cabinet makers use chisels differently. Because tools and the way they are used reflect the particular accumulated insights of communities, it is not possible to use a tool appropriately without understanding the community or culture in which it is used.

Conceptual tools similarly reflect the cumulative wisdom of the culture in which they are used and the insights and experience of individuals. Their meaning is not invariant but a product of negotiation within the community. Again, appropriate use is not simply a function of the abstract concept alone. It is a function of the culture and the activities in which the concept has been developed. Just as carpenters and cabinet makers use chisels differently, so physicists and engineers use mathematical formulae differently. Activity, concept, and culture are interdependent. No one can be totally understood without the other two. Learning must involve all three. Teaching methods often try to impart abstracted concepts as fixed, well-defined, independent entities that can be explored in prototypical examples and textbook exercises. But such exemplification cannot provide the important insights into either the culture or the authentic activities of members of that culture that learners need.

To talk about academic disciplines, professions, or even manual trades as communities or cultures will perhaps seem strange. Yet communities of practitioners are connected by more than their ostensible tasks. They are bound by intricate, socially constructed webs of belief, which are essential to understanding what they do (Geertz, 1983). The activities of many communities are unfathomable, unless they are viewed from within the culture. The culture and the use of a tool act together to determine the way practitioners see the world; and the way the world appears to them determines the culture's understanding of the world and of the tools. Unfortunately, students are too often asked to use the tools of a discipline without being able to adopt its culture. To learn to use tools as practitioners use them, a student, like an apprentice, must enter that community and its culture. Thus, in a significant way, learning is, we believe, a process of enculturation.

Learning and enculturation. Enculturating may, at first, appear to have little to do with learning. But it is, in fact, what people do in learning to speak, read, and write, or becoming school children, office workers, researchers, and so on. From a very early age and throughout their lives, people, consciously or unconsciously, adopt the behavior and belief systems of new social groups. Given the chance to

observe and practice in situ the behavior of members of a culture, people pick up relevant jargon, imitate behavior, and gradually start to act in accordance with its norms. These cultural practices are often recondite and extremely complex. Nonetheless, given the opportunity to observe and practice them, people adopt them with great success. Students, for instance, can quickly get an implicit sense of what is suitable diction, what makes a relevant question, what is legitimate or illegitimate behavior in a particular activity. The ease and success with which people do this (as opposed to the intricacy of describing what it entails) belie the immense importance of the process and obscures the fact that what they pick up is a product of the ambient culture rather than of explicit teaching.

Too often the practices of contemporary schooling deny students the chance to engage the relevant domain culture, because that culture is not in evidence. Although students are shown the tools of many academic cultures in the course of a school career, the pervasive cultures that they observe, in which they participate, and which some enter quite effectively are the cultures of school life itself. These cultures can be unintentionally antithetical to useful domain learning. The ways schools use dictionaries, or math formulae, or historical analysis are very different from the ways practitioners use them (Schoenfeld, in press). Thus, students may pass exams (a distinctive part of school cultures) but still not be able to use a domain's conceptual tools in authentic practice.

This is not to suggest that all students of math or history must be expected to become professional mathematicians or historians, but to claim that in order to learn these subjects (and not just to learn about them) students need much more than abstract concepts and self-contained examples. They need to be exposed to the use of a domain's conceptual tools in authentic activity --- to teachers acting as practitioners and using these tools in wrestling with problems of the world. Such activity can tease out the way a mathematician or historian looks at the world and solves emergent problems. The process may appear informal, but it is nonetheless full-blooded, authentic activity that can be deeply informative --- in a way that textbook examples and declarative explanations are not.

Authentic Activity

Our case so far rests on an undefined distinction between authentic and school activity. If we take learning to be a process of enculturation, it is possible to clarify this distinction and to explain why

much school work is inauthentic and thus not fully productive of useful learning.

The activities of a domain are framed by its culture. Their meaning and purpose are socially constructed through negotiations among present and past members. Activities thus cohere in a way that is, in theory, if not always in practice, accessible to members who move within the social framework. These coherent, meaningful, and purposeful activities are authentic, according to the definition of the term we use here. Authentic activities then, are most simply defined as the ordinary practices of the culture.

This is not to say that authentic activity can only be pursued by experts. Apprentice tailors (Lave, 1988a), for instance, begin by ironing finished garments (which tacitly teaches them a lot about cutting and sewing). Ironing is simple, valuable, and absolutely authentic. Students of Palincsar and Brown's (1984) reciprocal teaching of reading may read elementary texts, but they develop authentic strategies that are recognized by all readers. The students in Miller and Gildea's study, by contrast, were given a strategy that is a poor extrapolation of experienced readers' situated use of dictionaries.

School activity too often tends to be hybrid, implicitly framed by one culture, but explicitly attributed to another. Classroom activity very much takes place within the culture of schools, although it is attributed to the culture of readers, writers, mathematicians, historians, economists, geographers, and so forth. Many of the activities students undertake are simply not the activities of practitioners and would not make sense or be endorsed by the cultures to which they are attributed. This hybrid activity, furthermore, limits students' access to the important structuring and supporting cues that arise from the context. What students do tends to be ersatz activity.

Archetypal school activity is very different from what we have in mind when we talk of authentic activity, because it is very different from what authentic practitioners do. When authentic activities are transferred to the classroom, their context is inevitably transmuted; they become classroom tasks and part of the school culture. Classroom procedures, as a result, are then applied to what have become classroom tasks. The system of learning and using (and, of course, testing) thereafter remains hermetically sealed within the self-confirming culture of the school. Consequently, contrary to the aim of schooling, success within this culture often has little bearing on performance elsewhere.

Math word problems, for instance, are generally encoded in a syntax and diction that is common only to other math problems. Thus the word problems of a textbook of 1478 are instantly recognizable today (Lave, 1988c). But word problems are as foreign to authentic math practice as Miller and Gildea's example of dictionary learning is to the practices of readers and writers. By participating in such ersatz activities students are likely to misconceive entirely what practitioners actually do. As a result, students can easily be introduced to a formalistic, intimidating view of math that encourages a culture of math phobia rather than one of authentic math activity.

In the creation of classroom tasks, apparently peripheral features of authentic tasks --- like the extralinguistic supports involved in the interpretation of communication --- are often dismissed as "noise" from which salient features can be abstracted for the purpose of teaching. But the context of activity is an extraordinarily complex network from which practitioners draw essential support. The source of such support is often only tacitly recognized by practitioners, or even by teachers or designers of simulations. Classroom tasks, therefore, can completely fail to provide the contextual features that allow authentic activity. At the same time, students may come to rely, in important but little noticed ways, on features of the classroom context, in which the task is now embedded, that are wholly absent from and alien to authentic activity. Thus, much of what is learned in school may apply only to the ersatz activity, if it was learned through such activity.

Activities of students, practitioners, and just plain folks. The idea that most school activity exists in a culture of its own is central to understanding many of the difficulties of learning in school. Jean Lave's ethnographic studies of learning and everyday activity (1988b) reveal how different schooling is from the activities and culture that give meaning and purpose to what students learn elsewhere. Lave focuses on the behavior of JPFs (just plain folks) and records that the ways they learn are quite distinct from what students are asked to do.

Three categories primarily concern us here: JPFs, students, and practitioners. Put most simply, when JPFs aspire to learn a particular set of practices, they have two apparent options. First, they can enculturate through apprenticeship. Becoming an apprentice doesn't involve a qualitative change from what JPFs normally do. People enculturate into different communities all the time. The apprentices' behavior and the JPFs' behavior can thus be thought of as pretty much the same.⁵

The second, and now more conventional, option is to enter a school as a student. Schools, however, do seem to demand a qualitative change in behavior. What the student is expected to do and what a JPF does are significantly different. The student enters the school culture while ostensibly being taught something else. And the general strategies for intuitive reasoning, resolving issues, and negotiating meaning that people develop through everyday activity are superseded by the precise, well-defined problems, formal definitions, and symbol manipulation of much school activity.

We try to represent this discontinuity in Table 1, which compares salient features of JPF, practitioner, and putative student behavior.

TABLE 1.

JPF, Practitioner, and Student Activity

	JPFs	Students	Practitioners
reasoning with:	casual stories	laws	casual models
acting on:	situations	symbols	conceptual situations
resolving:	emergent problems and dilemmas	well-defined problems	ill-defined problems
producing:	negotiable meaning & socially constructed understanding	fixed meaning & immutable concepts	negotiable meaning & socially constructed understanding

This Table is intended mainly to make apparent that, in our terms, there is a great similarity between JPFs' and practitioners' activity. Both have their activities situated in the cultures in which they work, within which they negotiate meanings and construct understanding. The issues and problems that they face arise out of, are defined by, and are resolved within the constraints of the activity they are pursuing.

Lave's work (1988b) provides a good example of a JPF engaged in authentic activity using the context in which an issue emerged to help

find a resolution. The example comes from a study of a Weight Watchers class, whose participants were preparing their carefully regulated meals under instruction.

In this case they were to fix a serving of cottage cheese, supposing the amount laid out for the meal was three-quarters of the two thirds cup the program allowed. The problem solver in this example began the task muttering that he had taken a calculus course in college... Then after a pause he suddenly announced that he had "got it!" From then on he appeared certain he was correct, even before carrying out the procedure. He filled a measuring-cup two thirds full of cottage cheese, dumped it out on the cutting board, patted it into a circle, marked a cross on it, scooped away one quadrant, and served the rest.

Thus, "take three-quarters of two-thirds of a cup of cottage cheese" was not just the problem statement but also the solution to the problem and the procedure for solving it. The setting was part of the calculating process and the solution was simply the problem statement, enacted with the setting. At no time did the Weight Watcher check his procedure against a paper and pencil algorithm, which would have produced $\frac{3}{4} \text{ cup} \times \frac{2}{3} \text{ cup} = \frac{1}{2} \text{ cup}$. Instead, the coincidence of the problem, setting, and enactment was the means by which checking took place. (p. 165)

The dieter's solution path was extremely expedient and drew on the sort of inventiveness that characterizes the activity of both JPFs and practitioners. It reflected the nature of the activity, the resources available, and the sort of resolution required in a way that problem solving that relies on abstracted knowledge cannot.

This inventive resolution depended on the dieter seeing the problem in the particular context which itself was embedded in ongoing activity. And this again is characteristic of both JPFs and experts. The dieter's position gave him privileged access to the solution path he chose. (This probably accounts for the certainty he expressed before beginning his calculation.) He was thus able to see the problem and its resolution in terms of the measuring cup, cutting board, and knife. Activity-tool-culture (cooking-kitchen utensils-dieting) moved in step throughout this procedure because of the way the problem was seen and the task was performed. The whole micro-routine simply became one more step on the road to a meal.⁶ Knowing and doing were interlocked and inseparable.

This sort of problem solving is carried out in conjunction with the environment and is quite distinct from the processing solely inside heads that many teaching practices implicitly endorse. By off-loading part of the cognitive task onto the environment, the dieter automatically used his environment to help solve the problem. His actions were not in any way exceptional; they resemble many ordinary working practices. Scribner (1984) records, for instance, how complex calculations can be performed by practitioners using their environment directly. In the case she studied, dairy loaders used the configuration of crates they were filling and emptying almost like an elaborate abacus. Nor are such problem solving strategies limited to the physical or social environment. This sort of reliance on situations can be seen in the work of physicists, who see "through" formulae by envisioning a physical situation, which then provides support for inferences and approximations (deKleer & Brown, 1984). Hutchins' (in press) study of intricate collaborative naval navigation records the way people distribute the burden across the environment and the group as well. The resulting cognitive activity can then only be explained in relation to its context. "[W]hen the context of cognition is ignored," Hutchins observes, "it is impossible to see the contribution of structure in the environment, in artifacts, and in other people to the organization of mental processes.

Instead of taking problems out of the context of their creation and providing them with an extraneous framework, JPFs seem particularly adept at solving them within the framework of the context that produced them. This allows JPFs to share the burdens of both defining and solving the problem with the task environment as they respond in "real time." The adequacy of the solution they reach becomes apparent in relation to the role it must play in allowing activity to continue. The problem, the solution, and the cognition involved in getting between the two cannot be isolated from the context in which they are embedded.

Even though students are expected to behave differently, they inevitably do behave like the JPFs they are and solve most of their problems in their own situated way. Schoenfeld (in press) describes mathematics students using well-known but unacknowledged strategies, such as the position of a problem in a particular section of the book (e.g., the first questions at the end of chapters are always simple ones, and the last usually demand concepts from earlier chapters) or the occurrence of a particular word in the problem (e.g., "left" signals a subtraction problem), to find solutions quickly and efficiently. Such ploys indicate how thoroughly learners really are

situated, and how they always lean on whatever context is available for help. Within the practices of schooling this can obviously be very effective. But the school situation is extremely specialized. Viewed from outside, where problems do not come in textbooks, a dependency on such school-based cues makes the learning extremely fragile.

Furthermore, though schooling seeks to encourage problem solving, it disregards most of the inventive heuristics that students bring to the classroom. It thus implicitly devalues not just individual heuristics, which may be fragile, but the whole process of inventive problem solving. Lave (1988c) describes how some students feel it necessary to disguise effective strategies so that teachers believe the problems have been solved in the approved way.

Structuring activity. Authentic activity, as we have argued, is important for learners, because it is the only way they gain access to the standpoint that enables practitioners to act meaningfully and purposefully. It is activity that shapes or hones their tools. How and why remain to be explained. Activity also provides experience, which is plainly important for subsequent action. Here, we try to explain some of the products of activity in terms of idiosyncratic "indexicalized" representations.

Representations arising out of activity cannot easily (or perhaps at all) be replaced by descriptions. Plans, as Suchman argues (1987), are distinct from situated actions. Most people will agree that a picture of a complex machine in a manual is distinctly different from how the machine actually looks. (In an intriguing way you need the machine to understand the manual, as much as the manual to understand the machine.) The perceptions resulting from actions are a central feature in both learning and activity. How a person perceives activity may be determined by tools and their appropriated use. What they perceive, however, contributes to how they act and learn. Different activities produce different indexicalized representations not equivalent, universal ones. And, thus, the activity that led to those representations plays a central role in learning.

Representations are, we suggest, indexicalized rather in the way that language is. That is to say, they are dependent on context. In face-to-face conversations, people can interpret indexical expressions (containing such words as I, you, here, now, that, etc.), because they have access to the indexed features of the situation, though people rarely notice the significance of the surroundings to their understanding. The importance of the surroundings becomes apparent,

however, when they try to hold similar conversations at a distance. Then indexical expressions become problematic until ways are found to secure their interpretation by situating their reference (see, for instance, Rubin, 1980, on the difference between speech and writing).

Perhaps the best way to discover the importance and efficiency of indexical terms and their embedding context is to imagine discourse without them. Authors of a collaborative work such as this one will recognize the problem if they have ever discussed the paper over the phone. "What you say here" is not a very useful remark. Here in this setting needs an elaborate description (such as "page 3, second full paragraph, fifth sentence," beginning...) and can often lead to conversations at cross purposes. The problem gets harder in conference calls when you becomes as ambiguous as here is unclear. The contents of a shared environment make a central contribution to conversation.

When the immediacy of indexical terms is replaced by descriptions, the nature of discourse changes and understanding becomes much more problematic. Indexical terms are virtually transparent. They draw little or no attention to themselves. They do not necessarily add significantly to the difficulty of understanding a proposition in which they occur, but simply point to the subject under discussion, which then provides essential structure for the discourse. Descriptions, by comparison, are at best translucent and at worst opaque, intruding emphatically between speakers and their subjects. The audience has first to focus on the descriptions and try to interpret them and find what they might refer to. Only then can the proposition in which they are embedded be understood. (However elaborate, a description does not merely replace the indexical word.) The more elaborate the description is in an attempt to be unambiguous, the more opaque it is in danger of becoming. And in some circumstances, the indexical term simply cannot be replaced (Perry, 1979).

Knowledge, we suggest, similarly indexes the situation in which it arises and is used. The embedding circumstances efficiently provide essential parts of its structure and meaning. So knowledge, which comes coded by and connected to the activity and environment in which it is developed, is spread across its component parts, some of which are in the mind and some in the world much as the final picture on a jigsaw is spread across its component pieces.

As Hutchins (in press), Pea (1988), and others point out, the structure of cognition is widely distributed across the environment, both social

and physical. And we suggest that the environment, therefore, contributes importantly to indexical representations people form in activity. These representations, in turn, contribute to future activity. Indexical representations developed through engagement in a task may greatly increase the efficiency with which subsequent tasks can be done, if part of the environment that structures the representations remains invariant. This is evident in the ability to perform tasks that cannot be described or remembered in the absence of the situation. Recurring features of the environment may thus afford recurrent sequences of actions. Memory and subsequent actions, as knots in handkerchiefs and other aides memoires reveal, are not context-independent processes. Routines (Agre, 1985) may well be a product of this sort of indexicalization. Thus, authentic activity becomes a central component of learning.

One of the key points of the concept of indexicality is that it indicates that knowledge, and not just learning, is situated. A corollary of this is that learning methods that are embedded in authentic situations are not merely useful; they are essential.

Learning Through Cognitive Apprenticeship

We have been working toward a conception of human learning and reasoning that, we feel, it is important for school practices to honor. Though there are many innovative teachers, schools, and programs that act otherwise, prevalent school practices assume, more often than not, that knowledge is individual and self-structured, that schools are neutral with respect to what is learned, that concepts are abstract, relatively fixed, and unaffected by the activity through which they are acquired and used, and that JPF behavior should be discouraged.

Cognitive apprenticeship (Collins, Brown, & Newman, in press), whose mechanisms we have, to some extent, been trying to elucidate, embraces methods that stand in contradistinction to these practices. Cognitive apprenticeship methods try to enculturate students into authentic practices through activity and social interaction in a way similar to that evident --- and evidently successful --- in craft apprenticeship. In this section, we examine briefly two examples of mathematics teaching in an attempt to illustrate how some of the characteristics of learning that we have discussed can be honored in the classroom. We use examples from mathematics in part because that is where some of the most innovative work in teaching can be

found. But we firmly believe that this sort of teaching is not just possible in mathematics.

Schoenfeld's teaching of problem solving. Schoenfeld's teaching of problem solving (1985, in press) deliberately attempts to generate mathematical practice and to show college students how to think mathematically about the world, how to see the world through mathematicians' eyes, and, thus, how to use the mathematician's tools. His approach goes well beyond simply giving students problem-solving strategies. Much more importantly, it provides students with the opportunity to enter the culture of mathematical practice.


Schoenfeld's students bring problems to class that he and they investigate mathematically. His students can witness and participate in spontaneous mathematical thinking and see mathematics as a sense-making pursuit. This approach is distinctive because, before graduate school, few students get the opportunity to see their teachers engaged in mathematical practice, yet the students are expected to understand the nature of that practice.

In one case (Schoenfeld, in press), he and his class faced the problem of the magic square (see Figure 1). Though the problem is relatively straightforward, the collaborative work involved in solving it and, importantly, in analyzing the solution helped reveal to the class the way mathematicians look at problems. The class worked collectively through a number of strategies, which, on reflection, they recognized as more general and more powerful mathematical ideas. In discussing whether 9 can go in the center of the square, they developed the ideas of "focusing on key points that give leverage," and "exploiting extreme cases." Although Schoenfeld may appear to be teaching strategy rather than subject matter, he was, more fundamentally, building with his class a mathematical belief system around his own and the class's intuitive responses to the problem.

FIGURE 1

The Magic Square Problem

Can you place the digits 1,2,3,4,5,6,7,8,9 in the box below, so that the sum of the digits along each row, each column, and each diagonal is the same? The completed box is called a magic square.



Note: From Schoenfeld, in press.

As an indication that Schoenfeld's class was working in the culture of mathematics, not in the culture of schooling, he did not have the students stop at what, in culture of school practice, would mark the end: an answer.

Are we done? In most mathematics classes the answer is "yes. " Early in the semester, my students all say "yes," expecting me to go on to another problem. My answer, however, is a resounding "no." In most classes, so-called "problems" are exercises; you are done when you've shown that you've mastered the relevant technique by getting the answer. (Schoenfeld, in press)

His class's goal, by contrast, was to understand the mathematical nature of magic square, and it was in part by doing this that the belief system was exemplified. The class explored other possible magic squares and discovered general principles (e.g., an algebraic form for describing the squares). It also led to some further generalizable mathematical strategies that are less commonly seen in classroom

practice, such as working forwards from an initial solution; using systematic generating procedures; having more than one way to solve a problem. Schoenfeld is consistently careful to emphasize that all such strategies are illustrated in action, developed by the class, not declared by the teacher. In his classes, the belief system is instilled in the only way it can be, through practice in which the students actively take part.

Lampert's teaching of multiplication. Lampert (1986) also involves her students in mathematical exploration, which she tries to make continuous with their everyday knowledge. She has devised methods for teaching mathematics to fourth grade students that lead from students' implicit understanding of the world beyond the classroom, through activity and social construction in the culture, to the sort of robust learning that direct teaching of algorithms usually fails to achieve.

She starts teaching multiplication, for example, in the context of coin problems, because in the community of fourth grade students, there is usually a strong, implicit, shared understanding of coins. Next, the students create stories for multiplication problems, drawing on their implicit knowledge to delineate different examples of multiplication. Then, Lampert helps them toward the abstract algorithm that everyone learns for multidigit multiplication, in the context of the coin problems and stories the community has created. Thus, the method presents the algorithm as one more useful strategy to help them resolve community problems.

The first phase of teaching starts with simple coin problems, such as "using only nickels and pennies, make 82 cents." With such problems, Lampert helps her students explore their implicit knowledge. Then, in the second phase, the students create stories for multiplication problems (see Figure 2). They perform a series of decompositions and discover that there is no one, magically "right" decomposition decreed by authority, just more and less useful decompositions whose use is judged in the context of the problem to be solved and the interests of the problem solvers.

FIGURE 2

Story Problems for Teaching Multiplication

T: Can anyone give me a story that could go with this multiplication...
 12×4 ?

S1: There were 12 jars, and each had 4 butterflies in it.

T: And if I did this multiplication and found the answer, what would I know about those jars and butterflies?

S1: You'd know you had that many butterflies altogether.

T: Okay, here are the jars. [Draws a picture to represent the jars of butterflies --- see diagram.] The stars in them will stand for butterflies. Now, it will be easier for us to count how many butterflies there are altogether, if we think of the jars in groups. And as usual, the mathematician's favorite number for thinking about groups is?

S2: 10

T: Each of these 10 jars has 4 butterflies in it. [Draws a loop around the 10 jars.]...

T: Suppose I erase my circle and go back to looking at the 12 jars again altogether. Is there any other way I could group them to make it easier for us to count all the butterflies?

S6: You could do 6 and 6.

T: Now, how many do I have in this group?

S7: 24

T: How did you figure that out?

S7: 8 and 8 and 8. [He puts the 6 jars together into 3 pairs, intuitively finding a grouping that made the figuring easier for him.]

T: That's 3×8 . It's also 6×4 . Now, how many are in this group?

S6: 24. It's the same. They both have 6 jars.

T: And now how many are there altogether?

S8: 24 and 24 is 48.

T: Do we get the same number of butterflies as before? Why?

S8: Yeah, because we have the same number of jars and they still have 4 butterflies in each.

Note: From Lampert, 1986.

The third phase of instruction gradually introduces students to the standard algorithm, now that such an algorithm has a meaning and a purpose in their community. The students' procedure parallels the story problems they had created. Eventually they find ways to shorten the process, and they usually arrive at the standard algorithm, justifying their findings with the stories they created earlier.

Through this method, students develop a composite understanding of four different kinds of mathematical knowledge: (a) intuitive knowledge, the kind of short cuts people invent when doing multiplication problems in authentic settings; (b) computational knowledge, the basic algorithms that are usually taught; (c) concrete knowledge, the kind of concrete models of the algorithm associated with the stories the students created; and (d) principled knowledge, the principles such as associativity and commutativity that underlie the algorithmic manipulations of numbers. Lampert tries to inculcate an inseparable understanding of these kinds of knowledge and the connections between them, and thus to bridge the huge gap that emerges from much conventional teaching between conceptual knowledge and problem solving activity --- between, as we characterized them at the beginning, knowing and doing.

This approach fosters procedures that are characteristic of cognitive apprenticeship:

- By beginning with a task embedded in a familiar activity, it shows the students the legitimacy of their implicit knowledge and its availability as scaffolding in apparently unfamiliar tasks.
- By pointing to different decompositions, it stresses that heuristics are not absolute, but assessed with respect to a particular task --- and that even algorithms can be assessed in this way.
- By allowing students to generate their own solution paths, it helps make them conscious, creative members of the culture of problem-solving mathematicians. And, in enculturating through this activity, they acquire some of the culture's tools --- a shared vocabulary and the means to discuss, reflect upon, evaluate, and validate community procedures in a collaborative process.

Schoenfeld's approach differs principally in its strong emphasis on exposing students to the authentic ways of thinking of a culture and its conceptual viewpoint, as much as to its subject matter.

Figure 3 shows how, in the terms of cognitive apprenticeship, we can represent the progress of the students from embedded activity to general principles of the culture. In this sequence, apprenticeship and coaching in a domain begin by providing modeling in situ and scaffolding for students to get started in an authentic activity. As the students gain more self-confidence and control, they move into a more autonomous phase of collaborative learning, where they begin to participate consciously in the culture. The social network within the

culture helps them develop its language and the belief systems and promotes the process of enculturation. Collaboration also leads to articulation of strategies, which can then be discussed and reflected on. This, in turn, fosters generalizing, grounded in the students' situated understanding. From here, students can use their fledgling conceptual knowledge in activity, seeing that activity in a new light, which in turn leads to the further development of the conceptual knowledge.

In language learning, for instance, the original frail understanding of a word is developed and extended through subsequent use and social negotiation, though each use is obviously situated. Miller and Gildea (1978) describe two stages of this process. The first, in which people learn the word and assign it a semantic category (e.g., the word olive is first assigned to the general category of color words), is quickly done. The second, in which distinctions within this semantic category (e.g., between olive and other colors) are explored as the word occurs again and again, is a far more gradual process, which "may never be completely finished" (p. 95). This second phase of word learning corresponds to the development through activity of all conceptual knowledge. The threadbare concepts that initially develop out of activity are gradually given texture as they are deployed in different situations.

Apprenticeship and Cognition

The development of concepts out of and through continuing authentic activity is the approach of cognitive apprenticeship --- a term closely allied to our image of knowledge as a tool. Cognitive apprenticeship supports learning in a domain by enabling students to acquire, develop, and use cognitive tools in authentic domain activity. Similarly, craft apprenticeship enables apprentices to acquire and develop the tools and skills of their craft through authentic work at and membership in their trade. Through this process, apprentices enter the culture of practice. So the term apprenticeship helps to emphasize the centrality of activity in learning and knowledge and highlights the inherently context-dependent, situated, and enculturating nature of learning. And apprenticeship also suggests the paradigm of situated modeling, coaching, and fading (Collins, Brown, & Newman, in press), whereby teachers or coaches promote learning, first by making explicit their tacit knowledge or by modeling their strategies for students in authentic activity. Then, teachers and colleagues support students' attempts at doing the task. And finally they empower the students to continue independently. The progressive process of learning and

enculturation perhaps argues that Increasingly Complex Microworlds (see Burton, Brown, & Fischer, 1984) can be replaced by increasing complex enculturating environments.

Cognitive emphasizes that apprenticeship techniques actually reach well beyond the physical skills usually associated with apprenticeship to the kinds of cognitive skills more normally associated with conventional schooling. This extension is not as incompatible with traditional apprenticeship as it may at first seem. The physical skills usually associated with apprenticeship embody important cognitive skills, if our argument for the inseparability of knowing and doing is correct. Certainly many professions with generally acknowledged cognitive content, such as law, medicine, architecture, and business, have nonetheless traditionally been learned through apprenticeship.

Moreover, advanced graduate students in the humanities, the social sciences, and the physical sciences acquire their extremely refined research skills through the apprenticeships they serve with senior researchers. It is then that they, like all apprentices, must recognize and resolve the ill-defined problems that issue out of authentic activity, in contrast to the well-defined exercises that are typically given to them in text books and on exams throughout their earlier schooling. It is at this stage, in short, that students no longer behave as students, but as practitioners, and develop their conceptual understanding through social interaction and collaboration in the culture of the domain, not of the school.

In essence, cognitive apprenticeship attempts to promote learning within the nexus of activity, tool, and culture that we have described. Learning, both outside and inside school, advances through collaborative social interaction and the social construction of knowledge. Resnick has pointed out (1988) that throughout most of their lives people learn and work collaboratively, not individually, as they are asked to do in many schools. Lampert's and Schoenfeld's work, Scardamalia, Bereiter, and Steinbach's teaching of writing (1984), and Palincsar and Brown's (1984) work with reciprocal teaching of reading all employ some form of social interaction, social construction of knowledge, and collaboration.

Within a culture, ideas are exchanged and modified and belief systems developed and appropriated through conversation and narratives, so these must be promoted, not inhibited. Though they are often anathema to traditional schooling, they are an essential component of social interaction and, thus, of learning. They provide access to much

of the distributed knowledge and elaborate support of the social matrix (Orr, 1987). So learning environments must allow narratives to circulate and "war stories" to be added to the collective wisdom of the community.

The role of narratives and conversations is perhaps more complex than might first appear. An intriguing role in learning is played by "legitimate peripheral participation," where people who are not taking part directly in a particular activity learn a great deal from their legitimate position on the periphery (Lave & Wenger, in preparation). It is a mistake to think that important discourse in learning is always direct and declarative. This peripheral participation is particularly important for people entering the culture. They need to observe how practitioners at various levels behave and talk to get a sense of how expertise is manifest in conversation and other activities.

Cognitive apprenticeship and collaborative learning. If, as we propose, learning is a process of enculturating that is supported in part through social interaction and the circulation of narrative, groups of practitioners are particularly important, for it is only within groups that social interaction and conversation can take place. Salient features of group learning include:

- Collective problem solving. Groups are not just a convenient way to accumulate the individual knowledge of their members. They give rise synergistically to insights and solutions that would not come about without them (Schoenfeld, in preparation).
- Displaying multiple roles. Successful execution of most individual tasks requires students to understand the many different roles needed for carrying out any cognitive task. Getting one person to be able to play all the roles entailed by authentic activity and to reflect productively upon his or her performance is one of the monumental tasks of education. The group, however, permits different roles to be displayed and engenders reflective narratives and discussions about the aptness of those roles.
- Confronting ineffective strategies and misconceptions. We know from an extensive literature (diSessa, 1982, 1983, 1986; McCloskey, Caramazza, & Green, 1980; White, 1983) that students have many misconceptions about qualitative phenomena in physics. Teachers rarely have the opportunity to hear enough of what students think to recognize when the information that is offered back by students is only a surface retelling for school purposes (the handing back of an

uncomprehended tool, as we described it at the beginning) that may mask deep misconceptions about the physical world and problem solving strategies. Groups however, can be efficient in drawing out, confronting and discussing both misconceptions and ineffective strategies.

- Providing collaborative work skills. Students who are taught individually rather than collaboratively can fail to develop skills needed for collaborative work. In the collaborative conditions of the workplace, knowing how to learn and work collaboratively is increasingly important. If people are going to learn and work in conjunction with others, they must be given the situated opportunity to develop those skills.

In looking at Schoenfeld's and Lampert's teaching, in noting what we believe are important features of their methods, and in stressing social interaction and collaborative learning, we are trying to show how teaching through a form of apprenticeship can accommodate the new view of knowledge and learning we have been outlining. The increasing role of the teacher as a master to apprentices, and the teachers' use of authentic domain activity as a major part of teaching will perhaps, once and for all, dismiss George Bernard Shaw's scurrilous criticism of teachers, "He who can, does. He who cannot, teaches." His comment may then be replaced with Alexander Pope's hopeful "Let such teach others who themselves excel."

Conclusion ---Toward an Epistemology of Situated Cognition

Much research investigating situated features of cognition remains to be done. It is, however, already possible to begin serious reappraisal of the assumptions about learning that underlie current classroom practice (see, for example Resnick, 1988; Shanker, 1988).

One of the particularly difficult challenges for research, (which exceptional teachers may solve independently) is determining what should be made explicit in teaching and what should be left implicit. A common strategy in trying to overcome difficult pedagogic problems is to make as much as possible explicit. Thus, we have ended up with wholly inappropriate methods of teaching. Whatever the domain, explication often lifts implicit and possibly even nonconceptual constraints (Cussins, 1988) out of the embedding world and tries to make them explicit or conceptual. These now take a place in our ontology and become something more to learn about rather than simply something useful in learning. But indexical representations gain their efficiency by leaving much of the context underrepresented or

implicit. Future work into situated cognition, from which educational practices will benefit, must, among other things, try to frame a convincing account of the relationship between explicit knowledge and implicit understanding.

We have described here only a fragment of an agenda for a fully developed theory of situated cognition. There remains major theoretical work to shift the traditional focus of education. For centuries, the epistemology that has guided educational practice has concentrated primarily on conceptual representation and made its relation to objects in the world problematic by assuming that, cognitively, representation is prior to all else. A theory of situated cognition suggests that activity and perception are importantly and epistemologically prior --- at a nonconceptual level --- to conceptualization and that it is on them that more attention needs to be focused. An epistemology that begins with activity and perception, which are first and foremost embedded in the world, may simply bypass the classical problem of reference --- of mediating conceptual representations.

In conclusion, the unheralded importance of activity and enculturation to learning suggests that much common educational practice is the victim of an inadequate epistemology. A new epistemology might hold the key to a dramatic improvement in learning and a completely new perspective on education.

¹All work in this area is to a greater or lesser degree, built upon research of activity theorists such as Vygotsky, Leontiev, and others. For examples of recent work, see for instance, Rogoff and Lave, 1984; Scribner, 1984; Hutchins, in press; Engestrom, 1987; Lave and Wenger, in preparation; and in particular Lave, 1977, 1988a, 1988b, 1988c, in preparation. Anyone familiar with Jean Lave's work on learning, apprenticeship, and everyday cognition will realize at once that we are deeply indebted to her groundbreaking work,

²The dictionary definitions that the students used in writing these sentences are as follows: Correlate --- be related one to the other; meticulous --- very careful; stimulate --- stir up. They were given these definitions with little or no contextual help, so it would be unfair to regard the students as foolish for using the words as they did.

³In the linguistics literature, the term deixis is often used instead of indexicality. See, for example, J. Fillmore, Santa Cruz Lectures.

⁴This image is, of course, not original. For the way it is developed here, we are particularly indebted to Richard Burton, who explored it during a symposium on education organized by the Secretary of Education of Kentucky and to D. N. Perkins' book *Knowledge as Design* (1986).

⁵The JFP must, of course, have access to a culture and become what Lave and Wenger (in preparation) call a "legitimate peripheral participant." And, of course, an apprentice usually has to do a great deal of work. We are not trying to suggest that anything magical occurs in the process of enculturation. (Medical interns testify to how hard it can be.) But the process, we stress, is not qualitatively different from what people do all the time in adopting the behavior and belief systems of their peers.

⁶To get some sense of how foreign this is to school tasks, it might be useful to imagine the impropriety of a student's being given this problem and asked "Does the dieter have a measuring cup, cutting board, and knife at hand?" Though word problems are meant to ground theory in activity, the things that structure activity are denied to the problem solvers. Textbooks ask students to solve supposedly "real-life" questions about people who do very unreal things, such as driving at constant speeds in straight lines or filling leaking troughs with leaking buckets. Students are usually not allowed to indulge in real-life speculation. Their everyday inventiveness is constrained by prescribing and proscribing ways in which the solution must be found. The ubiquitous Mr. Smith might, after all, wisely repair the hole in his bucket or fill the trough with a hose. Sitting down and calculating how many journeys it will take with a leaking bucket is probably the very last thing he would do. (See also Lave, 1988c.)

Editor's Note: In an effort to encourage informed discussion and debate on the themes of this article, the new ER will publish a set of commentaries in the May 1989 issue.

Acknowledgment: Many of the ideas in this paper emerged from group discussions at the Institute for Research on Learning. We are

especially grateful to James Greeno, Jean Lave, Susan Newman, Roy Pea, and John Rheinfrank, who read earlier drafts and commented on them with great care. We are also grateful to Richard Burton, William Clancey, and Alan Schoenfeld for helpful and insightful contributions. More generally, we would like to acknowledge the influence of Brian Cantwell Smith's pioneering research into a theory of computation and semantics built on notions of situatedness, embeddedness, and embodiedness; of Susan Stucky's important new idealization of mind in terms of "radical" efficiency rather than rationality; and also of the work on indexicality of Philip Agre and David Chapman.

This research was supported in part by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research under Contract NO. NOO014-C-85-0026. Contract Authority Identification Number, NR 667-540.

An extended version of this article will appear as IRL report No. 88-0008 (available from the Institute for Research on Learning) and as BBN Research Report 6886 (available from Bolt Beranek & Newman Inc.).

Agre P. (1985). Routines. MIT AI Memo.

Barwise, K. J., & Perry, J. (1983). Situations and attitudes. Cambridge, MA: MIT Press.

Burton, R., Brown, J. S., & Fischer, G. (1984). Skiing as a model of instruction. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its development in social context*. (pp. 139-150) Cambridge, MA: Harvard University Press.

Collins, A., Brown, J. S., & Newman, S. E. (in press). Cognitive apprenticeship: Teaching the craft of reading, writing and mathematics. In L. B. Resnick (Ed.), *Knowing learning and instructions: Essays in honor of Robert Glaser*. Hillsdale, NJ: Erlbaum.

Cussins, A. (1988). The connectionist construction of concepts. (SSL Research Report). Palo Alto, CA: Xerox Palo Alto Research Center.

deKleer, J., & Brown, J. S. (1984). A qualitative physics based on confluences. *Artificial Intelligence Journal* 24, 1-3.

diSessa, A. (1982). Unlearning Aristotelian physics: A study of knowledge-based learning. *Cognitive Science* 6, 37-75.

diSessa, A. (1983). Phenomenology and the evolution of intuition. In D. Gentner and A. Stevens (Eds.), *Mental models*. (pp. 15-33) Hillsdale, NJ: Erlbaum.

diSessa, A. (1986). Knowledge in pieces. In G. Forman & P. Pufall (Eds.), *Constructivism in the computer age*. Hillsdale, NJ: Erlbaum.

Engestrom, Y. (1987). *Learning by expanding*. Helsinki: Orienta-Konsultit Oy.

Fillmore, J. (1974). *Santa Cruz lectures on deixis*. Bloomington, IN: Indiana University Linguistics Club.

Geertz, C. (1983). *Local knowledge*. New York: Basic Books.

Hutchins, E. (in press). Learning to navigate. In S. Chalkin and J. Lave (Eds.), *Situated Learning*.

Lampert, M. (1986). Knowing, doing, and teaching multiplication. *Cognition and Instruction*, 3, 305-342.

Lave, J. (1977). Tailor-made experiments and evaluating the intellectual consequences of apprenticeship training. *The Quarterly Newsletter of the Institute for Comparative Human Development* 1 1-3.

Lave, J. (1988a). The culture of acquisition and the practice of understanding. (IRL report 88-00087). Palo Alto, CA: Institute for Research on Learning.

Lave, J. (1988b). *Cognition in practice*. Boston, MA: Cambridge.

Lave, J. (1988c). Word problems: A microcosm of theories of learning. Paper presented at AERA annual conference, New Orleans, LA.

Lave, J. (in preparation). Tailored learning: Education and everyday practice among craftsmen in West Africa.

Lave, J., & Wenger, E. (in preparation). *Situated learning: Legitimate peripheral participation*.

McCloskey, M., Caramazza, A., & Green, B. (1980). Curvilinear motion in the absence of external forces: Naive beliefs about the motion of objects. *Science*, 210, 1139-1141.

Miller, G. A., Gildea, P. M. (1987). How children learn words. *Scientific American*, 257, (3), 94-99.

Nunberg, G. (1978). *The pragmatics of reference*. Bloomington, IN: Indiana University Linguistics Club.

Orr, J. (1987). *Talking about machines*. (SSL Report). Palo Alto, CA: Xerox Palo Alto Research Center.

Palincsar, A. S. (1986). Metacognitive strategy instruction. *Exceptional Children*, 53, 118-124.

Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and monitoring activities. *Cognition and Instruction*, 1, 117-175.

Pea, R. D. (1988). Distributed intelligence in learning and reasoning processes. Paper presented at the meeting of the Cognitive Science Society, Montreal.

Perkins, D. N. (1986). *Knowledge as design*. Hillsdale, NJ: Erlbaum.

Perry, J. (1979). The problem of the essential indexical. *Nous*, 13, 3-21.

Resnick, L. (1988). Learning in school and out. *Educational Researcher*, 16(9), 13-20.

Rogoff, B., & Lave, J. (Eds.). (1984). *Everyday cognition: Its development in social context*. Cambridge, MA: Harvard University Press.

Rubin, A. (1980). A theoretical taxonomy of the differences between oral and written language. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 411-438). Hillsdale, NJ: Erlbaum.

Scardamalia, M., Bereiter, C., & Steinbach, R. (1984). Teachability of reflective processes in written composition. *Cognitive Science*, 8, 173-190.

Schoenfeld, A. H. (1985). *Mathematical problem solving*. Orlando, FL: Academic Press.

Schoenfeld, A. H. (in press). On mathematics as sense-making: An informal attack on the unfortunate divorce of formal and informal mathematics. In D. N. Perkins, J. Segal, & J. Voss (Eds.), *Informal reasoning and education*. Hillsdale, NJ: Erlbaum.

Schoenfeld, A. H. (in preparation). *Ideas in the air*. (IRL report 88-0011). Palo Alto, CA: Institute for Research on Learning.

Scribner, S. (1984). Studying working intelligence. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its development in social context* (pp. 9-40). Cambridge, MA: Harvard University Press.

Shanker, A. (1988). Exploring the missing connection. *New York Times*, E7.

Suchman, L. (1987). *Plans and situated actions*. New York: Cambridge University Press.

White, B. (1983). Sources of difficulty in understanding Newtonian dynamics. *Cognitive Science*, 7, 41-65.

Whitehead, A. N. (1929). *The aims of education*. Cambridge: Cambridge University Press.

JOHN SEELY BROWN and PAUL DUGUID are at the Institute for Research on Learning, 2550 Hanover Street, Palo Alto, California 94304. ALLAN COLLINS is at BBN Inc, 10 Moulton Street, Cambridge, Massachusetts 02238.

Permission to reprint granted by American Educational Research Association.

Return to
[Museum Education](#)
[Institute for Inquiry Home](#)

© [Exploratorium](#), 3601 Lyon St., San Francisco, CA 94123