Fourth printing, 1997
First MIT Press Paperback edition, 1994

Set in Bembo by DEKR Corporation.

Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Bruer, John T., 1949–
Schools for thought : a science of learning in the classroom /
John T. Bruer.
p. cm.
"A Bradford book."
Includes bibliographical references and index.
ISBN 0-262-02352-0 (HB), 0-262-52196-2 (PB)
1. Learning, Psychology of. 2. Cognitive science. 3. Cognition
in children. 4. Teaching. 5. Education—United States—Aims and
objectives. I. Title.
LB1060.B78 1993
370.15'23—dc20 92-29484
CIP

Schools for Thought

A Science of Learning in the Classroom

John T. Bruer

A Bradford Book

The MIT Press
Cambridge, Massachusetts
London, England

3

Intelligent Novices: Knowing How to Learn

Imagine that a small, peaceful country is being threatened by a large, belligerent neighbor. The small country is unprepared historically, temperamentally, and militarily to defend itself; however, it has among its citizens the world's reigning chess champion. The prime minister decides that his country's only chance is to outwit its aggressive neighbor. Reasoning that the chess champion is a formidable strategic thinker and a deft tactician—a highly intelligent, highly skilled problem solver—the prime minister asks him to assume responsibility for defending the country. Can the chess champion save his country from invasion?

This scenario is not a plot from a Franz Lehar operetta, but a thought experiment devised by David Perkins and Gavriel Salomon (1989). As they point out, our predictions about the chess champion's performance as national security chief depend on what we believe intelligence and expertise are. If the goal of education is to develop our children into intelligent subject-matter experts, our predictions about the chess champion, based on what we believe about intelligence and expertise, have implications for what we should do in our schools.

Since the mid 1950s cognitive science has contributed to the formulation and evolution of theories of intelligence, and so to our understanding of what causes skilled cognitive performance and what should be taught in schools. In this chapter, we will review how our understanding of intelligence and expertise has evolved over the past two decades and see how these theories have influenced educational policy and practice.

Four theories will figure in this story.

The oldest theory maintains that a student builds up his or her intellect by mastering formal disciplines, such as Latin, Greek, logic, and maybe chess. These subjects build minds as barbells build muscles. On this theory the chess champion might succeed in the national security field. If this theory is correct, these formal disciplines should figure centrally in school instruction.

In the early years of the cognitive revolution, it appeared that general skills and reasoning abilities might be at the heart of human intelligence and skilled performance. If this is so, again the chess champion might succeed, and schools should teach these general thinking and problem-solving skills—maybe even in separate critical-thinking and study-skills classes.

By the mid 1970s, cognitive research suggested that general domain-independent skills couldn't adequately account for human expertise. Researchers then began to think that the key to intelligence in a domain was extensive experience with and knowledge about that domain. Expertise was domain specific. This suggested that the chess expert was doomed to failure, and that schools should teach the knowledge, skills, and representations needed to solve problems within specific domains.

In the early 1980s researchers turned their attention to other apparent features of expert performance. They noticed that there were intelligent novices—people who learned new fields and solved novel problems more expertly than most, regardless of how much domain-specific knowledge they possessed. Intelligent novices controlled and monitored their thought processes and made use of general, domain-independent strategies and skills where appropriate. This suggested that there was more to expert performance than just domain-specific knowledge and skills.

Perkins and Salomon call this latest theory or view the "new synthesis," because it incorporates what was correct about the earlier views, while pointing out that none of the earlier theories alone provides an adequate basis for effective educational practice. According to the new synthesis, we should combine the learning of domain-specific subject matter with the learning of general thinking skills, while also making sure that children learn to monitor and control their thinking and learning.

The new synthesis introduces an important new idea into discussions about educational reform. The first three theories of intelligence emphasize *what* we should teach in our schools—formal disciplines, general thinking and learning skills, or domain-specific knowledge and skills. The new synthesis, as we shall see, implies that we should be as concerned with *how* we teach as we traditionally have been concerned with what we teach. The most recent research shows that if we can apply the new synthesis in the classroom, we should be able to teach school subjects as high-order cognitive skills and help children become intelligent novices and expert learners.

Transfer

What connects the chess champion, theories of intelligence, and schooling is a phenomenon psychologists and educators call *transfer*. We generally believe that learning a certain skill or subject area can help us learn a related one. If we first learn tennis, we should be able to learn squash more easily. If we learn Spanish as a second language, we should be able to learn Italian as a third language more easily. Knowledge from the first skill or domain should transfer to the second, so there is less to learn. Notice, though, that in neither of these examples are we simply applying previously learned knowledge. Squash isn't tennis, and Italian isn't Spanish. In these situations, we are using old skills or knowledge in novel situations where we also have to learn new things. One cognitive scientist describes it this way: "Transfer means applying old knowledge in a setting sufficiently novel that it also requires learning new knowledge." (Larkin 1989, p. 283)

If this description is correct, we should be able to tell when transfer occurs. If knowledge transfers from task A to task B, then people who have learned A should be able to learn B more rapidly than people who did not first learn A. A tennis player should be able to learn squash more rapidly than a person with no prior experience at racquet sports.

Transfer is central to designing and developing effective instruction. Problems of transfer pervade schooling. Teachers want to teach lessons so that students can transfer what they have learned during class instruction to solve new problems at the end of a chapter. We want that learning to transfer to the unit, semester, or standardized test. Most important, we want school learning to transfer to real-world problem solving at home and on the job. If this is our goal, what and how should we teach?

If we want to teach so as to promote transfer of knowledge, we have to answer a prior question: What kinds of knowledge and skills, if any, transfer between tasks? What, if anything, might transfer from chess to national security? Theories of intelligence and expertise suggest answers to this question. Theories differ in their claims about what, whether, and when knowledge transfers from one task or knowledge domain to another. Of the theories outlined above, the first says that general mental strength transfers, the second that general skills and strategies transfer, and the third that expertise is domain specific (so that we might find some transfer within a domain, but little or none across domains). The new synthesis suggests that transfer can occur within and across domains, but only if we teach students appropriately.

Formal Disciplines and Mental Fitness

Our oldest theory of expertise and intelligence goes back to the classical Greeks, who believed that mastering formal disciplines, such as arithmetic and geometry, would improve general intelligence and reasoning ability. By the eighteenth century, scholars had added grammar, mnemonics, Greek, and Latin to the list of disciplines that build mental fitness. The theory was that these difficult formal disciplines would build general mental strength, just as rigorous physical exercise builds physical strength. On this theory, if we believe that chess is a formal discipline on a par with logic and geometry, we might favor the chess master's chances.

Edward Thorndike's careful studies of learning and of what knowledge transfers from one subject to another were among scientific psychology's first contributions to education. (See Thorndike and Woodworth 1901.) At the turn of the twentieth century, when Thorndike did his work, the prevailing view, derived from the ancient Greeks, was that learning formal disciplines improved general mental functioning. Thorndike, however, noted that no one had presented scientific evidence to support this view. Thorndike reasoned that if learning Latin strengthens general mental functioning, then students who had learned Latin should be able to learn other subjects more quickly. He found no evidence of this. Having learned one formal discipline did not result in more efficient learning in other domains. Mental "strength" in one domain didn't transfer to mental strength in others. Thorndike's results contributed to the demise of this ancient theory of intelligence and to a decline in the teaching of formal disciplines as mental calisthenics.

However, in some experiments where two subject domains shared surface similarity, Thorndike did observe faster learning in the second domain. He proposed a theory of "identical elements" to explain this. Thorndike suggested that where two domains share common *elements of knowledge*—not formal rigor—a person who has learned one of them might be able to learn the second more quickly. But because psychologists at the turn of the century had no precise way to describe and identify "elements," Thorndike couldn't test his theory rigorously. The methods he needed were those that cognitive psychologists developed more than 50 years later.

Elements, Productions, and Transfer

Once psychologists accepted the assumption that our minds process symbols, and once they realized they could study minds as information-processing devices, it became possible to test theories such as Thorndike's. Psychologists, using the framework of computational theory, could describe "elements" as symbol structures and devise problem-solving simulations and experiments to see which symbol structures two disciplines might share.

Production systems are among the things that allow psychologists to test modern versions of Thorndike's theory. If minds are devices that execute production systems, and if (as on the balance scale) learning occurs when we add new productions to long-term memory, then we might be able to formulate and test Thorndike's claim. We can think of each individual production rule as a piece of knowledge needed for a task; we can think of it as one of Thorndike's elements. If so, the transfer of learning from one task to another should be directly related to the number of productions the tasks share.

M. K. Singley and John R. Anderson (1985) performed an elegant study to test this hypothesis. They studied the way in which secretarial students learned to use three different text editors or word processors. Two of the editors, ED and EDT, were line editors that allowed the user to edit one line of text at a time. EMACS was a screen editor, more like a standard word processor, that allowed the user to edit a document one screen at a time.

As is typical of cognitive scientists, Singley and Anderson first did a careful task analysis of the three editors. The two screen editors used different names for the editing commands and differed in how the user located a line to edit. Once past these superficial differences, however, they found that the production systems used to simulate expert performance on the two line editors were nearly identical. However, the production system that simulated expert performance on EMACS, the screen editor, was almost entirely different from those for the line editors. Thus, there was considerable production overlap between the line editors and almost none between the line editors and the screen editor.

How did this affect learning? Students who learned either line editor first took as long to learn the screen editor as students who started out on the screen editor. Skill on the line editors didn't transfer to the screen editor. In the case of the two line editors, students who learned one learned the second much more quickly. There was considerable transfer between the two line editors. Anderson (1985, p. 241) estimates that learning one of the line editors eliminated up to 90 percent of the work normally needed to learn the second. Singley and Anderson concluded that the amount of production overlap between two skills predicts the amount of transfer between the skills.

Relying on computational theory, production systems, and task analysis allowed cognitive science to make precise scientific sense of Thorndike's hypothesis. The information-processing approach can give us fine-grained representations—in this case, productions—of Thorndike's common elements. Cognitive research gives us methods for stating and testing claims about the transfer of knowledge between tasks.

General Methods and Intelligent Behavior

Cognitive scientists started applying computational insights to issues of expertise, intelligence, and transfer in the late 1950s. To understand their initial approach, recall the model of problem solving presented in chapter 2. As we saw there, problems have initial states and goal states. The solver chooses operators that create a chain of knowledge states linking the initial state to the goal. The operators, themselves composed of basic information processes, combine to form procedures or programs that guide problem-solving behavior.

How do we choose operators when solving problems? Cognitive scientists say that we use *methods* or *strategies* to choose them. Imagine we are playing chess. One method I might use to choose a move, or an operator, is to pick a piece at random or move the first piece I happen to touch. At the other extreme, I might choose my moves by following an opening I have studied in a chess book. An intermediate option would be to use a method based on general chess principles: I might choose my moves so as to control the center, defend my pieces, and attack yours. The same spectrum of methods or strategies is available for balance-scale problems. I might generate a prediction by randomly choosing among left, right, and balance. At the other extreme, I might use Siegler's rule IV, having studied it in a book.

These methods differ in ways that are more interesting than how often I would win the game or make a correct balance-scale prediction. First, they differ in how widely applicable they are. Second, they differ in what I have to know to use them. The method of choosing an operator at random works for any problem: I don't have to know anything about chess or balance scales to use random choice. In contrast, following a line from a chess book or using rule IV works only for chess or balance-scale problems. Furthermore, I have to know a lot about chess or balance scales to use a book or Siegler's rule. Cognitive psychologists call methods that are widely applicable and that require little or no specific knowledge *weak methods*. They call methods that are situation-specific and domain-specific *strong methods*. Random choice is a weak method; rule IV is a strong one. Psychologists see all strategies, procedures, and skills as falling somewhere on the continuum between weak and strong methods.

In the early days of computer and cognitive science, there were divergent views about how to make computers or people more intelligent. Some thought the key to understanding intelligent behavior, for both machines and humans, lay in developing and understanding weak methods that were applicable across many problem domains. Others thought the better scientific bet was to study the knowledge needed in specific domains and find the specific strong methods that experts used.

Initial successes in artificial intelligence research suggested that weak methods were the way to go. Logic Theorist could prove logical theorems. A second-generation machine, called General Problem Solver, could solve problems in a variety of domains, including algebra, geometry, and chess (Ernst and Newell 1969). These programs used weak methods such as *hill climbing* and *means-end analysis.*

Hill climbing is a weak method that chooses intelligent next moves on a problem if the problem requires progress along a single dimension. If you were trying to find the top of a hill in the dark, you would keep taking steps that tended in an upward direction. When you couldn't take any more upward steps, you would stop, assuming you had reached the top of the hill. The children's game of helping a playmate find a hidden object by giving clues of "hotter" and "colder" as the playmate moves toward or away from the object is a hill-climbing game.

Means-end analysis, the method General Problem Solver used, is more complex. Hill climbing considers only one difference between the current and goal states—in the children's game, all that matters is distance from the hidden object. Means-end analysis identifies several differences between the current situation and the goal, then picks an action or an operator that will reduce one or more of those differences. If more than one action or operator could be used, means-end analysis chooses the one whose conditions of applicability best match the current situation. Sometimes, after choosing the action best suited to the situation, one still can't execute the action because the conditions aren't right. In this case, means-end analysis establishes a subgoal to create conditions that permit the chosen action.

The wide applicability of means-end analysis is suggested by the example Newell and Simon (1972, p. 416) used:

> I want to take my son to nursery school. What is the difference between what I have and what I want? One of distance. What changes distance? My automobile. My automobile won't work. What is needed to make it work? A new battery. What has new batteries? An auto repair shop. I want the shop to put in a new battery; but the shop doesn't know I need one. What is the difficulty? One of communication. What allows communication? A telephone . . . and so on.

The solver here looks at where he is and where he wants to be, then works back and forth between the ends and the means to achieve those ends until he has a set of actions, or operators, that achieve the goal. His finding a telephone starts a chain of events that results in his son's arriving at nursery school. Students often use means-end analysis to solve school math and science problems.

In the late 1960s and the early 1970s, programs such as General Problem Solver suggested that general skills might be fundamental to human expertise and intelligence. If we could identify and teach such general skills, maybe we could improve human problem-solving performance both in and out of school. If general methods, skills, and strategies are the basis of the chess champion's expertise, maybe he can succeed at solving diplomatic problems.

Experts' Domain-Specific Knowledge

Research to extend the initial insight about the contribution of weak methods to human intelligence met with frustration. The research community soon realized there was a serious limitation on the early successes. The early AI programs, including General Problem Solver, simulated intelligent performance on games or formal, logical problems (e.g., solving number puzzles or proving logical theorems)—but playing games and solving formal problems demand little factual knowledge about the world. To succeed at these tasks requires knowledge of little more than the rules of the game or the rules of the formal system. However, problem solving in other domains, such as physics or medical diagnosis, requires considerably more factual knowledge. Just think about what you have to know to play tic-tac-toe versus what you have to know to solve a physics problem. When cognitive scientists started to explore some of these other more knowledge-rich areas, physics and medicine among them, they found that weak methods didn't work so well. In knowledge-rich domains, they found that strong methods tailor-made to work on specific, well-organized knowledge bases worked better.

Weak methods might work in domains where there is little factual knowledge, but one can't generalize from these puzzle domains to a general theory of human intelligence and expertise. As Marvin Minsky and Seymour Papert (1974, p. 59) wrote, "It is by no means obvious that very smart people are that way directly because of the superior power of their general methods—as compared with average people." Maybe expert, intelligent behavior depends crucially on the knowledge people have, how they organize it, and the specific methods they learn or develop to process it.

Data from psychological experiments also undermined the primacy of weak, general methods for human expertise. One of the most influential experiments was William Chase and Herb Simon's (1973) study of novice and expert chess players, which followed on earlier work by A. D. De Groot (1965). Chase and Simon showed positions from *actual* chess games to subjects for 5 to 10 seconds and asked the subjects to reproduce the positions from memory. Each position contained 25 chess pieces. Expert players could accurately place 90 percent of the pieces, novices only 20 percent. Chase and Simon then had the subjects repeat the experiment, but this time the "positions" consisted of 25 pieces placed randomly on the board. These were generally not positions that would occur in an actual game. The experts were no better than the novices at reproducing the random positions: both experts and novices could place only five or six pieces correctly.

Other researchers replicated the Chase-Simon experiment in a variety of domains, using children, college students, and adults. The results were always the same: Experts had better memories for items in their area of expertise, but not for items in general. This shows, first, that mastering a mentally demanding game does not improve mental strength in general. The improved memory performance is domain specific. Chess isn't analogous to a barbell for the mind. Second, it shows that if memory *strategies* account for the expert's improved memory capacity, the strategies aren't general strategies or weak methods. Chess experts have better memories for genuine chess positions, but not for random patterns of chess pieces or for strings of words or digits. Thus, experts aren't using some general memory strategy that transfers from chess positions to random patterns of pieces or to digit strings.

From long experience at the game, chess experts have developed an extensive knowledge base of perceptual patterns, or chunks. Cognitive scientists estimate that chess experts learn about 50,000 chunks, and that it takes about 10 years to learn them. Chunking explains the difference between novice and expert performance. Reproducing chessboard configurations after 5 to 10 seconds of study is a working-memory task, because there is not enough time to code and store the information in long-term memory. When doing this task, novices see the chessboard in terms of individual pieces. They can store only the positions of five or six pieces in working memory—numbers close to what we found our own working memory spans to be on the basis of the experiment discussed in chapter 2. Experts see "chunks," or patterns, of several pieces. If each chunk contains four or five pieces and if the expert can hold five such chunks in working memory, then the expert can reproduce accurately the positions of 20 to 25 individual pieces. Chase and Simon even found that when experts reproduced the positions on the board, they did it in chunks. They rapidly placed four or five pieces, then paused before reproducing the next chunk.

Expertise, these studies suggest, depends on highly organized, domain-specific knowledge that can arise only after extensive experience and practice in the domain. Siegler's balance-scale study (chapter 2) is another example. Under normal conditions, it takes a child at least 17 years to become expert at balance-scale problems. More knowledge about and experience with the balance scale results in more sophisticated, expert-like performance. Chunking helps children develop more complex rules that contribute to their growing expertise on the balance-scale task.

Other studies of problem solving also argue against general strategies. Try to solve the two problems illustrated in figure 3.1. Philip Johnson-Laird (1983, p. 30) found an interesting difference in individuals' abilities to solve them. This is interesting because formally, or logically, they are the same problem. The same general strategy or formal rule solves both.

The correct answers are "E and 7" and "Manchester and car." Many people answer, incorrectly, that they have to turn only E, or else E and 4, in the first problem. You do have to turn E, because if that card has an odd number on the other side the rule is false. You don't have to turn 4, because even if that card had a consonant on the other side it doesn't matter; the rule doesn't say anything about what is on the other side of a consonant card. You have to turn 7, because a vowel on the other side of that card would make the rule false. The same problem-solving strategy works for the second problem, and for any "if-then" rule as logicians interpret such statements. (According to the laws of logic, an "if-then" statement is false only in the case where the "if" clause is true and the "then" clause is false; in every other case the statement is true.)

The two problems in figure 3.1 differ only in their subject matter. The first problem is an abstract one about letters and numbers, but the second one deals with a possible real-life situation. Johnson-Laird's subjects were much better at the second problem. Only 12 percent of them said they would turn over the "7" card to test the first rule, but over 60 percent said they would turn over the "car" card to test the second rule. Furthermore, he found that giving subjects experience with real-life if-then problems didn't improve their performance on more abstract versions. Apparently, most of

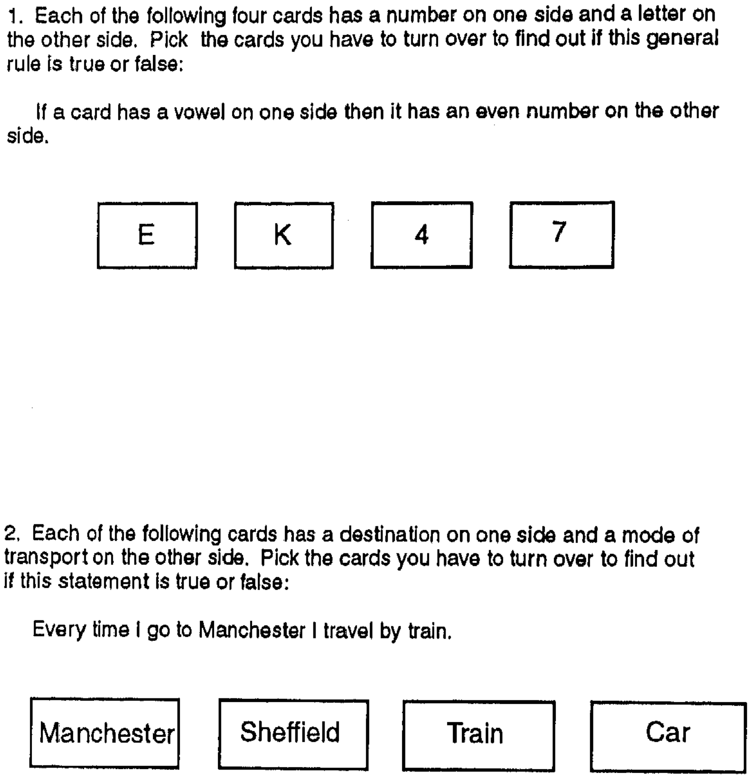

1. Each of the following four cards has a number on one side and a letter on the other side. Pick the cards you have to turn over to find out if this general rule is true or false:

If a card has a vowel on one side then it has an even number on the other side.

2. Each of the following cards has a destination on one side and a mode of transport on the other side. Pick the cards you have to turn over to find out if this statement is true or false:

Every time I go to Manchester I travel by train.

Figure 3.1
These two problems are logically the same; they differ only in their content. (From Johnson-Laird 1983.)

us don't use a general rule or strategy to solve these problems. If we did, we would use the same rule to solve all such problems. There would be no difference between performance on number-letter problems and on destination-transportation problems. In other experiments, researchers report that the ability to transfer a solution from one version of a problem to another occurs only if the experimenter *explicitly tells* subjects that the two problems are the same. For some reason, what the problem is about and how familiar we are with that content affects our problem solving. It seems as if domain-specific knowledge *does* contribute to expert performance.

It also appears that the ability to use general strategies at all depends on the subject's having a knowledge base on which the

strategies can work. When an experimenter asks subjects to memorize lists of words (e.g., "dog, gold, carrots, diamond, cat, peas"), subjects rarely repeat the words in that order. Usually, subjects say something like "dog, cat, carrots, peas, gold, diamond." To remember the words, subjects group them into meaningful categories—here animal, vegetable, and mineral. Psychologists call this often-unconscious strategy *clustering*. Clustering helps us remember things by exploiting the schema structure of long-term memory; we remember the words by associating them with the appropriate schemas.

When college students and young children were the subjects in such experiments, psychologists found that the college students recalled more words and did more clustering. Initially, psychologists attributed young children's poor performance to their inability to use the clustering strategy. Later it was discovered that if the word list included things young children know more about than college students, the results would be different (Lindberg 1980). If the experimenter used a 30-word list that included names of children's television celebrities, cartoon stars, and comic book characters, young children recalled more and used more clustering than the college students. Thus, there is an interaction between knowledge and strategy use—between facts and skills. Subjects are more likely to use a general memory strategy the more they know about a domain or a topic. Strategies can help us process knowledge, but first we have to have the knowledge to process.

While granting the possibility that strategies might play some role in problem solving, by the mid 1970s many cognitive scientists had come to believe that domain-specific knowledge and strong methods are the bases of expertise and intelligence—that the chess champion would fail to deter the belligerent neighbor nation. If they are right, there may not be a simple way to make people better general problem solvers. Siegler (1985, p. 184) sums up what this means for education: "Seen from this perspective, much of the task of education in problem solving may be to identify the encoding that we would like people to have on specific problems, and then to devise instructional methods to help them attain it." In other words, the educational challenge might be to identify the representations we want students to have in *specific domains* and then develop methods and curricula to teach those representations.

Weak Methods in the Schools

Despite this theoretical shift within the cognitive science community from favoring general strategies to arguing for domain-specific knowledge and skills, some educators still advocate teaching weak, general strategies. If children need thinking and learning skills but don't have them, they argue, the best educational strategy might be to teach those skills directly. Maybe teaching weak, general methods—skills that apply across the curriculum—can serve as a shortcut to higher intelligence and better school performance. Unfortunately, weak methods as they are traditionally taught seem to have little impact on student learning.

Traditional study and learning skills, though less general than means-end analysis, are general, weak methods. These skills are often the staples of junior high language-arts classes—taking notes, outlining, underlining, and figuring out words from context. We all believe that these strategies work, but research on traditional study skills shows that these skills are no more effective than simply reading and rereading a text. (See Anderson 1980.) Other researchers who have looked at the impact of teaching general reading-comprehension skills have found that teaching these skills increases students' awareness of the skills. However, even when students say they use the skills, the skills have little effect on their reading comprehension (Paris and Jacobs 1984). Similarly, the teaching of study or memorization skills in one subject has no impact on students' performance in other subjects. Often, it doesn't even occur to students to use the skills when studying a different subject. Typically, children will use strategies immediately after instruction, but will not use them on later occasions unless explicitly told to do so by the instructor.

Research shows that either the teaching of traditional study skills has no impact on learning or else the skills fail to transfer from the learning context to other situations. Either way, teaching these general skills is not the path to expertise and enhanced academic performance.

A wide variety of books and commercially available courses attempt to teach general cognitive and thinking skills. (For reviews and evaluations see Nickerson et al. 1985, Segal et al. 1985, and Chipman et al. 1985.) Analysis and evaluation of these programs again fail to support the belief that the teaching of general skills enhances students' overall performance.

Most of these programs teach general skills in stand-alone courses, separate from subject-matter instruction. The assumption is that students would find it too difficult to learn how to think and to learn subject content simultaneously. Like the early AI and cognitive science that inspire them, the courses contain many formal problems, logical puzzles, and games. The assumption is that the general, weak methods that work on these problems will work on problems in all subject domains.

A few of these programs, such as the Productive Thinking Program (Covington 1985) and Instrumental Enrichment (Feuerstein et al. 1985), have undergone extensive evaluation. The evaluations consistently report that students improve on problems like those contained in the course materials but show only limited improvement on novel problems or problems unlike those in the materials (Mansfield et al. 1978; Savell et al. 1986). The programs provide extensive practice on the specific kinds of problems that their designers want children to master. Children do improve on those problems, but this is different from developing *general* cognitive skills. After reviewing the effectiveness of several thinking-skills programs, one group of psychologists concluded that "there is no strong evidence that students in any of these thinking-skills programs improved in tasks that were dissimilar to those already explicitly practiced" (Bransford et al. 1985, p. 202). Students in the programs don't become more intelligent generally; the general problem-solving and thinking skills they learn do not transfer to novel problems. Rather, the programs help students become experts in the domain of puzzle problems.

The evaluations of these programs undercut the basic assumption about the power of weak methods in another way, too. If general skills, or weak methods, are the stuff of intelligence, then teaching these skills to students who had not previously used them should improve their performance. This doesn't happen. The programs don't help all students who were initially naive about the general skills taught. Typically, these programs help low-performing students most, average students some, and more able students hardly at all (Nickerson et al. 1985, p. 325).

Although we should not dismiss approaches that might help low-achieving students, this inverse pattern—low achievers benefiting most and high achievers hardly at all—is exactly what we would expect if school performance depends on domain-specific knowledge and strong methods. Low-performing students have neither general cognitive skills nor domain-specific knowledge. Teaching low achievers general skills can only help. The higher the level of initial performance, though, the more domain-specific knowledge a child has. If you have domain-specific knowledge and strong methods to go with it, why use weak ones? If you know all the standard variations on the Queen's Gambit Declined, why choose chess moves at random; why even rely on general chess principles? Teaching general cognitive skills to able students (even able students who haven't heard of those skills) doesn't improve their performance, because they are already relative experts. Able students already have domain-specific knowledge and use strong methods.

Evidence from the laboratory and the classroom argues against a fundamental role for weak methods and general skills in expertise and learning. Weak methods, in the guise of study skills, thinking-skills curricula, or critical-thinking programs, are not a short-cut to improved educational outcomes.

By the mid 1970s, then, most cognitive theorists recognized that domain-specific knowledge and strong methods were keys to expert performance and human intelligence. At that point, many would have bet against the chess champion's having a successful diplomatic career.

This message was picked up by some educators, and it has even reached the general public. E. D. Hirsch's *Cultural Literacy* (1987) is a thoughtful, sustained, and highly popular presentation of how domain-specific knowledge and skills are fundamental to literacy. Chapter 2 of Hirsch's book is an extended discussion of how cognitive research supports this educational philosophy. According to Hirsch, the research should make us skeptical of attempts to teach reading, writing, and critical thinking as general cognitive skills applicable to novel problems. Skilled performance in these subjects, like skilled performance in Simon and Chase's chess studies, demands an extensive store of domain-specific knowledge. "General programs contrived to teach general skills are ineffective," Hirsch argues (p. 61).

Hirsch characterizes the "critical thinking" movement—an attempt to teach weak methods—as a well-intentioned program "to take children beyond the minimal basic skills required by state guidelines and to encourage the teaching of 'higher order' skills" (1987, p. 132). The danger, as he sees it, is that advocates of higher-order thinking tend to ignore the importance of "mere facts." According to Hirsch, "we should direct our attention undeviatingly toward what schools teach" (p. 19).

There are also dangers associated with arguments, such as Hirsch's, for the primacy of domain-specific knowledge and skills. It is easy to oversimplify and misinterpret what the research means for educational practice. Certainly the research implies that we can't ignore "mere facts" in school instruction—domain knowledge is essential. But, conversely, curricula that merely transmit facts aren't desirable either. Cognitive research also implies that we have to be as concerned with how we teach as we are with what we teach. The danger with cultural literacy is embracing the what to the detriment of the how. Lists of proper nouns, such as appear in the appendix to Hirsch's book, might help outline curricular content, but they say nothing about how to teach that content effectively. Researchers have known for a long time that teaching word meanings to children can increase vocabulary knowledge, but more vocabulary knowledge doesn't necessarily improve reading comprehension. If better reading comprehension is the goal, how one teaches vocabulary matters. Similarly, current social studies texts may present the facts about geography or history, but fail to teach course content so that students have an understanding of geography or history. As we will see, how texts present the facts is vitally important.

Finally, the majority of the cognitive research Hirsch cites was done in the 1970s. But cognitive research didn't stop then. The prevailing view of 20 years ago was not the final word, nor should it necessarily guide educational practice. Research that started to appear in the early 1980s suggests that domain-specific knowledge and skills are necessary for expert performance but may not be sufficient. There is more to intelligence and expert performance than domain knowledge.

Metacognition

Around 1980, cognitive scientists introduced a new element, called *metacognition,* into discussions of intelligence and expert performance. Metacognition is the ability to think about thinking, to be consciously aware of oneself as a problem solver, and to monitor and control one's mental processing.

John Flavell, one of the developers of this notion, described metacognition as the fourth and highest level of mental activity

(Flavell and Wellman 1977). At the lowest level are the hard-wired, basic processes such as matching the contents of working memory to conditions on production rules. At the next level are things like knowing 9 × 7 = 63, being able to recall your mother's maiden name, and having command of sufficient schemas or facts to be culturally literate. At the third level are strategies, weak or strong methods, which we voluntarily and consciously use. For example, you might repeat a phone number silently to keep it active in working memory, or you might use Siegler's rule IV to solve a balance-scale problem. The fourth level is the metacognitive level—the knowledge, awareness, and control of the three lower levels. It is our conscious awareness of ourselves (and, by extension, others) as problem solvers.

Research on metacognition has shown that knowledge, awareness, and control of mental abilities develop with age and experience. As children mature, for example, they develop a much better sense of how many items they can hold in short-term memory. Four-year-old children can usually hold about three items, but will predict they can remember eight. Adolescents have a short-term memory capacity of about six items and accurately predict this capacity (Yussen and Levy 1975).

Children's performance on other memory tasks also provides evidence for metacognitive development. When experimenters give children a list of items to study and tell the children they can take as much time as they need to memorize the list, older children perform better than younger children (Flavell 1979). Although young children might lack effective memory strategies or lack background knowledge needed for the task, there is an independent metacognitive trend as well. Preschool children, after taking as much time as they want to study the list, think they have learned it completely but do very poorly when tested. In contrast, when elementary school children say they have learned the list, they can recall it accurately. The younger children don't know how to use study time effectively and have no idea if they have learned the list or not. It seems that the younger children don't know how to learn and don't know when they have learned.

Children's understanding of texts and stories shows a similar developmental trend. Even young children grasp the essential gist of a story. If given sufficient time to study a text, children at every age can remember more about it. Children younger than about 12, though, do not use the study time effectively. They remember more about the text, but tend to remember more details or isolated ideas from the text. They don't remember more about the text's themes or about how those themes interrelate. In short, before age 12 children don't seem to know what kinds of things are important for better understanding of texts and can't direct their mental energy to those things. The younger children lack important reading comprehension strategies, or, if they have the strategies, they lack control over them. They have weaknesses at Flavell's third and fourth cognitive levels. In contrast, children of age 12 and older usually remember more of the text's important ideas after additional study. The older children know what is important in texts, have strategies for reading texts and studying that are directed at those important features, know how and when to use the strategies, and can monitor their use of them. They can control their cognitive activity—they have metacognitive skills. Ann Brown and Judy Deloache, who reported some of these results, conclude that "one main aspect of 'what develops' is metacognition—the voluntary control an individual has over his own cognitive processes," and that "the growth of metacognitive abilities underlies many of the behavioral changes that take place with development" (Brown and DeLoache 1978, p. 26).

Hirsch emphasizes the necessity of domain-specific knowledge in learning and doesn't mention metacognition explicitly. Nonetheless, the importance of metacognition is implicit in his diagnosis of literacy problems. Although domain-specific knowledge contributes to expertise in all domains, in reading (as Hirsch carefully explains) background knowledge—knowledge that goes beyond what is literally printed on the page—is crucial for comprehension. Teaching the schemas of cultural literacy is intended to give students the background knowledge needed to be culturally literate. Note that such knowledge would fall into level 2 of Flavell's taxonomy: facts stored for recall in long-term memory.

But Hirsch alludes to knowledge that literate individuals have that would fall into Flavell's fourth level: "In effective reading, one must not only call up one's own schematic associations but also *monitor* [my italics] whether they are appropriate ones shared by the wider speech community." Literate adults do this automatically, but "young children and other semi-literates do not confidently know what other members of the speech community can be expected to know." They lack "readily accessible information about what is shared by others" (p. 68).

Calling up one's one schematic associations is a level-2 cognitive process. Monitoring their appropriateness, on the other hand, is a level-4, metacognitive process. Similarly, knowing about or estimating what other members of the speech community might know is also a metacognitive task; it involves the ability to envision other people as problem solvers whose minds work similarly to one's own. This is just to say that reading comprehension involves more than extensive cultural background knowledge. Minimally, reading comprehension also requires metacognitive monitoring skills. If students lack these skills, no amount of cultural knowledge on its own can make them literate.

Metacognition and Intelligent Novices

Metacognition is an important addition to a theory of expertise and intelligence. The results of the research discussed so far—the studies of memory, learning skills, and reading—are consistent with the contention that metacognitive skills are high-order skills but domain-specific skills nonetheless. Clearly, it is possible for a person to be expert and metacognitively sophisticated in one domain but not in others. Our cultural stereotype of the absent-minded professor—a scientist or scholar who is expert and metacognitively capable in an academic domain but inept and unaware outside that academic specialty, particularly in everyday life—derives from this possibility.

Other results, though, suggest that metacognitive skills are general skills—skills that some people can apply across domains and in domains where they have little prior background knowledge. Everyday experience suggests that there are *intelligent novices:* some novices learn new domains more quickly than other novices. Research tells us that one thing that makes some novices more intelligent than others is their metacognitive skills.

As part of an experiment, John Bransford, an expert cognitive psychologist who has done work on math learning, tried to learn physics from a textbook with the help of an expert physicist. He kept a diary of his learning experiences and recorded the skills and strategies most useful to him (Brown et al. 1983). Among the things he listed were (1) awareness of the difference between understanding and memorizing material and knowledge of which mental strategies to use in each case; (2) ability to recognize which parts of the text were difficult, which dictated where to start reading and how much time to spend; (3) awareness of the need to take problems and examples from the text, order them randomly, and then try to solve them; (4) knowing when he didn't understand, so he could seek help from the expert; and (5) knowing when the expert's explanations solved his immediate learning problem. These are all metacognitive skills; they all involve awareness and control of the learning problem that Bransford was trying to solve. Bransford might have learned these skills originally in one domain (cognitive psychology), but he could apply them as a novice when trying to learn a second domain (physics).

This self-experiment led Bransford and his colleagues to examine in a more controlled way the differences between expert and less-skilled learners. They found that the behavior of intelligent novices contrasted markedly with that of the less skilled. Intelligent novices used many of the same strategies Bransford had used to learn physics. Less-skilled learners used few, if any, of them. The less-skilled did not always appreciate the difference between memorization and comprehension and seemed to be unaware that different learning strategies should be used in each case (Bransford et al. 1986; Bransford and Stein 1984). These students were less likely to notice whether texts were easy or difficult, and thus were less able to adjust their strategies and their study time accordingly (Bransford et al. 1982). Less-able learners were unlikely to use self-tests and self-questioning as sources of feedback to correct misconceptions and inappropriate learning strategies (Brown et al. 1983; Stein et al. 1982).

Hirsch, in his discussion of reading, notes how expert readers "monitor" their schematic associations. Monitoring comprehension is also a metacognitive skill. Ellen Markman (1985) studied this skill by having students in grades 3 through 6 read short passages which they had never seen before and which contained obvious contradictions. For example, a passage about ants might say in one place that ants navigate by leaving a chemical trail which they can smell and in another place that ants have no sense of smell. Most of the younger children and even a few of the older ones were oblivious to the inconsistencies; they weren't monitoring their comprehension. Children did improve on the task with age, so Markman first interpreted the results in terms of developmental differences between younger and older children. Subsequent research supported a more general conclusion: that the ability to apply this metacognitive skill differentiated strong from weak learners at all ages.

The ability to monitor comprehension is an essential learning skill. Often poor students are totally unaware that they don't comprehend class material. If they aren't aware that they have a learning problem, they can't take steps to overcome it.

Everyday experience suggests there are intelligent novices. Research tells us that metacognitive skills contribute to these expert learning performances. Some people develop these skills naturally; others do not. Those who do can become intelligent novices; those who don't may have difficulty learning.

Metacognition and Education

The importance of metacognition for education is that a child is, in effect, a universal novice, constantly confronted with novel learning tasks. In such a situation it would be most beneficial to be an intelligent novice. What is encouraging is that the research also shows that it is possible to teach children metacognitive skills and when to use them. If we can do this, we will be able to help children become intelligent novices; we will be able to teach them how to learn.

Just as there are basic math and reading skills, there are basic metacognitive skills. Among the basic metacognitive skills are the abilities to predict the results of one's own problem-solving actions, to check the results of one's own actions (Did it work?), to monitor one's progress toward a solution (How am I doing?), and to test how reasonable one's actions and solutions are against the larger reality (Does this make sense?). For example, a metacognitively adept chess player tries to predict the consequences of a series of moves, checks the results of those moves, and monitors whether those moves might contribute to a possible checkmate. Such a player also checks possible strategies against the larger reality. In a game against a higher-rated opponent, a metacognitively aware player would not look for an easy mating combination early in the game; a quick reality check would convince him that such a strategy doesn't make sense. Brown and DeLoache (1978, p. 15) call these skills "the basic characteristics of efficient thought." To become efficient thinkers—intelligent novices—students have to learn the skills and learn when to use them. Although a student might first learn the skills in the context of some specific subject matter (as Bransford first learned them in psychology), once he or she has learned them, the student can apply the skills in any learning situation—*if* the student has also learned that these skills are applicable and useful in any learning situation. Cognitive scientists call instruction that teaches students metacognitive skills and when to use them *metacognitively aware* instruction.

How does one teach like this? One can think of metacognitive skills as the ability to be critical of one's own problem solving. Metacognitively aware instruction attempts to transfer the critic's role from the teacher to the student. The transfer occurs in stages. Initially the teacher models the critic's role for the students. Gradually, the students begin to share this critical, metacognitive role with the teacher; eventually they can take on the role themselves, with the teacher standing by to provide coaching when the students falter. As the children become more metacognitively adept—more self-critical—the teacher cedes the critic's role entirely to them. Researchers describe this transition from the teacher's modeling and control to the students' control as *scaffolding*. Instruction creates a scaffold to support learning, and then the scaffold is gradually dismantled as the students become increasingly self-critical.

One problem with metacognitive skills is that they are usually covert and implicit in expert performance. To teach these skills and when to use them, the instructor has to make metacognition overt and explicit. One effective way to do this is in group learning situations where teachers and students engage in dialogues about their joint learning and problem solving. Almost all children are able conversationalists and can play the dialogue game. Appropriately guided by the teacher, the dialogue can become a social, collaborative form of "thinking aloud" in which each member of the group makes his or her thinking overt. In such situations, first the teacher and then the students describe their problem-solving strategies, present their reasoning to the group, and then defend and justify it against criticisms. These group dialogues make reasoning, planning, and monitoring public and shared. The children begin to see cognitive processes in action and understand how they can be monitored and controlled.

In experimental situations, Brown, working with various collaborators, has used metacognitively aware instruction to teach memory strategies to mildly retarded students (Brown et al. 1979), text-summarizing skills to college students (Brown et al. 1981), and analogical reasoning to children as young as 3 and 4 years (Brown 1989), all with great success. As we will see, the approach also works in the classroom, where metacognitively aware instruction can improve students' understanding of scientific reasoning, reading comprehension, and the writing process.

The Final Element: General Skills Again

The final element leading to the new synthesis is general skills—they just won't go away. Besides metacognitive ability, there is another trait that intelligent novices show. Research scientists are an extreme but illustrative example. Scientists don't just apply their extensive domain knowledge to solve standard, textbook problems; rather, they formulate new problems and discover new solutions. Deep, domain-specific knowledge is fundamental to their performance, but so are general skills and strategies. Einstein attributed his interest in Brownian motion to childhood experiences watching the patterns formed by smoke rising from his uncle's pipe. My graduate adviser, a biophysicist, would help us graduate students understand an abstract geometrical theorem about how objects can move in space by referring to it as "the hatpin-through-the-grapefruit theorem."

On the frontiers of science, or in any creative endeavor, everyone is a novice in the sense that prior knowledge is not directly applicable. Here, too, it helps to be an intelligent novice. Scientists, scholars, artists, and skilled managers all have to take what they know and stretch it to pose and answer novel problems. They have to transfer prior learning to new situations. Intelligent novices, like Einstein and my adviser, often use general strategies to do this. They use strategies such as modeling and analogy—weak methods applicable across many domains—in their attempts to apply what they know flexibly and creatively in new ways. Their extensive domain-specific knowledge interacts with general strategies to help them acquire new knowledge. General strategies do seem to have a role in intelligence and in expert performance.

General strategies should also be helpful to those universal novices, children. Although research on the teaching of such study skills as memory strategies, underlining, and note taking has shown that after instruction children do not apply these skills spontaneously, it is still hard to accept that such skills are not useful. Maybe the problem is not with general skills but with how we have traditionally tried to teach them.

Let us review briefly what we think we know about intelligence and expertise. First, we have seen the importance of domain-specific knowledge for expert performance. The scientist's use of general strategies is based on deep understanding of at least one scientific domain. We can agree with the advocates of domain specificity that general programs contrived to teach general skills are ineffective, but that leaves open the possibility of teaching general skills within specific subject-matter instruction. General strategies do need a knowledge base on which to work, but once learned in a specific context they should be applicable in other domains. Second, both from the Johnson-Laird experiment (figure 3.1) and from research on study skills, we have seen that adults and children have difficulty transferring a skill or a strategy from one context to a similar context. In some cases, subjects could make the transfer between contexts only after the experimenter *told* them that the strategy applied in the new situation.

Perhaps, just as children have to be taught metacognitive skills and when to use them, they have to be taught general learning strategies and when to use them. Perhaps, then, previous attempts to teach general skills failed because course designers and instructors overestimated children's ability to generalize from one learning situation to another. Maybe children don't see how and why the situations are similar. In general-strategy and learning-skill instruction, rather than assume that students see the similarities between various learning situations, perhaps we should explicitly tell them how and why the situations are similar. This has led cognitive scientists to think that general-strategy instruction has a place in schools, but that strategy instruction has to be *informed*. By this they mean that strategy instruction should include explicit descriptions of the strategies, instruction about *when* the strategies are useful, and an explanation of *why* they are useful.

Paris et al. (1982) ran an experiment in which they compared informed instruction with a more traditional approach to strategy instruction. On each of the first two days of the experiment, they had 7- and 8-year-old children study sets of 24 pictures. After a period of study, they asked the children to remember as many of the pictures as they could. On average, the children could recall 12 or 13 pictures. On the third day, all the children were taught memory strategies: naming or labeling each item, sorting the items into related groups (clustering), learning the items one group at a time, and then testing themselves by trying to recall the items in groups (a metacognitive skill).

Half the children (the control group) saw the instructor demonstrate the memory strategies and were permitted to practice them, but were given no explanations of why the strategies worked and no feedback on their performance when they tried to use them. In short, the control group received traditional strategy instruction.

The other children (the experimental group) received the same instruction as the control group but in addition were told why the strategies worked and when to use them. Also, when the children in this group used a strategy they received immediate feedback on how successful they had been with it. This group received informed instruction.

Immediately after learning the strategies, children in the control group could recall on average 16 pictures and children in the experimental group 19. The experiment was continued for two more days. By the fifth day, children in the experimental group could still recall 16 items. These children continued to use the memory strategies spontaneously. Without being told, they continued to label, sort, cluster, and self-test. In contrast, by the fifth day children in the control group had fallen to the pre-instruction level of 12 to 13 items and had reverted to passive, pre-instruction learning strategies, such as looking at the pictures and trying to remember them. This experiment shows that children will use a strategy spontaneously—they will transfer it to a new situation—if they understand why it works and when it can help them learn. Informed strategy instruction works; traditional instruction doesn't.

General thinking, learning, and study strategies are important elements of intelligence and expert performance, and now it seems we may know how to teach them. According to Brown (1985, p. 335), "ideal cognitive skills training programs include practice in the specific task-appropriate strategies, direct instruction in the orchestrating, overseeing and monitoring of these skills, and information concerning the significance of those activities." Such instruction recognizes the necessity of domain-specific knowledge in that the strategies are specific, task appropriate, and integrated into subject-matter learning. The instruction is also metacognitively aware, in that the children receive direct instruction about how to

monitor and control their problem solving. It is informed in that children learn why the strategies work.

The New Synthesis and the Teaching of High-Order Skills

The new synthesis suggests that domain-specific knowledge, metacognitive skills, and general strategies are all elements of human intelligence and expert performance.

Where does this leave our chess champion? His detailed knowledge of chess won't help him in his new diplomatic career, because by definition it is domain specific. On the other hand, even if he has some potentially relevant general skills, he may not be able to transfer them readily to new problem domains—not everyone can do that. Much depends on whether the champion is an intelligent novice or not. Is his expertise narrowly confined to chess, or does he have the metacognitive insight to be an effective, rapid learner? If the latter, then maybe with some tutoring in foreign affairs and some on-the-job experience he could rapidly become a national security expert. Some chess champions—like some college professors and some school children—are naturally intelligent novices; others aren't. Whether our chess champion will succeed depends on what cognitive skills he has beyond his chess expertise.

We are just beginning to see what the new synthesis might mean for educational practice. Many of the examples in the following chapters illustrate how this latest theory of human intelligence might be applied in the classroom. For education, the most important implication of the theory is that how we teach is as important as what we teach. Domain-specific knowledge and skills are essential to expertise; however, school instruction must also be metacognitively aware and informed.

Most important, innovative classroom practices based on the new synthesis can help us achieve our goal of teaching high-order skills to all students. In chapter 1, we initially identified high-order skills as the skills that students need to achieve the higher NAEP proficiency levels. We observed that these higher proficiencies demand that students solve complex problems, for which there often are no standard solution procedures and no single correct answer. Students with high-order skills can use their knowledge flexibly to solve ill-structured, novel problems.

This characterization of high-order skills relies primarily on the kinds of problems students can solve and on students' observed behaviors. But "high-order" also refers to the thought processes needed to solve such problems and guide such behaviors. Susan Chipman, Program Manager for Cognitive Science at the Office of Naval Research, argues that "behind our choice of the term 'higher-order,' there are strong intuitions about the way in which our cognitive activities are structured and controlled" (Chipman 1992). These intuitions link high-order skills with our current theory of intelligence and expert performance.

First, Chipman points out that higher-order skills in a subject domain, such as those needed to solve ill-structured complex problems, are skills grounded in deep factual and procedural knowledge of the domain. As the new synthesis implies, high-order skills require extensive domain knowledge.

Second, she notes that students who genuinely possess high-order skills in a subject domain not only have the requisite factual and procedural knowledge, they also can recognize *when* the knowledge is applicable and can use it appropriately. It is this feature of high-order skills that accounts for the flexible, spontaneous use of knowledge in novel situations. This connects high-order skills with the notion of transfer. High-order skills should transfer from school learning to real-world situations and allow students to use what they already know to learn new things more rapidly. The key to transfer, and so to high-order skills, is knowing when to use knowledge. If we want to teach high-order skills, then, as the new synthesis says, the instruction should be informed.

Third, implicit in our use of "high-order," according to Chipman, are intuitions about how students control and monitor their cognitive skills. High-order skills, in this sense, involve awareness of what is happening in working memory, of how those processes determine eventual action, and of how to control those processes. Thus, metacognitive abilities are implicit in our notion of high-order skills. For this reason, if we want students to acquire high-order skills, instruction must be metacognitively explicit.

In short, high-order skills require extensive domain knowledge, understanding when to use the knowledge, and metacognitive monitoring and control. Students who have these things can solve novel, ambiguous problems; students who have high-order skills are intelligent novices.

For these reasons, instruction based on elements of the new synthesis is our best educational bet if we want all students to have the knowledge and skills that in past generations have been confined to the college-bound elite—if we want all students to acquire high-order skills. Educational practice grounded in cognitive theory, Lauren Resnick (1986, p. 43) writes, "would transform the whole curriculum in fundamental ways. It would treat the development of higher-order skills as the paramount goal of *all* schooling."

Transforming the curriculum to meet this goal won't be easy. We will have to rethink, or at least reevaluate, much of our received wisdom about educational policy, practice, and standards and about teacher training. We will have to restructure our schools—starting in the classrooms, where teachers interact with students. We will need teachers who can create and maintain learning environments where students can become intelligent novices. Many of us will have to change our representations of what schools and schooling are.

Admittedly, there is much we still don't know about how our minds work, how children best learn, and how to design better schools. Nonetheless, as a first step, we can start applying what we already know to improve instruction in what Resnick calls the "enabling" or "tool" domains: mathematics, science, reading, and writing. Mastery of these domains is necessary for advanced learning in more specialized subjects. We can teach these enabling domains as high-order cognitive skills to all students, as the following chapters will show.