

Critical Thinking

Why Is It So Hard to Teach?

By Daniel T. Willingham

Virtually everyone would agree that a primary, yet insufficiently met, goal of schooling is to enable students to think critically. In layperson's terms, critical thinking consists of seeing both sides of an issue, being open to new evidence that disconfirms your ideas, reasoning dispassionately, demanding that claims be backed by evidence, deducing and inferring conclusions from available facts, solving problems, and so forth. Then too, there are specific types of critical thinking that are characteristic of different subject matter: That's what we mean when we refer to "thinking like a scientist" or "thinking like a historian."

This proper and commonsensical goal has very often been translated into calls to teach "critical thinking skills" and "higher-order thinking skills"—and into generic calls for teaching students to make better judgments, reason more logically, and so forth. In a recent survey of human resource officials¹ and in testimony delivered just a few months ago before the Senate Finance Committee,² business leaders have repeatedly exhorted schools to do a better job of teaching students to think critically. And they are not alone. Organizations and initiatives involved in education reform, such as the National Center on Education and the Economy, the American Diploma Project, and the Aspen Institute, have pointed out the need for students

Daniel T. Willingham is professor of cognitive psychology at the University of Virginia and author of Cognition: The Thinking Animal as well as over 50 articles. With Barbara Spellman, he edited Current Directions in Cognitive Science. He regularly contributes to American Educator by writing the "Ask the Cognitive Scientist" column. His research focuses on the role of consciousness in learning.

to think and/or reason critically. The College Board recently revamped the SAT to better assess students' critical thinking. And ACT, Inc. offers a test of critical thinking for college students.

These calls are not new. In 1983, *A Nation At Risk*, a report by the National Commission on Excellence in Education, found that many 17-year-olds did not possess the "higher-order" intellectual skills" this country needed. It claimed that nearly 40 percent could not draw inferences from written material and only one-fifth could write a persuasive essay.

Following the release of *A Nation At Risk*, programs designed to teach students to think critically across the curriculum became extremely popular. By 1990, most states had initiatives designed to encourage educators to teach critical thinking, and one of the most widely used programs, Tactics for Thinking, sold 70,000 teacher guides.³ But, for reasons I'll explain, the programs were not very effective—and today we still lament students' lack of critical thinking.

After more than 20 years of lamentation, exhortation, and little improvement, maybe it's time to ask a fundamental question: Can critical thinking actually be taught? Decades of cognitive research point to a disappointing answer: not really. People who have sought to teach critical thinking have assumed that it is a skill, like riding a bicycle, and that, like other skills, once you learn it, you can apply it in any situation. Research from cognitive science shows that thinking is not that sort of skill. The processes of thinking are intertwined with the content of thought (that is, domain knowledge). Thus, if you remind a student to "look at an issue from multiple perspectives" often enough, he will learn that he ought to do so, but if he doesn't know much about

Critical thinking is not a set of skills that can be deployed at any time, in any context. It is a type of thought that even 3-year-olds can engage in—and even trained scientists can fail in.

an issue, he *can't* think about it from multiple perspectives. You can teach students maxims about how they ought to think, but without background knowledge and practice, they probably will not be able to implement the advice they memorize. Just as it makes no sense to try to teach factual content without giving students opportunities to practice using it, it also makes no sense to try to teach critical thinking devoid of factual content.

In this article, I will describe the nature of critical thinking, explain why it is so hard to do and to teach, and explore how students acquire a specific type of critical thinking: thinking scientifically. Along the way, we'll see that critical thinking is not a set of skills that can be deployed at any time, in any context. It is a type of thought that even 3-year-olds can engage in—and even trained scientists can fail in. And it is very much dependent on domain knowledge and practice.

Why Is Thinking Critically So Hard?

Educators have long noted that school attendance and even academic success are no guarantee that a student will graduate an effective thinker in all situations. There is an odd tendency for rigorous thinking to cling to particular examples or types of problems. Thus, a student may have learned to estimate the answer to a math problem before beginning calculations as a way of checking the accuracy of his answer, but in the chemistry lab, the same student

calculates the components of a compound without noticing that his estimates sum to more than 100 percent. And a student who has learned to thoughtfully discuss the causes of the American Revolution from both the British and American perspectives doesn't even think to question how the Germans viewed World War II. Why are students able to think critically in one situation, but not in another? The brief answer is: Thought processes are intertwined with what is being thought about. Let's explore this in depth by looking at a particular kind of critical thinking that has been studied extensively: problem solving.

Imagine a seventh-grade math class immersed in word problems. How is it that students will be able to answer one problem, but not the next, even though mathematically both word problems are the same, that is, they rely on the same mathematical knowledge? Typically, the students are focusing on the scenario that the word problem describes (its surface structure) instead of on the mathematics required to solve it (its deep structure). So even though students have been taught how to solve a particular type of word problem, when the teacher or textbook changes the scenario, students still struggle to apply the solution because they don't recognize that the problems are mathematically the same.

Thinking Tends to Focus on a Problem's "Surface Structure"

To understand why the surface structure of a problem is so distracting and, as a result, why it's so hard to apply familiar solutions to problems that appear new, let's first consider how you understand what's being asked when you are given a problem. Anything you hear or read is automatically interpreted in light of what you already know about similar subjects. For example, suppose you read these two sentences: "After years of pressure from the film and television industry, the President has filed a formal complaint with China over what U.S. firms say is copyright infringement. These firms assert that the Chinese government sets stringent trade restrictions for U.S. entertainment products, even as it turns a blind eye to Chinese companies that copy American movies and television shows and sell them on the black market." Background knowledge not only allows you to comprehend the sentences, it also has a powerful effect as you continue to read because it narrows the interpretations of new text that you will entertain. For example, if you later read the word "Bush," it would not make you think of a small shrub, nor would you wonder whether it referred to the former President Bush, the rock band, or a term for rural hinterlands. If you read "piracy," you would not think of eye-patched swabbies shouting "shiver me timbers!" The cognitive system gambles that incoming information will be related to what you've just been thinking about. Thus, it significantly narrows the scope of possible interpretations of words, sentences, and ideas. The benefit is that comprehension proceeds faster and more smoothly; the cost is that the deep structure of a problem is harder to recognize.

The narrowing of ideas that occurs while you read (or

How Do Cognitive Scientists Define Critical Thinking?

From the cognitive scientist's point of view, the mental activities that are typically called critical thinking are actually a subset of three types of thinking: reasoning, making judgments and decisions, and problem solving. I say that critical thinking is a subset of these because we think in these ways all the time, but only sometimes in a critical way. Deciding to read this article, for example, is not critical thinking. But carefully weighing the evidence it presents in order to decide whether or not to believe what it says is. *Critical* rea-

soning, decision making, and problem solving—which, for brevity's sake, I will refer to as critical thinking—have three key features: effectiveness, novelty, and self-direction. Critical thinking is effective in that it avoids common pitfalls, such as seeing only one side of an issue, discounting new evidence that disconfirms your ideas, reasoning from passion rather than logic, failing to support statements with evidence, and so on. Critical thinking is novel in that you don't simply remember a solution or a

situation that is similar enough to guide you. For example, solving a complex but familiar physics problem by applying a multi-step algorithm isn't critical thinking because you are really drawing on memory to solve the problem. But devising a new algorithm is critical thinking. Critical thinking is self-directed in that the thinker must be calling the shots: We wouldn't give a student much credit for critical thinking if the teacher were prompting each step he took.

—D.W.

listen) means that you tend to focus on the surface structure, rather than on the underlying structure of the problem. For example, in one experiment,⁴ subjects saw a problem like this one:

Members of the West High School Band were hard at work practicing for the annual Homecoming Parade. First they tried marching in rows of 12, but Andrew was left by himself to bring up the rear. Then the director told the band members to march in columns of eight, but Andrew was still left to march alone. Even when the band marched in rows of three, Andrew was left out. Finally, in exasperation, Andrew told the band director that they should march in rows of five in order to have all the rows filled. He was right. Given that there were at least 45 musicians on the field but fewer than 200 musicians, how many students were there in the West High School Band?

Earlier in the experiment, subjects had read four problems along with detailed explanations of how to solve each one, ostensibly to rate them for the clarity of the writing. One of the four problems concerned the number of vegetables to buy for a garden, and it relied on the same type of solution necessary for the band problem—calculation of the least common multiple. Yet, few subjects—just 19 percent—saw that the band problem was similar and that they could use the garden problem solution. Why?

When a student reads a word problem, her mind interprets the problem in light of her prior knowledge, as happened when you read the two sentences about copyrights and China. The difficulty is that the knowledge that seems relevant relates to the surface structure—in this problem, the reader dredges up knowledge about bands, high school, musicians, and so forth. The student is unlikely to read the problem and think of it in terms of its deep structure—using the least common multiple. The surface structure of the problem is overt, but the deep structure of the problem is not. Thus, people fail to use the first prob-

lem to help them solve the second: In their minds, the first was about vegetables in a garden and the second was about rows of band marchers.

With Deep Knowledge, Thinking Can Penetrate Beyond Surface Structure

If knowledge of how to solve a problem never transferred to problems with new surface structures, schooling would be inefficient or even futile—but of course, such transfer does occur. When and why is complex,⁵ but two factors are especially relevant for educators: familiarity with a problem's deep structure and the knowledge that one should look for a deep structure. I'll address each in turn.

When one is very familiar with a problem's deep-structure, knowledge about how to solve it transfers well. That familiarity can come from long-term, repeated experience with one problem, or with various manifestations of one type of problem (i.e., many problems that have different surface structures, but the same deep structure). After repeated exposure to either or both, the subject simply perceives the deep structure as part of the problem description. Here's an example:

A treasure hunter is going to explore a cave up on a hill near a beach. He suspected there might be many paths inside the cave so he was afraid he might get lost. Obviously, he did not have a map of the cave; all he had with him were some common items such as a flashlight and a bag. What could he do to make sure he did not get lost trying to get back out of the cave later?

The solution is to carry some sand with you in the bag, and leave a trail as you go, so you can trace your path back when you're ready to leave the cave. About 75 percent of American college students thought of this solution—but only 25 percent of Chinese students solved it.⁶ The experimenters suggested that Americans solved it because most grew up hearing the story of Hansel and Gre-

tel, which includes the idea of leaving a trail as you travel to an unknown place in order to find your way back. The experimenters also gave subjects another puzzle based on a common Chinese folk tale, and the percentage of solvers from each culture reversed. (To read the puzzle based on the Chinese folk tale, and the tale itself, go to www.aft.org/pubs-reports/american_educator/index.htm.)

It takes a good deal of practice with a problem type before students know it well enough to immediately recognize its deep structure, irrespective of the surface structure, as Americans did for the Hansel and Gretel problem. American subjects didn't think of the problem in terms of sand, caves, and treasure; they thought of it in terms of finding something with which to leave a trail. The deep structure of the problem is so well represented in their memory, that they immediately saw that structure when they read the problem.

Looking for a Deep Structure Helps, but It Only Takes You So Far

Now let's turn to the second factor that aids in transfer despite distracting differences in surface structure—knowing to look for a deep structure. Consider what would happen if I said to a student working on the band problem, "this one is similar to the garden problem." The student would understand that the problems must share a deep structure and would try to figure out what it is. Students can do something similar without the hint. A student might think "I'm seeing this problem in a math class, so there must be a math formula that will solve this problem." Then he could scan his memory (or textbook) for candidates, and see if one of them helps. This is an example of what psychologists call metacognition, or regulating one's thoughts. In the introduction, I mentioned that you can teach students maxims about how they ought to think.

Critical Thinking Programs: Lots of Time, Modest Benefit

Since the ability to think critically is a primary goal of education, it's no surprise that people have tried to develop programs that could directly teach students to think critically without immersing them in any particular academic content. But the evidence shows that such programs primarily improve students' thinking with the sort of problems they practiced in the program—not with other types of problems. More generally, it's doubtful that a program that effectively teaches students to think critically in a variety of situations will ever be developed.

As the main article explains, the ability to think critically depends on having adequate content knowledge; you can't think critically about topics you know little about or solve problems that you don't know well enough to recognize and execute the type of solutions they call for.

Nonetheless, these programs do help us better understand what can be taught, so they are worth reviewing briefly.

A large number of programs¹ designed to make students better thinkers are available, and they have

some features in common. They are premised on the idea that there is a set of critical thinking skills that can be applied and practiced across content domains. They are designed to supplement regular curricula, not to replace them, and so they are not tied to particular content areas such as language arts, science, or social studies. Many programs are intended to last about three years, with several hours of instruction (delivered in one or two lessons) per week. The programs vary in how they deliver this instruction and practice. Some use abstract problems such as finding patterns in meaningless figures (Reuven Feuerstein's Instrumental Enrichment), some use mystery stories (Martin Covington's Productive Thinking), some use group discussion of interesting problems that one might encounter in daily life (Edward de Bono's Cognitive Research Trust, or CoRT), and so on. However it is implemented, each program introduces students to examples of critical thinking and then requires that the students practice such thinking themselves.

How well do these programs work? Many researchers have tried

to answer that question, but their studies tend to have methodological problems.² Four limitations of these studies are especially typical, and they make any effects suspect: 1) students are evaluated just once after the program, so it's not known whether any observed effects are enduring; 2) there is not a control group, leaving it unclear whether gains are due to the thinking program, to other aspects of schooling, or to experiences outside the classroom; 3) the control group does not have a comparison intervention, so any positive effects found may be due, for example, to the teacher's enthusiasm for something new, not the program itself; and 4) there is no measure of whether or not students can transfer their new thinking ability to materials that differ from those used in the program. In addition, only a small fraction of the studies have undergone peer review (meaning that they have been impartially evaluated by independent experts). Peer review is crucial because it is known that researchers unconsciously bias the design and analysis of their research to favor the conclusions they hope to see.³

Cognitive scientists refer to these maxims as metacognitive strategies. They are little chunks of knowledge—like “look for a problem’s deep structure” or “consider both sides of an issue”—that students can learn and then use to steer their thoughts in more productive directions.

Helping students become better at regulating their thoughts was one of the goals of the critical thinking programs that were popular 20 years ago. As the sidebar below explains, these programs are not very effective. Their modest benefit is likely due to teaching students to effectively use metacognitive strategies. Students learn to avoid biases that most of us are prey to when we think, such as settling on the first conclusion that seems reasonable, only seeking evidence that confirms one’s beliefs, ignoring countervailing evidence, overconfidence, and others.⁷ Thus, a student who has been encouraged many times to see both sides of an issue, for example, is probably more likely to spontane-

ously think “I should look at both sides of this issue” when working on a problem.

Unfortunately, metacognitive strategies can only take you so far. Although they suggest what you ought to do, they don’t provide the knowledge necessary to implement the strategy. For example, when experimenters told subjects working on the band problem that it was similar to the garden problem, more subjects solved the problem (35 percent compared to 19 percent without the hint), but most subjects, even when told what to do, weren’t able to do it. Likewise, you may know that you ought not accept the first reasonable-sounding solution to a problem, but that doesn’t mean you know how to come up with alternative solutions or weigh how reasonable each one is. That requires domain knowledge and practice in putting that knowledge to work.

Since critical thinking relies so heavily on domain

Studies of the Philosophy for Children program may be taken as typical. Two researchers⁴ identified eight studies that evaluated academic outcomes and met minimal research-design criteria. (Of these eight, only one had been subjected to peer review.) Still, they concluded that three of the eight had identifiable problems that clouded the researchers’ conclusions. Among the remaining five studies, three measured reading ability, and one of these reported a significant gain. Three studies measured reasoning ability, and two reported significant gains. And, two studies took more impressionistic measures of student’s participation in class (e.g., generating ideas, providing reasons), and both reported a positive effect.

Despite the difficulties and general lack of rigor in evaluation, most researchers reviewing the literature conclude that some critical thinking programs do have some positive effect.⁵ But these reviewers offer two important caveats. First, as with almost any educational endeavor, the success of the program depends on the skill of the teacher. Second, thinking programs look good when the outcome measure is quite similar to the material in the program. As one tests for transfer to more and more dissimilar material, the apparent effectiveness of the program

Knowing that one should think critically is not the same as being able to do so. That requires domain knowledge and practice.

rapidly drops.

Both the conclusion and the caveats make sense from the cognitive scientist’s point of view. It is not surprising that the success of the program depends on the skill of the teacher. The developers of the programs cannot anticipate all of the ideas—right or wrong—that students will generate as they practice thinking critically, so it is up to the teacher to provide the all-important feedback to the students.

It is also reasonable that the programs should lead to gains in abilities that are measured with materials similar to those used in the program.

The programs that include puzzles like those found on IQ tests, for instance, report gains in IQ scores. In an earlier column,* I described a bedrock principle of memory: You remember what you think about. The same goes for critical thinking: You learn to think critically in the ways in which you practice thinking critically. If you practice logic puzzles with an effective teacher, you are likely to get better at solving logic puzzles. But substantial improvement requires a great deal of practice. Unfortunately, because critical thinking curricula include many different types of problems, students typically don’t get enough practice with any one type of problem. As explained in the main article, the modest benefits that these programs seem to produce are likely due to teaching students metacognitive strategies—like “look at both sides of an issue”—that cue them to try to think critically. But knowing that one should think critically is not the same as being able to do so. That requires domain knowledge and practice.

—D.W.

*See “Students Remember ... What They Think About” in the Summer 2003 issue of *American Educator*; online at www.aft.org/pubs-reports/american_educator/summer2003/cogsci.html.

(*Endnotes on page 19*)

Teaching students to think critically probably lies in large part in enabling them to deploy the right type of thinking at the right time.

knowledge, educators may wonder if thinking critically in a particular domain is easier to learn. The quick answer is yes, it's a *little* easier. To understand why, let's focus on one domain, science, and examine the development of scientific thinking.

Is Thinking Like a Scientist Easier?

Teaching science has been the focus of intensive study for decades, and the research can be usefully categorized into two strands. The first examines how children acquire scientific concepts; for example, how they come to forgo naive conceptions of motion and replace them with an understanding of physics. The second strand is what we would call thinking scientifically, that is, the mental procedures by which science is conducted: developing a model, deriving a hypothesis from the model, designing an experiment to test the hypothesis, gathering data from the experiment, interpreting the data in light of the model, and so forth.[†] Most researchers believe that scientific thinking is really a subset of reasoning that is not different in kind from other types of reasoning that children and adults do.⁸ What makes it *scientific* thinking is knowing when to engage in such reasoning, and having accumulated enough relevant knowledge and spent enough time practicing to do so.

Recognizing *when* to engage in scientific reasoning is so important because the evidence shows that being able to reason is not enough; children and adults use *and* fail to

use the proper reasoning processes on problems that seem similar. For example, consider a type of reasoning about cause and effect that is very important in science: conditional probabilities. If two things go together, it's possible that one causes the other. Suppose you start a new medicine and notice that you seem to be getting headaches more often than usual. You would infer that the medication influenced your chances of getting a headache. But it could also be that the medication increases your chances of getting a headache only in certain circumstances or conditions. In conditional probability, the relationship between two things (e.g., medication and headaches) is dependent on a third factor. For example, the medication might increase the probability of a headache *only* when you've had a cup of coffee. The relationship of the medication and headaches is conditional on the presence of coffee.

Understanding and using conditional probabilities is essential to scientific thinking because it is so important in reasoning about what causes what. But people's success in thinking this way depends on the particulars of how the question is presented. Studies show that adults sometimes use conditional probabilities successfully,⁹ but fail to do so with many problems that call for it.¹⁰ Even trained scientists are open to pitfalls in reasoning about conditional probabilities (as well as other types of reasoning). Physicians are known to discount or misinterpret new patient data that conflict with a diagnosis they have in mind,¹¹ and Ph.D.-level scientists are prey to faulty reasoning when faced with a problem embedded in an unfamiliar context.¹²

And yet, young children are sometimes able to reason about conditional probabilities. In one experiment,¹³ the researchers showed 3-year-olds a box and told them it was a "blicket detector" that would play music if a blicket were placed on top. The child then saw one of the two sequences shown below in which blocks are placed on the blicket detector. At the end of the sequence, the child was asked whether each block was a blicket. In other words, the child was to use conditional reasoning to infer which block caused the music to play.

Note that the relationship between each individual block (yellow cube and blue cylinder) and the music is the same in sequences 1 and 2. In either sequence, the child sees the yellow cube associated with music three times, and the blue cylinder associated with the absence of music once and the presence of music twice. What differs between the first and second sequence is the relationship between the blue and yellow blocks, and therefore, the conditional probability of each block being a blicket. Three-year-olds understood the importance of conditional probabilities.

[†] These two strands are the most often studied, but these two approaches—content and process of science—are incomplete. Underemphasized in U.S. classrooms are the many methods of scientific study, and the role of theories and models in advancing scientific thought.

[‡] Although this is not highly relevant for K-12 teachers, it is important to note that for people with extensive training, such as Ph.D.-level scientists, critical thinking does have some skill-like characteristics. In particular, they are better able to deploy critical reasoning with a wide variety of content, even that with which they are not very familiar. But, of course, this does not mean that they will never make mistakes.

“Teaching content alone is not likely to lead to proficiency in science, nor is engaging in inquiry experiences devoid of meaningful science content.”


—National Research Council

What’s going on? One issue is that the common conception of critical thinking or scientific thinking (or historical thinking) as a set of skills is not accurate. Critical thinking does not have certain characteristics normally associated with skills—in particular, being able to use that skill at any time. If I told you that I learned to read music, for example, you would expect, correctly, that I could use my new skill (i.e., read music) whenever I wanted. But critical thinking is very different. As we saw in the discussion of conditional probabilities, people can engage in some types of critical thinking without training, but even with extensive training, they will sometimes fail to think critically. This understanding that critical thinking is not a skill is vital.[‡] It tells us that teaching students to think critically probably lies in small part in showing them new ways of thinking, and in large part in enabling them to deploy the right type of thinking at the right time.

Returning to our focus on science, we’re ready to address a key question: Can students be taught when to engage in scientific thinking? Sort of. It is easier than trying to teach general critical thinking, but not as easy as we would like. Recall that when we were discussing problem solving, we found that students can learn metacognitive strategies that help them look past the surface structure of a problem and identify its deep structure, thereby getting them a step closer to figuring out a solution. Essentially the same thing can happen with scientific thinking. Students can learn certain metacognitive strategies that will cue them to think scientifically. But, as with problem solving, the metacognitive strategies only tell the students what they should do—they do not provide the knowledge that students need to actually do it. The good news is that within a content area like science, students have more context cues to help them figure out which metacognitive strategy to use, and teachers have a clearer idea of what

For sequence 1, they said the yellow cube was a blicket, but the blue cylinder was not; for sequence 2, they chose equally between the two blocks.

This body of studies has been summarized simply: Children are not as dumb as you might think, and adults (even trained scientists) are not as smart as you might think.



domain knowledge they must teach to enable students to do what the strategy calls for.

For example, two researchers¹⁴ taught second-, third-, and fourth-graders the scientific concept behind controlling variables; that is, of keeping everything in two comparison conditions the same, except for the one variable that is the focus of investigation. The experimenters gave explicit instruction about this strategy for conducting experiments and then had students practice with a set of materials (e.g., springs) to answer a specific question (e.g., which of these factors determine how far a spring will stretch: length, coil diameter, wire diameter, or weight?). The experimenters found that students not only understood the concept of controlling variables, they were able to apply it seven months later with different materials and a different experimenter, although the older children showed more robust transfer than the younger children. In this case, the students recognized that they were designing an experiment and that cued them to recall the metacognitive strategy, “When I design experiments, I should try to control variables.” Of course, succeeding in controlling all of the relevant variables is another matter—that depends on knowing which variables may matter and how they could vary.

Why Scientific Thinking Depends on Scientific Knowledge

Experts in teaching science recommend that scientific

reasoning be taught in the context of rich subject matter knowledge. A committee of prominent science educators brought together by the National Research Council¹⁵ put it plainly: “Teaching content alone is not likely to lead to proficiency in science, nor is engaging in inquiry experiences devoid of meaningful science content.”

The committee drew this conclusion based on evidence that background knowledge is necessary to engage in scientific thinking. For example, knowing that one needs a control group in an experiment is important. Like having two comparison conditions, having a control group in addition to an experimental group helps you focus on the variable you want to study. But knowing that you need a control group is not the same as being able to create one. Since it’s not always possible to have two groups that are *exactly* alike, knowing which factors can vary between groups and which must not vary is one example of necessary background knowledge. In experiments measuring how quickly subjects can respond, for example, control groups must be matched for age, because age affects response speed, but they need not be perfectly matched for gender.

More formal experimental work verifies that background knowledge is necessary to reason scientifically. For example, consider devising a research hypothesis. One could generate multiple hypotheses for any given situation. Suppose you know that car A gets better gas mileage than car

B and you’d like to know why. There are many differences between the cars, so which will you investigate first? Engine size? Tire pressure? A key determinant of the hypothesis you select is plausibility. You won’t choose to investigate a difference between cars A and B that you think is unlikely to contribute to gas mileage (e.g., paint color), but if someone provides a reason to make this factor more plausible (e.g., the way your teenage son’s driving habits changed after he painted his car red), you are more likely to say that this now-plausible factor should be investigated.¹⁶ One’s judgment about the plausibility of a factor being important is based on one’s knowledge of the domain.

Other data indicate that familiarity with the domain makes it easier to juggle different factors simultaneously, which in turn allows you to construct experiments that simultaneously control for more factors. For example, in one experiment,¹⁷ eighth-graders completed two tasks. In one, they were to manipulate conditions in a com-

Did Sherlock Holmes Take a Course in Critical Thinking?

No one better exemplifies the power of broad, deep knowledge in driving critical thinking than Sherlock Holmes. In his famous first encounter with Dr. Watson, Holmes greets him with this observation: “You have been in Afghanistan, I perceive.” Watson is astonished—how could Holmes have known? Eventually Holmes explains his insight, which turns not on incredible intelligence or creativity or wild guessing, but on having relevant knowledge. Holmes is told that Watson is a doctor; everything else he deduces by drawing on his knowledge of, among other things, the military, geography, how injuries heal, and current events. Here’s how Holmes explains his thought process:

I knew you came from Afghanistan. From long habit the train of thoughts ran so swiftly through my mind, that I arrived at the conclusion without being conscious of intermediate steps. There were such steps, however. The train of reasoning ran, “Here is a gentleman of a medical type, but with the air of a military man. Clearly an army doctor, then. He has just come from the tropics, for his face is dark, and that is not the natural tint of his skin, for his wrists are fair. He has undergone hardship and sickness, as his haggard face says clearly. His left arm has been injured. He holds it in a stiff and unnatural manner. Where in the tropics could an English army doctor have seen much hardship and got his arm wounded? Clearly in Afghanistan.” The whole train of thought did not occupy a second. I then remarked that you came from Afghanistan, and you were astonished.

Source: *A Study in Scarlet* by Sir Arthur Conan Doyle.

—EDITORS

Subjects who started with more and better integrated knowledge planned more informative experiments and made better use of experimental outcomes.

puter simulation to keep imaginary creatures alive. In the other, they were told that they had been hired by a swimming pool company to evaluate how the surface area of swimming pools was related to the cooling rate of its water. Students were more adept at designing experiments for the first task than the second, which the researchers interpreted as being due to students' familiarity with the relevant variables. Students are used to thinking about factors that might influence creatures' health (e.g., food, predators), but have less experience working with factors that might influence water temperature (e.g., volume, surface area). Hence, it is not the case that "controlling variables in an experiment" is a pure process that is not affected by subjects' knowledge of those variables.

Prior knowledge and beliefs not only influence which hypotheses one chooses to test, they influence how one interprets data from an experiment. In one experiment,¹⁸ undergraduates were evaluated for their knowledge of electrical circuits. Then they participated in three weekly, 1.5-hour sessions during which they designed and conducted experiments using a computer simulation of circuitry, with the goal of learning how circuitry works. The results showed a strong relationship between subjects' initial knowledge and how much subjects learned in future sessions, in part due to how the subjects interpreted the data from the experiments they had conducted. Subjects who started with more and better integrated knowledge planned more informative experiments and made better use of experimental outcomes.

Other studies have found similar results, and have found that anomalous, or unexpected, outcomes may be particularly important in creating new knowledge—and particularly dependent upon prior knowledge.¹⁹ Data that seem odd because they don't fit one's mental model of the phenomenon under investigation are highly informative.

They tell you that your understanding is incomplete, and they guide the development of new hypotheses. But you could only recognize the outcome of an experiment as anomalous if you had some expectation of how it would turn out. And that expectation would be based on domain knowledge, as would your ability to create a new hypothesis that takes the anomalous outcome into account.

The idea that scientific thinking must be taught hand in hand with scientific content is further supported by research on scientific problem solving; that is, when students calculate an answer to a textbook-like problem, rather than design their own experiment. A meta-analysis²⁰ of 40 experiments investigating methods for teaching scientific problem solving showed that effective approaches were those that focused on building complex, integrated knowledge bases as part of problem solving, for example by including exercises like concept mapping. Ineffective approaches focused exclusively on the strategies to be used in problem solving while ignoring the knowledge necessary for the solution.

What do all these studies boil down to? First, critical thinking (as well as scientific thinking and other domain-based thinking) is not a skill. There is not a set of critical thinking skills that can be acquired and deployed regardless of context. Second, there are metacognitive strategies that, once learned, make critical thinking more likely. Third, the ability to think critically (to actually do what the metacognitive strategies call for) depends on domain knowledge and practice. For teachers, the situation is not hopeless, but no one should underestimate the difficulty of teaching students to think critically. □

Endnotes

¹Borja, R.R. (2006). "Work Skills of Graduates Seen Lacking" *Education Week*, 26, 9, 10.

²Green, W.D. (2007). "Accenture Chairman and CEO William D. Green Addresses Senate Finance Committee," Accenture, www.accenture.com.

³Viadero, D. (1991). "Parents in S.C. Attack Alleged 'New Age' Program." *Education Week*, www.edweek.org.

⁴Novick, L.R. and Holyoak, K.J. (1991). "Mathematical problem-solving by analogy," *Journal of Experimental Psychology: Learning, Memory and Cognition*, 17, 398-415.

⁵For reviews see: Reeves, L.M. and Weisberg, R.W. (1994), "The role of content and abstract information in analogical transfer," *Psychological Bulletin*, 115, 381-400; Barnett, S.M. and Ceci, S.J. (2002), "When and where do we apply what we learn? A taxonomy for far transfer," *Psychological Bulletin*, 128 (4), 612-637.

⁶Chen, Z., Mo, L., and Honovich, R. (2004). "Having the memory of an elephant: Long-term retrieval and the use of analogues in problem solving," *Journal of Experimental Psychology: General*, 133, 415-433.

⁷For a readable review see: Baron, J. (2000). *Thinking and Deciding*, Cambridge, UK: Cambridge University Press.

⁸For example see: Klahr, D. (2000). *Exploring science: The cognition and development of discovery processes*, Cambridge, Mass.: MIT press.

⁹Spellman, B. A. (1996). "Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes,"

Teaching Critical Thinking

Teaching students to think critically is high on any teacher's to-do list. So what strategies are consistent with the research?

■ **Special programs aren't worth it.** In the sidebar on page 12, I've mentioned a few of the better known programs. Despite their widespread availability, the evidence that these programs succeed in teaching students to think critically, especially in novel situations, is very limited. The modest boost that such programs may provide should be viewed, as should all claims of educational effectiveness, in light of their opportunity costs. Every hour students spend on the program is an hour they won't be learning something else.

■ **Thinking critically should be taught in the context of subject matter.** The foregoing does not mean that teachers shouldn't teach students to think critically—it means that critical thinking shouldn't be taught on its own. People do not spontaneously examine assumptions that underlie their thinking, try to consider all sides of an issue, question what they know, etc. These things must be modeled for students, and students must be given opportunities to practice—preferably in the context of normal classroom activity. This is true not only for science (as discussed in the main article), but for other subject matter. For example, an important part of thinking like a historian is considering the source of a document—who wrote it, when, and why. But teaching students to ask that question, independent of subject matter knowledge, won't do much good. Knowing that a letter was written by a Confederate private to his wife in New Orleans just after the Battle of Vicksburg won't help the student interpret the letter unless he knows something of Civil War history.

■ **Critical thinking is not just for advanced students.** I have sometimes heard teachers and administrators suggest that critical thinking exercises make a good enrichment activity for the best students, but struggling students should just be expected to understand and master more basic material. This argument sells short the less advanced students and conflicts with what cognitive scientists know about thinking. Virtually everyone is capable of critical thinking and uses it all the time—and, as the conditional probabilities research demonstrated (see p. 15), has been capable of doing so since they were very young. The difficulty lies not in thinking critically, but in recognizing when to do so, and in knowing enough to do so successfully.

■ **Student experiences offer entrée to complex concepts.** Although critical thinking needs to be nested in subject matter, when students don't have much subject matter knowledge, introducing a concept by drawing on student experiences can help. For example, the importance of a source in evaluating a historical document is familiar to even young children; deepening their understanding is a matter of asking questions that they have the knowledge to grapple with. Elementary school teachers could ask: Would a letter to a newspaper editor that criticized the abolishment of recess be viewed differently if written by a school principal versus a third-grader? Various concepts that are central to scientific thinking can also be taught with examples that draw on students' everyday knowledge and experience. For example, "correlation does not imply causation" is often illustrated by the robust association between the consumption of ice cream and the number of crimes committed on a given day. With a little prodding, students soon realize that ice cream consumption doesn't

Knowing that a letter was written by a Confederate private to his wife in New Orleans just after the Battle of Vicksburg won't help the student interpret the letter—unless he knows something of Civil War history.

cause crime, but high temperatures might cause increases in both.

■ **To teach critical thinking strategies, make them explicit and practice them.** Critical thinking strategies are abstractions. A plausible approach to teaching them is to make them explicit, and to proceed in stages. The first time (or several times) the concept is introduced, explain it with at least two different examples (possibly examples based on students' experiences, as discussed above), label it so as to identify it as a strategy that can be applied in various contexts, and show how it applies to the course content at hand. In future instances, try naming the appropriate critical thinking strategy to see if students remember it and can figure out how it applies to the material under discussion. With still more practice, students may see which strategy applies without a cue from you.

—D.W.

- ¹⁰For example see: Kuhn, D., Garcia-Mila, M., and Zohar, A. (1995). "Strategies of knowledge acquisition," *Monographs of the Society for Research in Child Development*, 60, 1-128.
- ¹¹Groopman, J. (2007). *How Doctors Think*, New York: Houghton Mifflin.
- ¹²Tweney, R. D. and Yachanin, S. A. (1985), "Can scientists rationally assess conditional inferences?" *Social Studies of Science*, 15, 155-173; Mahoney, M.J. and DeMonbreun, B.G. (1981), "Problem-solving bias in scientists," in R. D. Tweney, M. E. Doherty, and C. R. Mynatt (eds.) *On Scientific Thinking*, 139-144, New York: Columbia University Press.
- ¹³Gopnik, A., Sobel, D.M., Schulz, L.E., and Glymour, C. (2001). "Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation," *Developmental Psychology*, 37(5), 620-629.
- ¹⁴Chen, Z. and Klahr, D. (1999). "All Other Things Being Equal: Acquisition and Transfer of the Control of Variables Strategy," *Child Development*, 70 (5), 1098-1120.
- ¹⁵National Research Council (2007). *Taking Science to School*, Washington, D.C.: National Academies Press.
- ¹⁶Koslowski, B. (1996). *Theory and Evidence: The Development of Scientific Reasoning*, Cambridge, Mass.: MIT Press.
- ¹⁷Friedler, Y., Nachmias, R., and Linn, M. C. (1990). "Learning scientific reasoning skills in microcomputer-based laboratories," *Journal of Research in Science Teaching*, 27, 173-191.
- ¹⁸Schauble, L., Glaser, R., Raghavan, K., and Reiner, M. (1991). "Causal models and experimentation strategies in scientific reasoning," *The Journal of Learning Sciences*, 1, 201-238.
- ¹⁹For example see: Dunbar, K. N. and Fugelsang, J. A. (2005), "Causal thinking in science: How scientists and students interpret the unexpected," in M. E. Gorman, R. D. Tweney, D. C. Gooding, and A. P. Kincannon (eds.) *Scientific and Technological Thinking*, 57-79, Mahwah, N.J.: Erlbaum; Echevarria, M. (2003), "Anomalies as a catalyst for middle school students' knowledge construction and scientific reasoning during science inquiry," *Journal of Educational Psychology*, 95, 357-374.
- ²⁰Taconis, R., Ferguson-Hessler, M.G.M., and Broekkamp, H., (2001). "Teaching science problem solving: An overview of experimental work," *Journal of Research in Science Teaching*, 38(4), 442-468.

Sidebar Endnotes (p. 12)

- ¹Adams, M. J. (1989), "Thinking skills curricula: Their promise and progress," *Educational Psychologist*, 24, 25-77; Nickerson, R. S., Perkins, D. N., and Smith, E. E. (1985), *The Teaching of Thinking*, Hillsdale, N.J.: Erlbaum; Ritchart, R. and Perkins, D. N. (2005), "Learning to think: The challenges of teaching thinking," in K. J. Holyoak and R. G. Morrison (eds.) *The Cambridge Handbook of Thinking and Reasoning*, Cambridge, UK: Cambridge University Press.
- ²Sternberg, R. J. and Bhana, K. (1986). "Synthesis of research on the effectiveness of intellectual skills programs: Snake-oil remedies or miracle cures?" *Educational Leadership*, 44, 60-67.
- ³Mahoney, M.J. and DeMonbreun, B.G. (1981). "Problem-solving bias in scientists," in R. D. Tweney, M. E. Doherty, and C. R. Mynatt (eds.) *On Scientific Thinking*, 139-144, New York: Columbia University Press.
- ⁴Trickey, S. and Topping, K. J. (2004). "Philosophy for Children: A Systematic Review," *Research Papers in Education* 19, 365-380.
- ⁵Adams, M. J. (1989). "Thinking skills curricula: Their promise and progress," *Educational Psychologist*, 24, 25-77; Nickerson, R. S., Perkins, D. N., and Smith, E. E. (1985), *The Teaching of Thinking*, Hillsdale, N.J.: Erlbaum; Ritchart, R. and Perkins, D. N. (2005), "Learning to think: The challenges of teaching thinking," in K. J. Holyoak and R. G. Morrison (eds.) *The Cambridge Handbook of Thinking and Reasoning*, Cambridge, UK: Cambridge University Press.