Confessions of an Accidental Theorist

ALAN H. SCHOENFELD

David Wheeler had both theoretical and pragmatic reasons for inviting me to write this article. On the theoretical side, he noted that my ideas on "understanding and teaching the nature of mathematical thinking" have taken some curious twists and turns over the past decade. Originally inspired by Polya's ideas and intrigued by the potential for implementing them using the tools of artificial intelligence and information-processing psychology, I set out to develop prescriptive models of heuristic problem solving — models that included descriptions of how, and when, to use Polya's strategies. (In moments of verbal excess I was heard to say that my research plan was to "understand how competent problem solvers solve problems, and then find a way to cram that knowledge down students' throats.") Catch me talking today, and you'll hear me throwing about terms like metacognition, belief systems, and "culture as the growth medium for cognition"; there's little or no mention of prescriptive models. What happened in between? How were various ideas conceived, developed, modified, adapted, abandoned, and sometimes reborn? It might be of interest, suggested David, to see where the ideas came from. With regard to pragmatic issues, David was blunt. Over the past decade I've said a lot of stupid things. To help keep others from re-inventing square theoretical or pedagogical wheels, or to keep people from trying to ride some of the square wheels I've developed, he suggested, it might help if I recanted in public. So here goes ...

The story begins in 1974, when I tripped over Polya's marvelous little volume *How to solve it*. The book was a *tour de force*, a charming exposition of the problem solving introspections of one of the century's foremost mathematicians. (If you don't own a copy, you should.) In the spirit of Descartes, who had, three hundred years earlier, attempted a similar feat in the *Rules for the direction of the mind*, Polya examined his own thoughts to find useful patterns of problem solving behavior. The result was a general description of problem solving processes: a four-phase model of problem solving (understanding the problem, devising a plan, carrying out the plan, looking back), the details of which included a range of problem solving *heuristics*, or rules of thumb for making progress on difficult problems. The book and Polya's subsequent elaborations of the heuristic theme (in *Mathematics and plausible reasoning*, and *Mathematical discovery*) are brilliant pieces of insight and mathematical exposition.

A young mathematician only a few years out of graduate school, I was completely bowled over by the book. Page after page, Polya described the problem solving techniques that he used. Though I hadn't been taught them, I too used those techniques; I'd picked them up then pretty much by accident, by virtue of having solved thousands of problems during my mathematical career (that is, I'd been "trained" by the discipline, picking up bits and pieces of mathematical thinking as I developed). My experience was hardly unique, of course. In my excitement I joined thousands of mathematicians who, in reading Polya's works, had the same thrill of recognition. In spirit I enlisted in the army of teachers who, inspired by Polya's vision, decided to focus on teaching their students to think mathematically instead of focusing merely on the mastery of mathematical subject matter.

To be more accurate, I thought about enlisting in that army. Excited by my readings, I sought out some problem-solving experts, mathematics faculty who coached students for the Putnam exam or for various Olympiads. Their verdict was unanimous and unequivocal: Polya was of no use for budding young problem-solvers. Students don't learn to solve problems by reading Polya's books, they said. In their experience, students learned to solve problems by (starting with raw talent and) solving lots of problems. This was troubling, so I looked elsewhere for (either positive or negative) evidence. As noted above, I was hardly the first Polya enthusiast: By the time I read *How to solve it* the math-ed literature was chock full of studies designed to teach problem-solving *via* heuristics. Unfortunately, the results — whether in first grade, algebra, calculus, or number theory, to name a few — were all depressingly the same, and confirmed the statements of the Putnam and Olympiad trainers. Study after study produced "promising" results, where teacher and students alike were happy with the instruction (a typical phenomenon when teachers have a vested interest in a new program) but where there was at best marginal evidence (if any!) of improved problem solving performance. Despite all the enthusiasm for the approach, there was no clear evidence that the students had actually learned more as a result of their heuristic instruction, or that they had learned any general problem solving skills that transferred to novel situations.

Intrigued by the contradiction — my gut reaction was still that Polya was on to something significant — I decided to trade in my mathematician's cap for a mathematics

For the Learning of Mathematics 7, 1 (February 1987)
FLM Publishing Association, Montreal, Quebec, Canada

educator's and explore the issue Well, not exactly a straight mathematics educator's; as I said above, math ed had not produced much that was encouraging on the problem solving front. I turned to a different field, in the hope of blending its insights with Polya's and those of mathematics educators.

The first task I faced was to figure out why Polya's strategies didn't work If I succeeded in that, the next task was to make them work — to characterize the strategies so that students could learn to use them. The approach I took was inspired by classic problem solving work in cognitive science and artificial intelligence, typified by Newell and Simon's [1972] *Human problem solving*. In the book Newell and Simon describe the genesis of a computer program called General Problem Solver (GPS), which was developed to solve problems in symbolic logic, chess, and "cryptarithmetic" (a puzzle domain similar to cryptograms, but with letters standing for numbers instead of letters). GPS played a decent game of chess, solved cryptarithmetic problems fairly well, and managed to prove almost all of the first 50 theorems in Russell and Whitehead's *Principia Mathematica* — all in all, rather convincing evidence that its problem solving strategies were pretty solid.

Where did those strategies come from? In short, they came from detailed observations of people solving problems. Newell, Simon, and colleagues recorded many people's attempts to solve problems in chess, cryptarithmetic, and symbolic logic. They then explored those attempts in detail, looking for uniformities in the problem solver's behavior. If they could find those regularities in people's behavior, describe those regularities precisely (i.e. as computer programs), and get the programs to work (i.e. to solve problems) then they had pretty good evidence that the strategies they had characterized were useful. As noted above, they succeeded. Similar techniques had been used in other areas: for example, a rather simple program called SAINT (for Symbolic Automatic INTegrator) solved indefinite integrals with better facility than most M.I.T. freshment. In all such cases, AI produced a set of *prescriptive* procedures - problem solving methods described in such detail that a machine, following their instructions, could obtain pretty spectacular results.

It is ironic that no one had thought to do something similar for human problem solving. The point is that one could turn the man-machine metaphor back on itself. Why not make detailed observations of expert human problem solvers, with an eye towards abstracting regularities in their behavior — regularities that could be codified as prescriptive guides to human problem solving? No slight to students was intended by this approach, nor was there any thought of students as problem solving machines Rather, the ideas was to pose the problem from a cognitive science perspective: "What level of detail is needed so that students can actually use the strategies one believes to be useful?" Methodologies for dealing with this question were suggested by the methodologies used in artificial intelligence. One could make detailed observations of individuals solving problems, seek regularities in their problem solving behavior, and try to characterize those regularities with enough precision, and in enough detail, so that students could take those characterizations as guidelines for problem solving. That's what I set out to do.

The detailed studies of problem solving behavior turned up some results pretty fast. In particular, they quickly revealed one reason that attempts to teach problem solving *via* heuristics had failed The reason is that Polya's heuristic strategies weren't really coherent strategies at all. Polya's characterizations were broad and descriptive, rather than prescriptive. Professional mathematicians could indeed recognize them (because they knew them, albeit implicitly), but novice problem solvers could hardly use them as guides to productive problem solving behavior. In short, Polya's characterizations were *labels* under which families of related strategies were subsumed. There isn't much room for exposition here, but one example will give the flavor of the analysis. The basic idea is that when you look closely at any single heuristic "strategy," it explodes into a dozen or more similar, but fundamentally different, problem-solving techniques. Consider a typical strategy, "examining special cases:"

> To better understand an unfamiliar problem, you may wish to exemplify the problem by considering various special cases. This may suggest the direction of, perhaps the plausbility of, a solution.

Now consider the solutions to the following three problems.

> *Problem 1*. Determine a formula in closed form for the series
>
> $$\sum_{i=1}^{n} k/(k+1)!$$
>
> *Problem 2*. Let $P(x)$ and $Q(x)$ by polynomials whose coefficients are the same but in "backwards order":
>
> $$P(x) = a_0 + a_1x + a_2x^2 + \ldots a_nx^n \text{ and}$$
> $$Q(x) = a_n + a_{n-1}x + a_{n-2}x^2 + \ldots a_0x^n$$
>
> What is the relationship between the roots of $P(x)$ and Q(x)? Prove your answer.
>
> *Problem 3*. Let the real numbers a_0 and a_1 be given. Define the sequence $\{a_n\}$ by
>
> $$a_n = 1/2\ (a_{n-2} + a_{n-1}) \text{ for each } n \geq 2.$$
>
> Does the sequence $\{a_n\}$ converge? If so, to what value?

I'll leave the details of the solutions to you. However, the following observations are important. For problem 1, the special cases that help are examining what happens when there the integer parameter n takes on the values 1, 2, 3, ... in sequence; this suggests a general pattern that can be confirmed by induction. Yet if you try to use special cases in the same way on the second problem, you may get into trouble: Looking at values $n=1, 2, 3, \ldots$ can lead to a wild goose chase. It turns out that the right special cases of $P(x)$ and $Q(x)$ you need to look at for problem 2 are easily factorable polynomials. If, for example, you consider

$$P(x) = (2x + 1)(x + 4)(3x - 2),$$

you will discover that its "reverse," Q, is easily factorable. The roots of the P and Q are easy to compare, and the result (which is best proved another way) is obvious. And again, the special cases that simplify the third problem are different in nature. If you choose the values $a_0=0$ and $a_1=1$, you can see what happens for that particular sequence. The pattern in that case suggests what happens in general, and (especially if you draw the right picture!) leads to a solution of the original problem.

Each of these problems typifies a large class of problems, and exemplifies a different special cases strategy. We have:

> Strategy 1. When dealing with problems in which an integer parameter n plays a prominent role, it may be of use to examine values of $n=1, 2, 3, \ldots$ in sequence, in search of a pattern.
>
> Strategy 2. When dealing with problems that concern the roots of polynomials, it may be of use to look at easily factorable polynomials.
>
> Strategy 3. When dealing with problems that concern sequences or series that are constructed recursively, it may be of use to try initial values of 0 and 1 — if such choices don't destroy the generality of the processes under investigation.

Needless to say, these three strategies hardly exhaust "special cases." At this level of analysis — the level of analysis necessary for implementing the strategies — one could find a dozen more. This is the case for almost all of Polya's strategies. In consequence the two dozen or so "powerful strategies" in *How to solve it* are, in actuality, a collection of two or three hundred less "powerful", but actually usable strategies. The task of teaching problem solving *via* heuristics — my original goal — thus expanded to (1) explicitly identifying the most frequently used techniques from this long list, (2) characterizing them in sufficient detail so that students could use them, and (3) providing the appropriate amount and degree of training.

(*Warning:* It is easy to underestimate both the amount of detail and training that are necessary. For example, to execute a moderately complex "strategy" like "exploit an easier related problem" with success, you have to (a) think to use the strategy (non-trivial!); (b) know which version of the strategy to use; (c) generate *appropriate* and *potentially useful* easier related problems; (d) make the right choice of related problem; (e) solve the problem; and (f) find a way to exploit its solution to help solve the original problem. Students need instruction in all of these.)

Well, this approach made progress, but it wasn't good enough. Fleshing out Polya's strategies did make them implementable, but it revealed a new problem. An arsenal of a dozen or so powerful techniques may be manageable in problem solving. But with all the new detail, our arsenal comprised a couple of hundred problem solving techniques. This caused a new problem, which I'll introduce with an analogy.

A number of years ago, I deliberately put the problem

$$\int \frac{x\ dx}{x^2 - 9}$$

as the first problem on a test, to give my students a boost as they began the exam. After all, a quick look at the fraction suggests the substitution $u = x^2 - 9$, and this substitution knocks the problem off in just a few seconds. 178 students took the exam. About half used the right substitution and got off to a good start, as I intended. However, 44 of the students, noting the factorable denominator in the integrand, used partial fractions to express x/x^2-9 in the form $A/(x-3) + B/(x+3)$ — correct but quite time-consuming. They didn't do too well on the exam. And 17 students, noting the $(u^2 - a^2)$ form of the denominator, worked the problem using the substitution $x = 3\sin\theta$. This too yields the right answer — but it was even more time-consuming, and the students wound up so far behind that they bombed the exam.

Doing well, then, is based on more than "knowing the subject matter"; it's based on knowing which techniques to use and when. If your strategy choice isn't good, you're in trouble. That's the case in techniques of integration, when there are only a dozen techniques and they're all algorithmic. As we've seen, heuristic techniques are anything but algorithmic, and they're much harder to master. In addition, there are hundreds of them — so strategy selection becomes even more important a factor in success. My point was this. Knowing the strategies isn't enough. You've got to know when to use which strategies.

As you might expect by now, the AI metaphor provided the basic approach. I observed the good problem solvers with an eye towards replicating their heuristic strategy selection. Generalizing what they did, I came up with a prescriptive scheme for picking heuristics, called a "managerial strategy." It told the student which strategies to use, and when (unless the student was sure he had a better idea). Now again, this approach is not quite as silly as it sounds. Indeed, the seeds of it are in Polya ("First. You have to understand the problem.") The students weren't forced to follow the managerial strategy like little automata. But the strategy suggested that heuristic techniques for understanding the problem should be used first, planning heuristics next, exploration heuristics in a particular order (the metric was that the further the exploration took you from the original problem the later you should consider it), and so on. In class we talked about which heuristic technique we might use at any time, and why. Was the approach reductive? Maybe so. But the bottom line is that this combination of making the heuristics explicit, and providing a managerial strategy for students, was gloriously successful

The final examinations for my problem solving courses had three parts. Part 1 had problems similar to problems we had worked in the course Part 2 had problems that could be solved by the methods we had studied, but the problems did not resemble ones we had worked. Part 3 consisted of problems that had stumped me. I had looked through contest problem books, and as soon as I found a problem that baffled me, I put it on the exam! The students did quite well even on part 3; some solved problems on which I had not made progress, in the same amount of time

Thus ended Phase 1 of my work At that point — the late 1970's to 1980 — I was pretty happy with the instruction, and was getting pretty good results Then something happened that shook me up quite a bit. Thanks to a National Science Foundation grant I got a videotape machine, and actually looked at students' problem solving behavior. What I saw was frightening.

Even discounting possible hyperbole in the last sentence, one statement in the previous paragraph sounds pretty strange I'd been teaching for more than a decade and doing research on problem solving for about half that time. How can I suggest that, with all of that experience, I had never really looked at students' problem solving behavior?

With the videotape equipment, I brought students into my office, gave them problems (before, after, and completely independently of my problem solving courses), and had them work on the problems at length. Then, at leisure, I looked at the videotapes and examined, in detail, what the students actually did while they worked on the problems. What I saw was nothing like what I expected, and nothing like what I saw as a teacher That's because as teachers (and often as researchers) we look at a very narrow spectrum of student behavior Generally speaking, we only see what students produce on tests; that's the product, but focusing on the product leaves the process by which it evolved largely invisible. (There's a substantial difference between watching a 20-minute videotape of a student working a problem and reading the page or two of "solution" that student produced in those 20 minutes The difference can be mind boggling.) In class, or in office hours, we have conversations with the students, but the conversations are directed toward a goal — explaining something the student comes prepared to understand, and knows is coming The student is primed for what we have to say. And that's the point. When we give students a calculus test and there's a max-min problem in it, students know it's a max-min problem They've just finished a unit on max-min problem, and they expect to see a max-min problems on the exam. In other words, the context tells the students what mathematics to use. We get to see them at their very best, because (a) they're prepared, and (b) the general context puts them in the right ballpark and tells them what procedures to use. By way of analogy, you don't discover whether kids can speak grammatically (or think on their feet) when you given them a spelling test, after they've been given the list of words they'll be tested on. (Even when I taught the problem solving class, I was showing students techniques that they knew were to be used in the context of the problem solving class Hence they came to my final prepared to use those techniques)

In my office, problems come out of the blue and the context doesn't tell students what methods are appropriate. The result is that I get to see a very different kind of behavior. One problem used in my research, for example, is the following:

> *Problem 4.* Three points are chosen on the circumference of a circle of radius R, and the triangle with those points as vertices is drawn. What choice of points results in the triangle with largest possible area? Justify your answer as well as you can.

Though there are clever solutions to this problem (see below), the fact is that you can approach it as a standard multivariate max-min problem Virtually none of my students (who had finished 3rd-semester calculus, and who knew more than enough mathematics to knock the problem off) approached it that way One particular pair of students had just gotten A's in their 3rd-semester calculus class, and each had gotten full credit on a comparably difficult problem on their exam. Yet when they worked on this problem they jumped into another (and to me, clearly irrelevant) approach altogether, and persisted at it for the full amount of allotted time. When they ran out of time, I asked them where they were going with that approach and how it might help them. They couldn't tell me. That solution attempt is best described as a twenty-minute wild goose chase.

Most of my videotapes showed students working on problems that they "knew" enough mathematics to solve. Yet time and time again, students never got to use their knowledge They read the problem, picked a direction (often in just a second or two), and persevered in that direction no matter what Almost sixty percent of my tapes are of that nature. But perhaps the most embarrassing of the tapes is one on which I recorded a student who had taken my problem solving course the year before.

There is an elegant solution to Problem 4, which goes as follows. Suppose the three vertices are A, B, and C Hold A and B fixed, and ask what choice of C gives the largest area. It's clearly when the height of the traingle is maximized — when the triangle is isosceles. So the largest triangle must be isosceles Now you can either maximize isosceles triangles (a one-variable calculus problem), or finish the argument by contradiction Suppose the largest triangle, ABC, isn't equilateral. Then two sides are unequal; say $AC \neq BC$. If that's the case, however, the isosceles triangle with base AB is larger than ABC — a contradiction. So ABC must be equilateral.

The student sat down to work the problem. He remembered that we'd worked it in class the previous year, and that there was a elegant solution. As a result, he approached the problem by trying to do something clever. In an attempt to exploit symmetry he changed the problem he was working on (without acknowledging that this might have serious consequences). Then, pursuing the goal of a slick solution he missed leads that clearly pointed to a straightforward solution. He also gave up potentially fruitful approaches that were cumbersome because "there must be an easier way". In short, a cynic would argue that he was worse off after my course then before. (That's how I felt that afternoon.)

In any case, I drew two morals from this kind of experience. The first is that my course, broad as it was, suffered from the kind of insularity I discussed above. Despite the fact that I was teaching "general problem solving strategies", I was getting good results partly because I had narrowed the context: students knew they were supposed to be using the strategies in class, and on my tests. If I wanted to affect the students' behavior in a lasting way, outside of my classroom, I would have to do something different. (Note: I had plenty of testimonials from students that my course had "made me a much better problem solver", "helped me do much better in all of my other courses", and "changed my life". I'm not really sanguine about any of that). Second and more important, I realized that there was a fundamental mistake in the approach I had taken to teaching problem solving — the idea that I could, as I put it so indelicately in the first paragraph of this paper, cram problem solving knowledge down my students' throats.

That kind of approach makes a naive and very dangerous assumption about students and learning. It assumes, in essence, that each student comes to you as a *tabula rasa*, that you can write your problem solving "message" upon that blank slate, and that the message will "take." And it just ain't so. The students in my problem solving classes were the successes of our system. They were at Hamilton College, at Rochester, or at Berkeley because they were good students; they were in a problem solving class (which was known as a killer) because they liked mathematics and did pretty well at it. They came to the class with well engrained habits — the very habits that have gotten them to the class in the first place, and accounted for their success. I ignore all of that (well, not really; but a brief caricature is all I've got room for) and show them "how to do it right". And when they leave the classroom and are on their own ... well, let's be realistic. How could a semester's worth of training stack up against an academic lifetime's worth of experience, especially if the course ignores that experience? (Think of what it takes to retrain a self-taught musician or tennis player, rather than teach one from scratch. Old habits die very very hard, if they die at all.)

Well, the point is clear. If you're going to try to affect students' mathematical problem solving behavior, you'd better understand that behavior. That effort was the main thrust of what (linear type that I am) I'll call phase 2. Instead of trying to do things to (and with) students, the ideas was to understand what went on in their heads when they tried to do mathematics. Roughly speaking, the idea was this. Suppose I ask someone to solve some mathematics problems for me. For the sake of a permanent record, I videotape the problem solving session (and the person talks out loud as he or she works, giving me a verbal "trace" as well). My goal is to understand what the person did, why he or she did it, and how those actions contributed to his or her success or failure at solving the problem. Along the way I'm at liberty to ask any questions I want, give any tests that seem relevant, and perform any (reasonable) experiments. What do I have to look at, to be reasonably confident that I've focused on the main determinant of behavior, and on what caused success or failure?

The details of my answer are xvi+409 pages long. The masochistic reader may find them, as well as the details of the brief anecdotes sketched above, in my [1985] *Mathematical problem solving*. In brief, the book suggested that if you're going to try to make sense of what people do when they do mathematics, you'd better look at:

A. "Cognitive resources," one's basic knowledge of mathematical facts and procedures stored in LTM (long term memory). Most of modern psychology, which studies what's in a person's head and how that knowledge is accessed, is relevant here.

B. Problem solving strategies or heuristics. I've said enough about these.

C. Executive or "control" behavior. (For the record, this behavior is often referred to as "metacognition.) I discussed this above as well. It's not just what you know (A+B above), it's how you use it. The issue in the book was how to make sense of such things. It's tricky, for the most important thing in a problem solving session may be something that doesn't take place — asking yourself if it's really reasonable to do something, and thereby forestalling a wild good chase.

D. Belief systems. I haven't mentioned these yet, but I will now

Beliefs have to do with your mathematical *weltanschauung*, or world view. The idea is that your sense of what mathematics is all about will determine how you approach mathematical problems. At the joint CMS/-CMESG meetings in June 1986, Ed Williams told me a story that illustrates this category. Williams was one of the organizers of a problem solving contest which contained the following problem:

> "Which fits better, a square peg in a round hole or a round peg in a square hole?"

Since the peg-to-hole ratio is $2/\pi$ (about 0.64) in the former case and $\pi/4$ (about 0.79) in the latter, the answer is "the round peg". (Since the tangents line up in that case and not in the other, there's double reason to choose that answer.) It seems obvious that you have to answer the question by invoking a computation. How else, except with analytic support, can you defend your claim?

It may be obvious to us that an analytic answer is called for, but it's not at all obvious to students. More than 300 students — the cream of the crop — worked the problem. Most got the right answer, justifying it on the basis of a rough sketch. *Only four students* out of more than 300 justified their answer by comparing areas. (I can imagine a student saying "you just said to say which fit better. You didn't say to prove it.") Why? I'm sure the students could have done the calculations. They didn't think to, because they didn't realize that justifying one's answer is a necessary part of doing mathematics (from the mathematician's point of view).

For the sake of argument, I'm going to state the students' point of view (as described in the previous paragraph) in more provocative form, as a belief:

> *Belief 1:* If you're asked your opinion about a mathematical question, it suffices to give your opinion, although you might back it up with evidence if that evidence is readily available. Formal proofs or justifications aren't necessary, unless you're specifically asked for them — and that's only because you have to play the rules of the game.

We've seen the behavioral corollary of this belief, as Williams described it. Unfortunately, this belief has lots of company. Here are two of its friends, and their behavioral corollaries.

> *Belief 2:* All mathematics problems can be solved in ten minutes or less, if you understand the material. *Corollary:* students give up after ten minutes of work on a "problem."

> *Belief 3:* Only geniuses are capable of discovering, creating, and understanding mathematics. *Corollary.* students expect to take their mathematics passively, memorizing without hope or expectation of understanding.

An anecdote introduces one last belief. A while ago I gave a talk describing my research on problem solving to a group of very talented undergraduate science majors at Rochester. I asked the students to solve Problem 5, given in Figure 1. The students, working as a group, generated a correct proof; I wrote the proof (Figure 2) on the board. A few minutes later I gave the students Problem 6, given in Figure 3.

In the figure below, the circle with center C is tangent to the top and bottom lines at the points P and Q respectively.

a. *Prove that* PV = QV.

b. *Prove that the line segment* CV *bisects angle* PVQ.

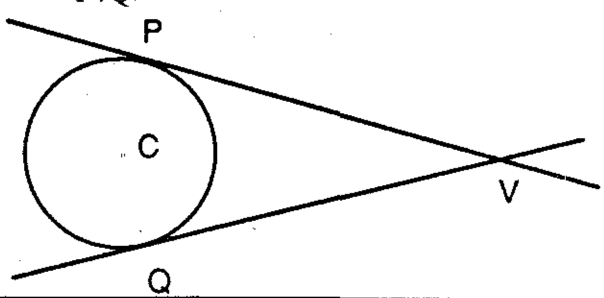

Figure 1

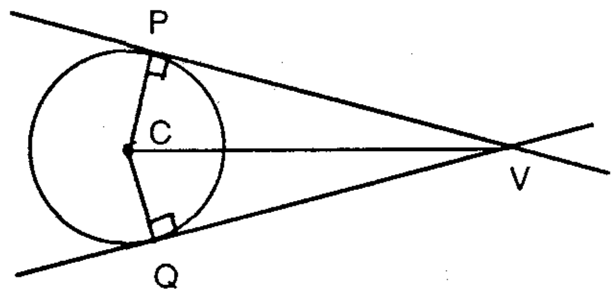

Proof:

Draw in the line segments CP, CQ, and CV. Since CP and CQ are radii of circle C, they are equal; since P and Q are points of tangency, angles CPV and CQV are right angles. Finally since CV = CV, triangles CPV and CQV are congruent.

a. Corresponding parts of congruent triangles are congruent, so PV = QV.

b. Corresponding parts of congruent triangles are congruent, so angle PVC = angle QVC. Thus CV bisects angle PVQ.

Figure 2

You are given two intersecting straight lines and a point P *marked on one of them, as in the figure below. Show how to construct, using straightedge and compass, a circle that is tangent to both lines and that has the point* P *as its point of tangency to the top line.*

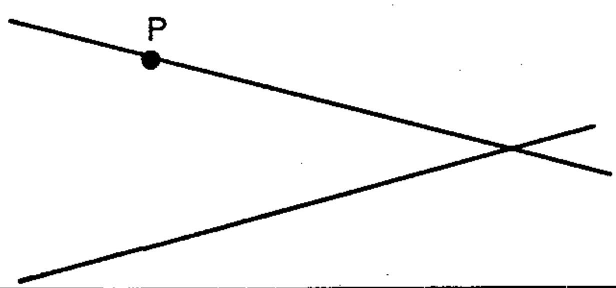

Figure 3

Students came to the board and made the following conjectures, in order:

a. Let *Q* be the point on the bottom line such that *QV* = *PV*. The center of the desired circle is the midpoint of line segment *PQ*. (Figure 4a)

b. Let *A* be the segment of the arc with vertex *V*, passing through *P*, and bounded by the two lines. The center of the described circle is the midpoint of the arc *A* (Figure 4b)

c. Let R be the point on the bottom line that intersects the line segment perpendicular to the top line at P. The ceter of the desired circle is the midpoint of line segment PR (Figure 4c)

d. Let L_1 be the line segment perpendicular to the top line at P, and L_2 the bisector of the angle at V. The center of the desired circle is the point of intersection of L_1 and L_2 (Figure 4d)

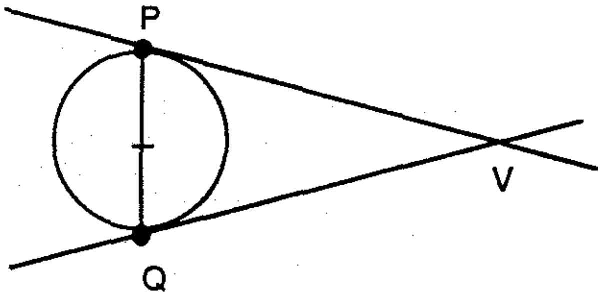

(a)

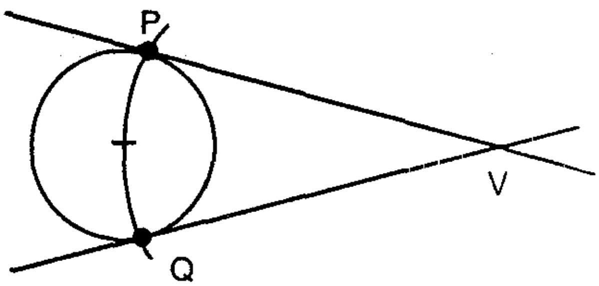

(b)

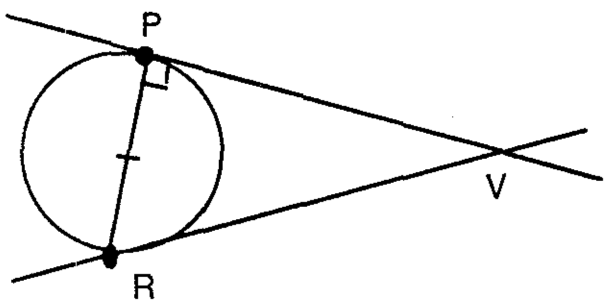

(c)

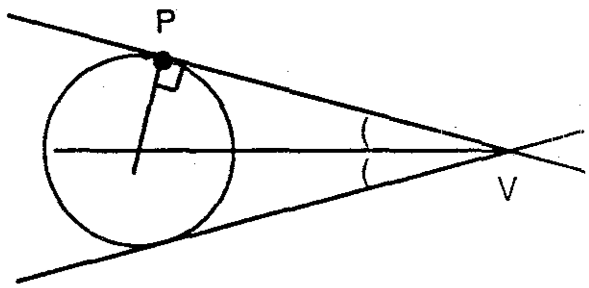

(d)

Students' conjectured solutions
(Short horizontal lines denote midpoints.)

Figure 4

The proof that the students had generated — which both provides the answer and rules out conjectures a, b, and c — was still on the board. Despite this, they argued for more than ten minutes about which construction was right. The argument was on purely empirical grounds (that is, on the grounds of which construction looked right), and it was not resolves. How could they have this argument, with the proof still on the board? I believe that this scene could only take place if the students simply didn't see the proof problem as being relevant to the construction problem. Or again in provocative form,

> *Belief 4*: Formal mathematics, and proof, have nothing to do with discovery or invention. *Corollary*: the results of formal mathematics are ignored when students work discovery problems.

Since we're in "brief survey mode," I don't want to spend too much time on beliefs *per se*. I think the point is clear. If you want to understand students' mathematical behavior, you have to know more than what they "know" These students "knew" plane geometry, and how to write proofs; yet they ignored that knowledge when working construction problems. Understanding what went on in their heads was (and is) tricky business. As I said, that was the main thrust of phase 2.

But enough of that; we're confronted with a real dilemma. The behavior I just described turns out to be almost universal. Undergraduates at Hamilton College, Rochester, and Berkeley all have much the same mathematical world view, and the (U.S.) National Assessments of Educational Progress indicate that the same holds for highs school students around the country. How in the world did those students develop their bizarre sense of what mathematics is all about?

The answer, of course, lies in the students' histories. Beliefs about mathematics, like beliefs about anything else — race, sex, and politics, to name a few — are shaped by one's environment. You develop your sense of what something is all about (be that something mathematics, race, sex, or politics) by virtue of your experiences with it, within the context of your social environment. You may pick up your culture's values, or revel against them — but you're shaped by them just the same.

Mathematics is a formal discipline, to which you're exposed mostly in schools. So if you want to see where kids' views about mathematics are shaped, the first place to go is into mathematics classrooms. I packed up my videotape equipment, and off I went. Some of the details of what I saw, and how I interpreted it, are given in the in-press articles cited in the references. A thumbnail sketch of some of the ideas follows.

Borrowing a term from anthropologists, what I observed in mathematics classes was the *practice* of schooling — the day-to-day rituals and interactions that take place in mathematics classes, and (*de facto*) define what it is to do mathematics. One set of practices has to do with homework and testing. The name of the game in school mathematics is "mastery:" Students are supposed to get their facts and procedures down cold. That means that most homework problems are trivial variants of things the students have already learned. For example, one "required" construction in plane geometery (which students memorize) it to construct a line through a given point, parallel to a given line. A homework assignment given a few days later contained the following problem: Given a point on a scale of a triangle, construct a line through that point parallel to the base of the triangle. This isn't a problem; it's an exercise. It was one of 27 "problems" given that night; the three previous assignments had contained 28, 45, and 18 problems respectively. The test on locus and constructions contained 25 problems, and the students were expected to finish (and check!) the test in 54 minutes — an average of two minutes and ten seconds per problem. Is it any wonder that students come to believe that any problem can be solved in ten minutes or less?

I also note that the teacher was quite explicit about how the students should prepare for the test. His advice — well intentioned — to the students when they asked about the exam was as follows: "You'll have to know all your constructions cold so that you don't spend a lot of time thinking about them." In fact, he's right. Certain skills should be automatic, and you shouldn't have to think about them. But when this is the primary if not the only message that students get, they abstract it as a belief: mathematics is mostly, if not all, memorizing.

Other aspects of what I'll call the "culture of schooling" shape students' view of what mathematics is all about. Though there is now a small movement toward group problem solving in the schools, mathematics for the most part is a solitary endeavor, with individual students working alone at their desks. The message they get is that mathematics is a solitary activity.

They also get a variety of messages about the nature of the mathematics itself. Many word problems in school tell a story that requires a straightforward calculation (for example, "John had twenty-eight candy bars in seven boxes. If each box contained the same number of candy bars, how many candy bars are there in each box?"). The students learn to read the story, figure out which calculation is appropriate, do the calculation, and write the answer. An oftquoted problem on the third National Assessment of Educational Progress (secondary school mathematics) points to the dangers of this approach. It asked how many buses were needed to carry 1128 soliders to their training site, if each bus holds 36 soldiers. The most frequent response was "31 remainder 12" — an answer that you get if you follow the *practice* for word problems just described, and ignore the fact that the story (ostensibly) refers to a "real world" situation.

Even when students deal with "applied" problems, the mathematics that they learn is generally clean, stripped of the complexities of the real world. Such problems are usually cleaned up in advance — simplified and presented in such a way that the techniques the students have just studied in class will provide a "solution." The result is that the students don't learn the delicate art of mathematizing — of taking complex situations, figuring out how to simplify them, and choosing the relevant mathematics to do the task. Is it any surprise that students aren't good at this, and that they don't "think mathematically" in real world situations for which mathematics would be useful?

I'm proposing here that thorny issues like the "transfer problem" (why students don't transfer skills they've learned in one context and use them in other, apparently related ones) and the failure of a whole slew of curriculum reform movements (e.g. the "applications" movement a few years back) have, at least in part, cultural explanations. Suppose we accept that there is such a thing as school culture, and it operates in ways like those described above. Curricular reform, then, means taking new curricula (or new ideas, or....) and shaping them so that they fit into the school culture. In the case of "applications," it means cleaning problems up so that they're trivial little exercises — and when you do that, you lose both the power, and the potential transfer, of the applications. In that sense, the culture of schooling stands as an obstacle to school reform. Real curricular reform, must in part involve a reform of school culture. Otherwise it doesn't stand a chance.

Well, here I am arguing away in the midst of — as though you haven't guessed — phase 3. There are two main differences from phase 2. The first is that I've moved from taking snapshot views of students (characterizing what's in a student's head when the student sits down to work some problems) to taking a motion picture. The question I'm exploring now is: how did what's in the student's head evolve the way it did? The second is that the explanatory framework has grown larger. Though I still worry about "what's going on in the kid's head," I look outside for some explanations — in particular, for cultural ones.

And yet *plus ça change, plus ça reste le méme*. I got into this business because, in Halmos's phrase, I thought of problem solving as "the heart of mathematics" — and I wanted students to have access to it. As often happens, I discovered that things were far more complex than I imagined. At the micro-level, explorations of students' thought processes have turned out to be much more detailed (and interesting!) than I might have expected. I expect to spend a substantial part of the next few years looking at videotapes of students learning about the properties of graphs. Just how do they make sense of mathematical ideas? Bits and pieces of "the fine structure of cognition" will help me to understand students' mathematical understandings. At the macro-level, I'm now much more aware of knowledge acquisition as a function of cultural context. That means that I get to play the role of amateur anthropologist — and that in addition to collaborating with mathematicians, mathematics educators, AI researchers, and cognitive scientists, I now get to collaborate with anthropologists and social theorists. That's part of the fun, of course. And that's only phase 3. I can't tell you what phase 4 will be like, but there's a good chance there will be one. Like the ones that preceded it, it will be based on the wish to understand and teach mathematical thinking. It will involve learning new things, and new coleagues from other disciplines. And it's almost certain to be stimulated by my discovery that there's something not right about the way I've been looking at things.

Are there any morals to this story — besides the obvious one, that I've been wrong so often in the past that you should be very skeptical about what I'm writing now? I think there's one. My work has taken some curious twists and turns, but there has been a strong thread of continuity in its development; in many ways, each (so-called) phase enveloped the previous ones. What caused the transitions? Luck, in part. I saw new things, and pursued them. But I saw them because they were there to be seen. Human problem solving behavior is extraordinarily rich, complex, and fascinating — and we only understand very little ofit. It's a vast territory waiting to be explored, and we've only explored the tinest part of that territory. Each of my "phase shifts" was precipitated by observations of students (and, at times, their teachers) in the process of grappling with mathematics. I assume that's how phase 4 will come about, for I'm convinced that — putting theories and methodologies, and tests, and just about everything else aside — if you just keep your eyes open and take a close look at what people do when they try to solve problems, you're almost guaranteed to see something of interest.

References

Descartes, Rene [1980]. *Rules for the direction of the mind.* In R. Descartes, *Philosophical essays.* Indianapolis, In: Bobbs-Merrill, 1980

Polya, G. *How to solve it* (2nd ed.) [1957] Princeton, N.J.: Princeton University Press

Polya, G. [1954]. *Mathematics and plausible reasoning.* (Two volumes: *Induction and analogy in mathematics* and *Patterns of plausible inference.*) Princeton, N.J.: Princeton University Press

Polya, George [1981]. *Mathematical discovery* (Combined paperback edition) New York: Wiley

Schoenfeld, Alan H. [1985] *Mathematical problem solving.* New York: Academic Press

Schoenfeld, Alan H. (in press). When good teaching leads to bad results: The disasters of "well taught" mathematics classes. *Educational Psychologist*

Schoenfeld, Alan H. (in press). What s all the fuss about metacognition? In Alan H. Schoenfeld (Ed.), *Cognitive science and mathematics education.* Hillsdale, N.J.: Erlbaum

There is a major reformulation of the major concerns of the philosophy of mathematics away from the logical and towards the phenomenological-historical-experiential. This takes mathematics "down from the sky", so to speak, and says, in effect, *we* are thinking these thoughts, *we* are writing these symbols, *we* are doing these mathematical things, and as a result of *our* activity, the consequences to *us* are such and such. Part of mathematics is beyond our control, just as part of the universe is. We may not be able entirely to separate the parts. Nonetheless a heightened sense of the interaction between mathematical thoughts and arguments and human life is emerging. Whether or not one has more or less mathematization may not really be the issue. What may be crucial is whether society develops a self-awareness that in its ordinary mathematical usages it is arranging itself in certain ways and hence is doing something to itself. In this way, mathematics becomes a human institution.

Philip J. Davis and Reuben Hersh