DOCUMENT RESUME

ED 228 045 SE 040 921

AUTHOR Schoenfeld, Alan H.
TITLE Theoretical and Pragmatic Issues in the Design of Mathematical "Problem Solving" Instruction.
PUB DATE Apr 83
NOTE 36p.; Paper presented at the annual meeting of the American Educational Research Association (Montreal, Quebec, Canada, April 11-14, 1983).
PUB TYPE Viewpoints (120) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *College Mathematics; Critical Thinking; *Educational Theories; Higher Education; *Mathematics Education; *Mathematics Instruction; *Problem Solving; *Teaching Methods

ABSTRACT

This paper considers the nature of "understanding" in mathematics, and of instruction designed to foster it. The first part is theoretical, presenting an argument that there are (at least) three qualitatively different components to competent mathematical performance: (1) possession of the appropriate set of cognitive "resources," (2) the ability to select appropriately from the resources potentially at one's disposal, and use them with some efficiency (i.e., good "control" behavior), and (3) possession of a "mathematical perspective" that establishes the context within which resources are selected and used. Examples are given to show that the absence of any of these can cause failure in students. The second part is practical, and the focus turns to "positive" behavior. Suggestions are made for (1) solidifying resources including a focus on representations and the use of heuristic strategies, (2) inducing more efficient "control" behaviors, and (3) uncovering inappropriate (anti-mathematical) perspectives in students, and working towards replacing them with perspectives that support the development and utilization of mathematical skills. (Author)

ED228045

Theoretical and Pragmatic Issues in the Design of Mathematical "Problem Solving" Instruction

Alan H. Schoenfeld
Graduate School of Education
and Department of Mathematics
The University of Rochester
Rochester, NY 14627

Running Head: Issues in Instruction

Paper presented at the 1983 Annual Meeting of the American Educational Research Association, Montreal, Canada, April 1983.

SE040921

Theoretical and Pragmatic Issues

in the design of mathematical "problem solving" instruction

Alan H. Schoenfeld
Graduate School of Education
and Department of Mathematics
The University of Rochester

Abstract

This paper considers the nature of "understanding" in mathematics, and of instruction designed to foster it. The first part is theoretical, presenting an argument that there are (at least) three qualitatively different components to competent mathematical performance: (1) possession of the appropriate set of cognitive "resources," (2) the ability to select appropriately from the resources potentially at one's disposal, and use them with some efficiency (i.e., good "control" behavior), and (3) possession of a "mathematical perspective" that establishes the context within which resources are selected and used. Examples are given to show that the absence of any of these can cause failure in students. The second part is practical, and the focus turns to "positive" behavior. Suggestions are made for (1) solidifying resources including a focus on representations and the use of heuristic strategies (2) inducing more efficient "control" behaviors, and (3) uncovering inappropriate (anti-mathematical) perspectives in students, and working towards replacing them with perspectives that support the development and utilization of mathematical skills.

Theoretical and Pragmatic Issues In the Design of Mathematical "Problem Solving" Instruction

Introduction

I first offered an undergraduate course in mathematical problem solving at Berkeley in 1976, and have periodically offered various incarnations of it since then. The most recent version is in progress this semester at Rochester. While many of the ideas underlying the course and many of the problems I use are the same as they were seven years ago, many things are different as well. I would like here to describe the major change in my perspective, which establishes the context for the balance of this paper.

When I began teaching "problem solving" I took that phrase, broadly construed, as an operational definition of understanding: you understand how to think mathematically when you are resourceful, flexible, and efficient in your ability to deal with new problems in mathematics. To be resourceful and flexible, students needed (I thought) to be familiar with a broad range of general problem solving strategies, known as heuristics. To be efficient, they needed coaching in how to "manage" the resources at their disposal. My course tried to provide both - and in some ways, still does. Consistent with the operational view expressed above, the "acid test" for its success was that, after instruction, the students should show marked improvement on a collection of problems not related to the ones they had studied in the course. (A later version of the course passed the test, with flying colors; see Schoenfeld, 1982.)

The early versions of my problem solving course reflected a positivist, "cognitive engineering" perspective: one proceeds by modeling an "ideal" way to solve problems (that model usually, but not necessarily, based on experts' performance), and then training students to perform in accordance with the model. As the acid test indicates, such courses can be successful.

My notion of success has changed in recent years. More accurately, my notion of mathematical understanding (and. therefore, of what I was teaching in the course) has changed. This change was induced largely by the part of my research that called for detailed examinations of my students' problem solving performance before they entered my course. My students were, by most measures, the successes of our educational system. Virtually all of them had completed at least one semester of calculus (many, three) with grades of B or better before enrolling in my course. Their enrollment itself indicated that they were partial to mathematics (it was an optional course, fulfilling no requirements) and relatively confident about their abilities (it had a reputation for being difficult). Videotapes of these students' problem solving performance revealed some unpleasant realities. They pointed to serious misunderstandings about mathematics, and often to deeply held anti-mathematical perspectives in these "successful" students. In some cases they indicated that students could work quite functionally in domains about which they understood virtually nothing. In others they indicated that the students "knew" a great deal that went unused because they felt that knowledge to be useless. At a theoretical level, the research indicated that my earlier notion of "problem solving" as an operational definition of "understanding" was a gross simplification: one's behavior in problem solving is a complex combination of (I) the cognitive resources potentially at one's disposal, (II) the ability one has to "oversee" the selection and deployment of those resources, and (III) the set of beliefs one has about the discipline, the environment, the task, and oneself - the beliefs that, in essence, determine the context within which one selects and deploys the cognitive resources potentially accessible in Long Term Memory. The next section of this paper provides a brief characterization of these categories, and of the roles that they play in "understanding." With this as theoretical backdrop, we then discuss some practical issues related to instruction.

Three Categories of Understanding*

Figure 1 outlines the contents of this section. There is an enormous literature on the first category, a substantial and growing literature on the second, and a sparse literature on the third. I believe that they are all critical components of understanding, and of problem solving performance.

Insert Figure 1 about here

I: Resources

This category is quite broad, comprising the range of facts and procedures potentially available to an individual problem solver for implementation. In a characterization of experts' problem solving, Simon (1979) describes the key issues as follows: "(a) how much knowledge does an expert or professional in the domain have stored in LTM (Long Term Memory), and (b) how is that knowledge organized and accessed so that it can be brought to bear on specific problems?" A third question is implicit: (c) how is that knowledge represented, and how does the nature of the representation affect access and implementation?

To begin with, there is the question of what domain-specific knowledge is accessible to the problem solver. Does a student trying to solve a

*This section takes liberally and extensively from my three in-press articles. My purpose here is to be summary-descriptive, and I have not provided extensive documentation for the positions summarized here; bibliographies are given in the in-press papers. There are, however, three recent reviews worth mention that provide broad and deep coverage of the relevant literatures: see the papers by the Federation of Behavioral, Psychological, and Cognitive Sciences; by Greeno and Simon; and by Nickerson.

Three Qualitatively Different Categories of Knowledge and Behavior Required for a Characterization of Human Problem Solving

Category I: Resources (Knowledge possessed by the individual, that can be brought to bear on the problem at hand)

- -- Facts and algorithms
- -- Relevant competencies, including the use of routine procedures, "local" decision-making, and implementing "local" heuristics.

Category II: Control (Selection and Implementation of Tactical Resources)

- -- Monitoring
- -- Assessment
- -- Decision making
- -- Conscious metacognitive acts

Category III: Belief Systems (Not necessarily conscious determinants of an individual's behavior)

- -- About self
- -- About the environment
- -- About the topic
- -- About mathematics

figure 1

straightedge-and-compass construction problem from plane geometry know that the radius of a circle is perpendicular to the tangent line at the point of tangency? Whether the student chooses to use that fact is another matter (see Category III). But, clearly, solutions depending on that knowledge would evolve in radically different ways if the students did or did not "know" it. The same holds for relevant procedures. Does the student know how to construct a perpendicular to a line through a given point? If not, does the student know that it can be done, so that deriving the construction is a possibility? Or must that too be "discovered?"

Beyond "possession" of factual and procedural knowledge comes the question of access to it. The student may "know" a particular fact, but will the student "see" that it is relevant in a particular problem? The literature indicates that much "routine" expert performance in a variety of domains is due to the possession of problem "schemata," which provide more or less "automatic" responses to generic situations; see, e.g., deGroot (1966). Experimental results in pysics (Chi, Feltovich, and Glaser, 1981) and mathematics (Schoenfeld and Herrmann, 1982) indicate that experts see through the "surface structure" of problems to perceive "deep structure" similarities and approach the problems accordingly. Moreover, students develop problem schemata that may or may not be consistent with those experts (Hinsley, Hayes, and Simon, 1977; Silver, 1979), and these schemata change with experience (Schoenfeld and Herrmann, 1982). For a characterization of the role of schemata in students' mathematical problem solving performance, see Silver, 1982. Again, issues of representation are a consistent thread throughout this literature (Greeno and Simon, 1983).

One further class of resources consists of having access to, and being able to implement, various problem solving heuristics. I shall not address here the differences between mathematics and other disciplines (e.g. physics) that accord heuristics a special status in mathematics, but will simply note that there is a fair body of research in mathematics indicating the role of heuristic fluency as a component of mathematical competence. The point I wish to make here is that the domain-specific implementation of many such strategies is very much on a par with the implementation of domain-specific schemata. Consider, for example, the following general strategy: "To discover useful information about the object you are trying to find, assume that you have that object and determine the properties that it must have.". The "plane geometry" version of the strategy is "Draw a figure to see what properties it has," which is quite similar to a domain-specific schema such as "look for congruent triangles when trying to prove certain quantities are equal." Such heuristics, like the facts, procedures, and schemata discussed above, comprise the set of tools potentially accessible to the problem solver.

An inventory of these resources provides a characterization of what the problem solver _might_ be able to use in solving a problem. Since we are dealing with the "real" behavior of students rather than the idealized behavior of experts, there are no guarantees that these resources will be called upon, even if it is appropriate for the problem solver to do so: this is where this discussion diverges from Simon's, as quoted above. Whether the tools potentially accessible to the problem solver are selected or discarded, how such decisions are made, and how such choices affect the problem solving process as a whole, are another category of behavior.

II: Control

Consider the following problem. "Three points are chosen on the circumference of a circle or radius R, and the triangle containing them is drawn. What choice of points results in the triangle of largest possible area? Justify your answer as best as you can." If space permitted, I would present here the full transcripts of the dialogues produced by my students as they worked this problem: one best sees the effects of "control" failures by examining the effects of those failures on the full solutions. See my (in-press, a) for typical transcripts and the details of an analytical framework for analyzing them. One typical transcript is summarized below.

The students K and A had completed one and three semesters of calculus respectively, with A just having finished a course including the multivariate techniques that provide the analytical (rather than heuristic) solution to the problem. They read the problem and quickly sketched out a representative case, indicating that they had understood it. After a total of 35 seconds elapsed, the following dialogue ensued:

<u>K</u>: I think the largest triangle should probably be equilateral...

<u>A</u>: So we have to divide the circumference of the three equal arcs to get this length (the length of a side) here. So, 60-120 arc degrees...

<u>K</u>: Do we have...(reads statement) justify your answer as best you can. Justify why this triangle...justify why you...right.

<u>A</u>: OK, let's somehow take a right triangle and see what we get. We'll get a right angle.

They made a brief (15 second) calculation on the right triangle, and then returned to the equilateral:

<u>A</u>: OK, but what we'll need is to say things like -- OK, let's go back to the angle -- probably we can do something with the angle.

It is worth taking a close look at this dialogue, which shaped the rest of their solution. <u>K</u> began with the obvious conjecture, that the equilateral triangle is the answer. <u>A</u> began calculating its area immediately, although (1) the problem does not ask for the area of the largest triangle, and (2) <u>it</u> is not clear that knowing that area will help. At this point <u>K</u> reread the problem statement, obviously worried about the phrase "justify your answer." They decided to look at the right triangle, a minor digression that represents, I think, a weak attempt at justification: if the area of the equilateral turned out to be larger than the area of the right triangle -- another archetype -- they would have had slightly more reason to believe it to be the largest. After this brief digression, they moved wholeheartedly into the calculation of the equilateral's area.

Note that the decision to undertake this calculation, which was to occupy them for the balance of the solution, was made without any overt discussion of its utility or relevance to the solution. They simply began calculating. Now there are times when it is appropriate to jump into a solution attempt without reflecting upon it, for example when one proceeds along familiar lines in a schema-driven solution. But when the territory is somewhat unfamiliar, some prudence in establishing one's directions seems to be called for. If a bad decision is made, and then not revoked, that one decision dooms the entire solution to failure. That is, in fact, what happened to K and A.

Having decided to determine the area of the equilateral triangle, the students began calculations. These became far more tortuous than they had expected. About seven minutes into the solution, as their energies flagged a bit, there was the following dialogue:

A: There used to be a problem...about the square being the biggest part of the area...

K: the largest area of...something in a circle, may be a rectangle, something like that...

A: Ah, well...

K: So this is R, and this is going to be 120 degrees, and...

This brief interaction provided the students with the first clear opportunity to pause and take stock. They might have asked: How well have we done? Are there other possibilities? Should we take those? Should we look for others? Instead the alternative faded out of the picture and, with renewed energy, they embarked once again on the calculation. (Note, incidentally, that K was the one to resume the calculations this time; A began them earlier. They both share responsibility for the one-dimensionality of their solution approach.) There were two similar occurrences later in the solution. The possibilities of approaching the problem by way of the calculus, and also by means of an heuristic variational argument, were mentioned -- and dropped -- in passing. When the twenty-minute video-cassette recording the students ran out of tape and clicked off, they were still calculating. When I asked them what good it would do to have the area of the triangle, they could not say.

I wish to argue that the students' failure on this problem was not, as in the examples in the previous section, a result of their lack of basic skills. (A had solved an equally complex multivariate problem on a final exam less than two weeks before he was videotaped; he clearly had the requisite skills.) Rather, their failure was strategic. First, they made no attempt to generate or consider plausible approaches to the problem. The first one that came to mind was embarked upon. Second, they did not evaluate the approach they took: after having pursued it for a full twenty minutes, they could not say what value it might have had. Third, they did not actively monitor and assess their progress (or lack of it) during the solution. An active "executive" might have curtailed their wild goose chase, and provided the opportunity for more productive behavior. Such an executive might have considered the three other possibilities that arose (a related calculus problem, maximization via calculus, a variational argument) and, at minimum, proposed that they be explored and evaluated before they were abandoned. The point is not that this kind of executive guarantees success, for success clearly depends on a variety of skills. Rather, the point is that such an executive (strategist) helps to avoid what is otherwise a guaranteed failure. Moreover, such failures are far from infrequent. The transcript summarized above was one of twelve recorded in a 1981 experiment. Of those twelve, seven could be categorized as being of the "read and then go on a wild goose chase" variety. The absence of planning, monitoring, assessing, -- in general, overseeing a solution -- can do students great harm. So can the imposition of inappropriate control decisions. If space permitted, I would discuss a protocol where students literally throw away the elements of a solution that is at their fingertips. Bad "control" is a consistent source of failure. See my in-press, a for details.

III: Beliefs

Consider the following problem. You are given two intersecting straight lines and a point P on one of them, as in the figure below. Show how to construct, with straightedge and compass, a circle that is tangent to both lines and has the point P as its point of tangency to one of them."

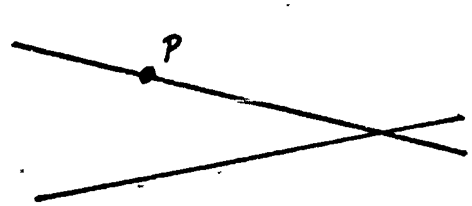

If one thinks to draw in the desired circle and to derive the properties it must have, the problem is nearly trivial (provided, of course, one has access to the relevant geometric knowledge): the center of the circle lies on the intersection of any two of (a) the perpendicular to the top line through P, (b) the perpendicular to the bottom line through P's direct "opposite," and (c) the bisector of the vertex angle. That is not how students proceeded, however. The following describes a typical session.

The students S and B read the problem, drew a rough sketch, and conjectured that the diameter of the desired circle was the line segment between P and its "opposite," P' (see Figure 2a). They reached off to the side for a straightedge and compass, in order to test their conjecture empirically. They took great care with the construction, working on it for five minutes or so. It didn't look right when they finished it, so they rejected the conjecture. At that point one of them recalled that the radius of the circle was supposed to be perpendicular to the tangent at P. As a result they revised their conjecture: now the diameter of the circle was, most likely, the segment of the perpendicular through P that lies between the two lines (Figure 2b). This, too, was tested empirically. Five more minutes elapsed in constructions, and when they examined the results of this construction, they decided that this guess was also incorrect. They re-examined their original sketch, and one of the students noticed that the center of the circle seemed to be half-way between the two given lines. Perhaps, then, the center lay on the angle bisector. This gave rise to a third conjecture (serendipitously correct; see Figure 2c), which was again tested by construction. This time they were "successful." Sketch in hand, they reported to me -- more than sixteen minutes after they had begun the problem -- that they had solved the problem. When I told them that their construction was indeed correct, and asked them if they could tell me why it worked, they said they had no idea; it just did.

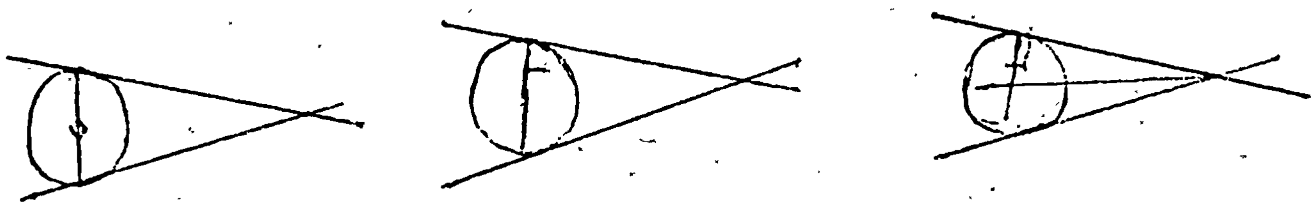

a. first conjecture

b. second conjecture

c. third conjecture

figure 2

The behavior of these two students was more typical than not. (See my in-press, b for another protocol and a detailed model of students' performance on this kind of problem.) Without exception, my students guessed at the nature of the solution and then tested their guesses by performing their hypothesized constructions. If the guesses were wrong (that is, the constructions did not "look right") they would try again, with another construction. Most often the students took fifteen to twenty minutes to finish the problem, the majority of that time spent with straightedge and compass in hand. (They either found a construction that worked and reported success, or ran out of hypotheses and reported failure.) Not one pair of students reasoned their way to a solution. Only one pair justified their solution afterwards, and this justification was an after-the-fact observation: "We got it, that is, it looks pretty close....I think, if you center it right, they touch...OK. These (PCV and P'CV in Figure 3) are similar triangles...Yeah...They are equal triangles too." In sum, their behaviors were purely empirical: the students' classroom knowledge of geometry was, with the sole exception just quoted, nowhere to be seen.

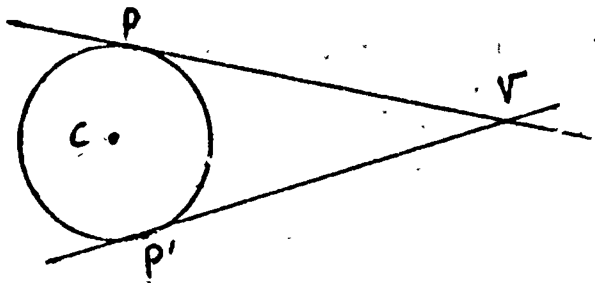

figure 3

It would be depressing enough to think that these students -- again, the "successes" of our system -- had forgotten all of the geometry that they had learned in high school. What is more depressing is that they remembered it. After the students finished their work on the problem, I asked them to work two "standard high school problems" for me: to prove that $\overline{PV} = \overline{P'V}$, and that CV bisected vertex angle V, in Figure 3. Most of the students were able to do so without difficulty.

Why, then, had the students not used this knowledge in working the Problem ? I submit that it is because they had no idea that geometric deduction ("proof") would be of the least value to them. During the students' careers, deductive arguments like the ones in the "standard high school problems" given above were not perceived as being useful at all. From the students' point of view, you only "proved" in class what you already know to be true; that's the classroom game you play, to satisfy the teacher. Proof never shows you something new, but only confirms what you know already. Writing proofs in geometry is seen as being similar to conjugating verbs in Latin. That is, it may be good for "discipline" and for "training the mind," but it has nothing whatsoever to do with thinking or solving problems. As a result, this kind of "classroom knowledge" is left behind in the classroom: perceived as useless, it goes unused. While it may seem melodramatic, I submit that the student who believes that deductive reasoning has no real value is about as likely to think of using it when confronted with a "real problem" as an atheist is likely to look for "divine intervention" as an explanation for a particular phenomenon, or a creationist is likely to think of evolutionary mechanisms as a way to explain the current configuration of a particular anatomic structure. We are confronted with an uncomfortable paradox: as a result of their mathematical training, these students have adopted an anti-mathematical stance. This stance, in a very real sense, bars them from using mathematics effectively. The belief underlying this stance (which, I should stress, may not be consciously held) and two other beliefs that significantly affect students' behavior, are as follows.

Belief 1: Formal mathematics has little or nothing to do with real thinking or problem solving.

Belief 2: Mathematics problems are always solved in less than ten minutes, if they are solved at all. Corollary: Give up after ten minutes.

Belief 3: Only geniuses are capable of discovering or creating mathematics. First corollary: If you forget something, too bad. After all, you're not a genius and you won't be able to derive it on your own. Second corollary: Accept procedures at face value, and don't try to understand why they work. After all, they are derived knowledge passed on "from above."

Practical Issues

The discussion in the previous section suggests the perspective that I shall take here:

> To be successful in training students to think mathematically, instruction in mathematics must (a) provide the students with a solid collection of resources, the more integrated the better; (b) assist the student in developing an effective set of "control" (planning, monitoring, assessment, decision-making) behaviors, and (c) serve to develop in students a set of perspectives (beliefs) that promote the students' propensities and abilities to use mathematics where it is appropriate to do so.

Note that I have not used the phrase "problem solving" in this statement, or suggested that such activities must or should take place in a separate course. The deeper issue is mathematical understanding, although the ability to solve problems is clearly a large component of that understanding. In the best of all possible worlds, a separate course with a heavy emphasis on problem solving would be unnecessary; (a) through (c) above would be developed in ordinary coursework. The research clearly indicates, however, that that does not take place. Separate courses have two primary virtues. The first is the freedom for teacher and students to explore their understanding of mathematics, without feeling the pressure to "move on" in order to cover other material listed in the course syllabus.* The second is that problems out of context are more

*It should not be inferred from this comment that separate problem solving or "understanding" courses are easier; only that they are different with regard to coverage: As the research cited above indicates, much of what students "learn" when standard material is "covered" may be illusory.

likely to reveal students' misconceptions. My comments will be made in the context of a separate course, but they apply in general. Where suggestions of how to deal with such issues have received adequate attention elsewhere, or when they are likely to be covered in some depth at this symposium, I shall be very brief.

I. Resources

I suspect that the talks by Joan Heller and Fred Reif, and by Jim Greeno, will say a great deal about promoting effective problem solving behavior and understanding. Speaking broadly, the general idea is to be as explicit as one possibly can in

(1) elucidating a set of effective procedures and behaviors that are within students' capacity to learn (often but not necessarily modeled on experts' performance), while

(2) helping students to develop cognitive "support structures" for those procedures and behaviors (often the appropriate representations; also the appropriate perceptions of stereotypical circumstances, for schema-based actions).

There is, fortunately, a growing literature on these topics. The best quick summary of useful suggestions can be obtained by taking the subsection titles from Section 3 (How might understanding be facilitated?) of Nickerson's (1982) "Understanding Understanding:" "Start where the student is; promote discovery; insist on active processing; use representations; use analogies and metaphors; simulate; explore the reasons for misconceptions and errors; teach evoking conditions for procedures; vary the context; make connections among concepts being taught; relate classroom teaching to real-world problems;

encourage the acquisition of domain-specific knowledge; reinforce efforts to understand; provide a generally supportive atmosphere."

As far as this list goes, I can only say "amen." I should point out, however, that most of the suggestions (an accurate reflection of the literature) deal with resources. There is little about control behavior or (save for misconceptions, and relating teaching to real-world problems) about beliefs.

II. Control

The presentation by Joan Heller and Fred Reif has something to say about the "positive" aspects of control behavior: the paper deals with the results of encouraging students in the domain of mechanics to obtain appropriate representations, facilitating good decision-making and search, and promoting the assessment of solutions' correctness and optimality. The more domain-specific one is, the more one can specify and encourage these "positive" control behaviors. Clearly more needs to be done in elucidating domain-specific (and generalizable) control strategies. Yet as we saw in the discussion of "control" in the previous section, a major function of the "executive" in decision making in general problem solving consists of making certain that explorations are justified and that wild goose chases are terminated before they become debilitating. This discussion will focus on those general, domain-independent behaviors. Since there is an extensive "how to" section that focuses largely on control decisions in (Schoenfeld, 1983), I will focus here on the rationale for such discussions in the classroom and make a few brief suggestions. The following three points form the basis of the rationale.

1. The very notion of serving as your own cognitive "manager" or "coach" - monitoring and assessing your progress as you work on a problem, and altering your solution as a result of those assessments - is completely alien to virtually all of my (college freshman) students. Their perception is that their minds are more or less autonomous entities when it comes to problem solving: they just do "what comes to mind."

2. While the costs of bad "resources" are immediate and apparent, the costs of bad control are usually not. Students may go off on a "wild goose chase," in effect prohibiting themselves from using what they know to solve a problem. They may make unwarranted assumptions (that could easily be caught by the appropriate "monitor") that invalidate their attempt from the beginning. In abandoning an unsuccessful attempt, they may throw away the elements of a solution. (Examples of all of these are in my in-press, a.) In all of these cases, failure is induced at the control level. But unless there is immediate feedback and evaluation, the failure is likely to be attributed elsewhere: "I didn't see that it could be done that way," or "I really had no idea of what to do."

3. Control processes are generally invisible. Most often when a student sees a problem "explained," the student sees a discussion of what "works." Most classroom "solutions" of problems are schema-based (on the part of the instructor), so students do not see the teacher "think." On those occasions where there is a real "problem,"

one usually sees the following. The teacher says "wait a minute, let me think about that;" there is a period of silence as the teacher thinks; the silence is broken with "All right, let's look at it this way." The solution proceeds. If the teacher gets "stuck" while solving a problem, the same generally happens - unless the teacher has reached a real impasse, in which case the attempt is usually terminated with a promise to present the solution at the next class meeting. Thus students have no models for the control processes discussed above.

The best way to bring these matters out in the open, and to provide useful models of control decisions for the students, is to engage in "real" problem solving in the classroom. At minimum, I recommend that the teacher elucidate the full decision process when presenting solutions to problems: that is, to explain what takes place during the silences alluded to in (3) above. ("All right, what are the options here? We might look at A, or B, or maybe C. It looks like A might be worth a try for a few minutes, because...) It is useful to work through problems on this kind of "blow by blow" basis, even when they are familiar. While the teacher knows what to do and why, the student does not: seeing the decision process modeled will, at least, "legitimize" it for the students. Even so, this legitimization is only the first step. I suggest two methods of group problem solving during class sessions (and use each with equal frequency). They are both described extensively, with some sample classroom sessions, in Schoenfeld (1983).

A: THE CLASS WORKS THE PROBLEM AS A WHOLE, WITH TEACHER AS "MANAGER."

I pose a problem for the class to solve, and invite suggestions for its solution.* Often suggestions will come quickly - too quickly, and frequently barely relevant to the problem (the first step of what might well be a wild goose chase). My role is not to judge the suggestions, but to point out the speed with which they were raised:

"Is everyone sure that they understand the problem, before we proceed with the solution?"

If the answer is "no," the class takes whatever steps are appropriate to remedy that: examining the conditions of the problem, looking at special cases, drawing a diagram or finding another appropriate representation, etc. Having done so, we return to the original suggestions. The class is asked if they seem reasonable. (If one or more of the suggestions now appears unreasonable, this may occasion a "sermon" about making sure you understand before proceeding, and about the dangers of wild goose chases.)

"Is there anything else we ought to look at or try?"

If there is only one plausible approach, we take it - after making sure that what we are doing is reasonably well defined, and that we have a sense of how we will use it in the solution. If there is more than one (as is often the case when you work "problems" rather than "exercises) we discuss the relative merits of the approaches, and what we might gain from them. The three "generic"

*Two comments. First, these are reasonably difficult problems that may take the class as a whole anywhere from ten to fifty minutes (or longer) to solve. Second, this kind of interaction only works well when the students feel free to make suggestions. There is a delicate tension between the wish to critique suggestions and the fear that students, once "scolded," will cease to participate. The way out of this is for the teacher to stay strictly in the role of "monitor," raising questions about the efficacy of suggested steps (both when they are useful, and not!). When the suggestions are ill-founded, the students will discover that for themselves.

questions for this discussion are given in (B) below. Having made our choice, we proceed with the solution. After five minutes or so have elapsed, we pause.

> "All right, we've been doing this for a while. Is it working? Are things going according to plan, or should we reconsider?"

Note that it is important to ask these questions even when things seem to be going well. Otherwise, they become a "cue" from the teacher that directions should be changed. The idea is that these questions should always be in the back of one's mind; they advance to the front when solutions seem to bog down. The class may decide to proceed, to proceed with caution ("We'll give it another three minutes, and then reconsider"), or to change directions.

> "Before we abandon this approach, is there anything in it that should be salvaged? Are there any ideas in it that we might want to return to, or related topics suggested by this approach, that we might want to explore if our new approach doesn't work out?"

The discussion continues in this vein until (with luck) the problem is solved. (We will occasionally stop at an impasse, and continue with the problem another day. On those occasions when, after such an impasse, it still appears that the class is ignorant of some relevant knowledge and is unlikely to derive it, I may then provide them with it.) We may then pursue some of the other suggestions, and solve the problem two or three different ways.

While the class works on the problem, my contributions are kept to a minimum. If the class decides to pursue a direction that I know leads to a dead end, I will let them - so long as the decision was reasonably made.

That happens, after all, in good problem solving; an effective "monitor/assessor" keeps such decisions from being fatal. In general, the role of the "external manager" played by the teacher is to help the students to get the most out of what they know: to ensure that they have fully understood a problem before embarking on its solution, that they have looked for good representations, that they generate and select approaches to the problem with care, that they capitalize on opportunities that arise during the solution, that they employ the resources at their disposal, and avoid squandering their energies on where it is clearly inappropriate. Nothing the external manager does depends on knowledge of the problem not accessible to the students; everything this manager does could be done (with the same positive effects) by the students themselves. In other words, all of the "control" functions performed in the classroom discussion could be internalized by the students, without additional knowledge.

After the problem has been dispatched with, I step back into the role of teacher to do a "post mortem" of the solution. This includes a discussion of problem representations, of related knowledge that might or should have been called into play, of the students' effective or ineffective use of control strategies (I occasionally let the class go on a wild goose chase, to point out what happens when one fails to exert executive control), and of elements in the students's approach that could, if pursued or exploited differently, have yielded insights into the problem or different solutions of it.

B. THE CLASS BREAKS INTO SMALL GROUPS TO WORK ON PROBLEMS, WITH THE TEACHER AS ROVING "CONSULTANT."

About a third of the time in my problem solving course is spent with the class divided into groups of four, working on problems that I have just handed out. As the students work on the problems, I circulate through the class as a possible source of help. Again, my role is not simply to provide information or hints, although I will if the situation calls for it. More often than not, my response to a request for a hint will be in the form of a (heuristic) question: Does that problem remind you of anything? Have you done something similar recently? Can you reduce it to something simpler?, etc.

In the small groups the emphasis on "control" decision remains, but the responsibility for it shifts from the external manager to the students. There is a large poster in the classroom with the following three "executive" questions:

> What (exactly) are you doing?
> (Can you describe it precisely?)
>
> Why are you doing it?
> (How does it fit into the solution?)
>
> How does it help you?
> (What will you do with the outcome when you obtain it?)

Their frequent inability, at the beginning of the course, to answer these questions - coupled with the realization that if they cannot, they have most likely been wasting their time* - serves as a strong catalyst for the internalization of these questions.

To cast the preceding discussions in a (one of a number of possible) theoretical framework, we can see the large group and small group discussions as the social mediating factors that help move the student through a "zone of proximal development" (a la Vygotsky) to the point where the appropriate behaviors are internalized. At a more pragmatic level, the following four reasons justify the use of class time for this admittedly time-consuming process.

1. This format affords the teacher the unique opportunity to intervene directly as the students solve problems, rather than being faced with a "finished product." That intervention (as indicated above) is much stronger than it could be in any other instructional format. Moreover, group formats bring these ordinarily covert control processes "out in the open," where they can be examined.

2. Solving problems with a small number of one's peers provokes discussions of plausible choices. When a student works a problem alone, the first plausible option is often the one taken. When different students have proposed approaches to the problem and they must settle on one, there must be a discussion of the merits of the approaches they have proposed - precisely the kind of discussion that the students should be having, internally.

*One point of clarification: these questions are not meant to rule out exploration. "I'm mucking around for a few minutes hoping to find some inspiration" is a perfectly good answer to them - so long as the student is aware of what's taking place, and doesn't let the solution process degenerate into an extensive and directionless series of random explorations. The purpose of the questions is simply to make sure the student is "in control" of what's taking place.

3. Problem solving is not always a solitary endeavor. This opportunity to engage in collaborative efforts does them no harm.

4. Students are remarkably insecure about their abilities, especially in a course of this nature. Working on problems with other students is reassuring: one sees that one's fellow students must also struggle to learn.

III: Belief

Dealing with students' beliefs and the effects of those beliefs on cognition is a far more subtle issue than either of the preceding ones, and my suggestions for it are far more primitive. We are just coming to understand the importance and impact of beliefs on cognition. One major aspect of belief systems has received some attention and will, I believe, be addressed in this symposium by Jack Lochhead and Lillian McDermott: Students' misconceptions about physics (their "naive physics") may interfere with their learning of the principles of physics, or may simply render the students' "book knowledge" meaningless (see Caramazza, McCloskey and Green, 1981; McCloskey, Caramazza, and Green, 1980; Trowbridge and McDermott, 1980). There is clear evidence that students enter their physics courses replete with a collection of naive, Aristotelian (pre-Newtonian) views of physical phenomena. Many of these students, in spite of doing well in those courses, emerge from them with their naive physics intact: asked to interpret a "real world" phenomenon, they invoke their pre-instruction models.

The research on misconceptions in physics indicates that students will not invoke certain kinds of formal reasoning if they believe that they have better explanations at their disposal. The research described on geometry in the previous section is similar, and perhaps more troubling: if students have decided that certain kinds of knowledge are "useless academic classroom tasks" they may not think to employ that knowledge, even when (1) they have access to it and could use it to solve the given problems easily, and (2) they are stymied without it.

The "moral" of this research is significant: students are not tabula rasae, waiting for knowledge to be printed on their "mental blackboards." Those blackboards have been extensively written upon, and what we try to write on them will only be meaningful if it connects with, or replaces, what is there. Probing for misconceptions is essential: one must "clean the slate," to continue the metaphor. There are two aspects to this. In "making connections," it is important for teachers to do so overtly; students may not see them otherwise. Nickerson's (1983) suggestions "simulate; make connections among concepts being taught; and (especially) relate classroom teaching to real world problems" all serve this end. The second aspect, helping students to remove inappropriate beliefs or ideas, is much more difficult: Those beliefs must be discovered before they can be dealt with.

We are most unlikely to see evidence of misconceptions or "misbeliefs" as long as (a) we are presenting material to the students, or (b) they are presenting to us what they believe we want to see - e.g. formal mathematics or formal physics in an obviously "formal" setting. Misbeliefs are only likely to surface if students are given the opportunity to show us what they "know."

In the classroom I have found that the most effective way to find out what lies beneath the surface of students' performance is to repeat (in different forms) one simple question: WHY? I shall give one example here.

The discussions took place this semester at Rochester, in my "problem solving" class. It has eighteen students, all of whom studied geometry in high school, and most of whom took calculus at the college level (and did well in it). I introduced geometric constructions with the simplest problem I could think of: Suppose you had line segments of lengths A, B, C respectively. How do you (using straightedge-and-compass) construct the triangle whose sides have lengths A, B, C? The class (acting as a whole group) solved the problem in less than a minute (mark off a line segment of length A; from its endpoints swing arcs of length B and C respectively; draw the line segments from the endpoint of the first segment to the point of intersection), and agreed unanimously that they had the solution.

"Why does that work?"

Silence.(for a long time)

"Why do you swing the arcs?"

Ditto

"When you swing the arc of length B around this endpoint, what do you get?" There were some answers to this, and we proceeded from there. The class eventually decided that the point (singular) of intersection had the property that it was simultaneously at distance B from one vortex of the desired triangle (whose base had been established as A), and distance C from the other: this must be the triangle we wanted.

"Are there any others?"

In response to the puzzled looks, I pointed out that the construction was not unique. Extended, the arcs (circles) intersected twice; moreover, one might use arcs of lengths C and B, respectively, instead of B and C. This construction led to four solutions. Response: "They're all the same."

"Why?"

After a long pause, one student said "All the triangles you get are congruent. They're the same." The class agreed.

"Are there any different-looking triangles with sides A, B, and C?"

This question left them nonplussed. The sentiment was "no," but there was no coherent argument to support that sentiment.

"Suppose you had two different triangles, each of which had sides A, B, and C. What could you say?"

They're congruent, of course...Hmm, maybe congruence has something to do with this idea of uniqueness (my language)...Forty minutes had elapsed when we closed this discussion of the problem they had "solved" in less than a minute.

We began the next session of class with the following question: "How do you bisect an angle?" Again, the class produced the construction in short order.

"Why does the construction work?"

I will spare you the details and report the result: It took longer than a half hour for the class to see that the (standard) construction yielded two congruent triangles, and that the line that resulted lay between two equal angles that had been created; thus it was the angle bisector(!).

In sum, my class spent a week (at the college level) uncovering the reasons for two constructions that they had been able to produce from memory in less than two minutes. Was this a waste of time? I believe just the opposite. The discussions not only explained "where the constructions came from" - a minor goal - but also served to legitimize geometric reasoning (notions of proof and congruence) as useful tools in thinking geometrically. In subsequent discussions "proof" and "congruence" were invoked frequently both as reasons that things worked and (MUCH more importantly) as ways of finding out what might work. This week of discussions had "unlocked" for use the content of a full year of study, that had lain stagnant in long term memory - not because it could not be accessed, but because it had been deemed worthless. The shame is that these discussions were "remedial," and took place in a collegiate problem solving course. Proper attention to the context of learning, to making knowledge meaningful, and to making sure that students "understand" should make such remediation unnecessary.

I am sorry that this example is anecdotal, for I consider it important; I intend to find "rigorous" documentation in the near future. But I hope the following two points have emerged clearly from my discussion:

(1) Teaching students to solve problems (a.k.a. "think" or "understand" calls for attention to resources, control, and beliefs;

(2) While the task is far from easy, it can be done (and is rewarding).

References

Caramazza, A., McCloskey, M., and Green, B. Naive beliefs in "sophisticated subjects: misconceptions about trajectories of objects. Cognition, 1981, 9, 117-123.

Chi, M., Feltovich, P., and Glaser, R. Categorization and representation of physics problems by experts & novices. Cognitive Science, 5, (3), 121-152.

deGroot, A. Perception and memory versus thinking. In B. Kleinmintz (ed.), Problem Solving: research method, and theory. New York: Wiley, 1966.

Federation of Behavioral, Psychological, and Cognitive Science. Report on research relevant to education in mathematics, science, and technology, submitted to the National Science Board Commission on Precollege Education in Mathematics, Science, and Technology (Draft, January, 1983). Cited with trepidation (CWT).

Greeno, J., and Simon, H. Problem Solving and Reasoning. (Draft, February, 1983). CWT.

Hinsley, D., Hayes, J., and Simon, H. From words to equations: meaning and representation in algebra word problems. In P. A. Carpenter and M. A. Just (Eds.), Cognitive Processes in Comprehension. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1977.

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

Nickerson, R. Understanding Understanding. (Draft, March, 1982). CWT.

Schoenfeld, A. Measures of problem solving performance and of problem solving instruction. Journal for Research in Mathematics Education, 13(1), January, 1982, pp. 31-49.

Schoenfeld, A., and Herrmann, D. Problem perception and knowledge structure in expert and novice mathematical problem solvers. Journal of Experimental Psychology, 1982, Vol. 8, No. 5, 484-494.

Schoenfeld, A. Problem solving in the mathematics curriculum: a report, recommendations, and an annotated bibliography. Washington, D.C.: Mathematical Association of America, 1983.

Schoenfeld, A. Episodes and executive decisions in mathematical problem solving. In R. Lesh and M. Landau (Eds.), Acquisition of Mathematical Concepts and Processes. New York: Academic Press, in press, a.

Schoenfeld, A. Beyond the purely cognitive: belief systems, social cognitions, and metacognitions as driving forces in intellectual performance. Cognitive Science, in press, b.

Schoenfeld, A. On failing to think mathematically: some causes of mathematical malfunctions in college students. Focus on learning problems in mathematics, in press, c.

Silver, E. Knowledge organization and mathematical problem solving. In F. Lester (Ed.), Mathematical Problem Solving. Philadelphia: Franklin Inst. Pr., 1982

Simon, H. Information processing models of cognition. Annual Review of Psychology, 1979, 30, pp. 363-396.

Trowbridge, D., and McDermott, L. Investigation of student understanding of the concept of velocity in one dimension. American Journal of Physics, 1980, 48(12), 1020-1028.