

ALAN H. SCHOENFELD

PROBLEM SOLVING FROM CRADLE TO GRAVE

Abstract. This rather speculative paper proposes an overarching theoretical perspective for characterizing human decision-making and problem solving “in the moment.” The scope is deliberately broad. My intention is to address the following question: “How and why do people make the decisions they do, as they are engaged in acts of problem solving?”

Some fundamental assumptions in this enterprise are:

1. “Problem solving” is used in a deliberately broad way here. It includes a child’s actions in interacting with its parents, a student working on a mathematics problem in class or in the laboratory, and a teacher’s decision-making while teaching a mathematics (or other) lesson. More broadly, I assume that almost all human action is goal-oriented – and that attaining high-priority goals can be characterized as a “problem.”
2. Most human behavior is rational, in the following sense. The actions people take in any particular context are fundamentally aimed at solving problems that are important to them. (These may or may not be the problems they have been “assigned” to solve!) If one is capable of understanding what problem a person is trying to solve at any given time, that person’s actions will often be seen to be rational and consistent. In certain contexts, such as teaching and problem solving (by the standard definition), that consistency in behavior can be strong enough to allow the individuals’ actions to be modeled.
3. In any given context, decision-making is a function of beliefs, goals, and knowledge. In brief outline: an individual’s beliefs, in interaction with the context, shape the formation and prioritization of goals. Given a particular constellation of goals, the individual looks for and implements knowledge that is consistent with his or her belief systems and is designed to satisfy one or more high-priority goals. As goals are satisfied (or not), or as the context changes, new goals take on high priority, and actions are then taken in the pursuit of these goals.

Examples are given to suggest the way in which this theoretical perspective can play out.

Résumé. Résolution de problèmes du berceau au tombeau.

Cet article assez spéculatif propose une mise en perspective théorique globale de la prise de décision et de la résolution de problème "en temps réel". Le domaine est volontairement très large. Mon but est de traiter de la question suivante : "Comment et pourquoi les gens prennent-ils les décisions qu'ils prennent lorsqu'ils se soit engagés dans une activité de résolution de problèmes."

Quelques hypothèses fondamentales de ce travail sont les suivantes :

1. L'expression “Résoudre un problème” est utilisée dans un sens très large. Elle comprend les actions d'un enfant dans ses relations avec ses parents, un étudiant travaillant sur un

problème de mathématique dans une classe ou dans un laboratoire, et la prise de décision par un enseignant lors d'une leçon, par exemple de mathématique. Plus généralement, je suppose que presque toute action humaine est orientée vers la réalisation d'un objectif et que réaliser les objectifs de haute priorité peut être caractérisé comme un "problème".

2. La plupart des comportements humains sont rationnels au sens suivant. Les actions entreprises par un individu dans n'importe quel contexte particulier ont pour but fondamental de résoudre les problèmes qui sont importants pour lui. (Ces problèmes peuvent être, mais ne sont pas nécessairement ceux qui lui ont été "assignés"!.) Si l'on est capable de comprendre quel problème une personne essaye de résoudre à un moment donné, les actions de cette personne paraîtront souvent rationnelles et cohérentes. Dans certains contextes, tels que l'enseignement et la résolution de problèmes (au sens standard du terme), cette cohérence de comportement peut être assez forte pour que les actions de l'individu puissent faire l'objet d'une modélisation.

3. Dans tout contexte, la prise de décision dépend des croyances, des objectifs et des connaissances. En raccourci schématique: les croyances d'un individu, en interaction avec le contexte, modélisent la formation et la hiérarchisation des objectifs. Devant une constellation particulière d'objectifs, l'individu recherche et implémente la connaissance qui est cohérente avec son système de croyances et qui lui permet d'atteindre un ou plusieurs objectifs de première priorité. Lorsque ces objectifs sont atteints (ou ne le sont pas), ou lorsque le contexte change, de nouveaux objectifs acquièrent une priorité élevée, et les actions sont alors dirigées vers la réalisation de ceux-ci.

Des exemples sont donnés pour indiquer de quelle façon cette perspective théorique peut être exploitée.

Mots-Clés. Résolution de problème, rationalité, prise de décision, heuristique, croyance, objectif.

1. Introduction

The purpose of this paper is to outline a very broad theoretical framework for thinking about human decision-making, and to provide suggestive evidence from a number of different domains (childhood behavior, mathematical problem solving, teaching) that the framework can be useful.

I start with a statement about the theoretical enterprise and my goals, not just for this paper, but for the more than thirty years of research it represents, and the decades of research that it proposes for the future. In my 1985 book *Mathematical Problem Solving* I presented a framework for the analysis of mathematical problem-solving behavior. There I claimed that in order to understand why someone is successful or unsuccessful in an attempt to solve a mathematical problem, one must examine that person's (a) knowledge base, (b) use of problem-solving strategies, (c) metacognitive aspects of behavior such as monitoring and

self-regulation, and (d) beliefs. The argument – since confirmed in multiple fields (see, e.g., deCorte, Greer, & Verschaffel, 1996) – was that all these aspects of knowledge and behavior are fundamental determinants of success or failure in problem solving.

I make two observations about that work. First, in it I restricted myself to a discussion of “non-routine” problem solving. I was interested in what people did when they worked on problems that, in some way or other, were new to them – problems that they did not know how to solve. Second, what was lacking from that work was a theory at a level of mechanism. There was no theory of how and why people made the decisions they did – why they chose one option over another, for example. Nor was there a theory of how the various aspects of performance (knowledge, strategies, metacognition, beliefs) interacted with each other.

In the research I have done since then, I have tried to explore those issues. On the surface, that work may look different: I have studied the behavior of tutors working with individual students, and of teachers in the midst of interacting with their classes. My goal has been to explain how and why the tutors and teachers made the decisions they did.

At a deep level, the research on teaching is an extension of the work on problem solving. In problem solving, there is one over-arching task: to obtain a solution to the particular goal or goals the individual has set for him-or-herself. (As we will see below, that goal may or may not be to find a solution to the mathematical problem that the individual has been asked to solve!) I posit that teaching is also a goal-directed activity: the teacher is using his or her knowledge, strategies, and metacognitive skills in the service of trying to achieve some high-priority goals. Those goals are shaped, of course, by the teacher’s beliefs and knowledge. Hence the studies of tutoring and teaching I have conducted over the past two decades are, in fact, studies of problem-solving – but at a level of mechanism, where there is an explicit focus on how and why each decision is made.

My goal in this paper is to begin to unify these two strands of work theoretically – to develop, if you will, a grand theory of problem solving that addresses the question of how and why, and with what success, people make the choices they do as they try to solve problems (that is, to achieve goals they have set for themselves). This is an extremely broad and ambitious goal. It includes not only non-routine problems, but all problems; it includes not only mathematical problem solving but all goal-directed behavior. As such, this paper is a theoretical manifesto. I point to a theoretical synthesis of my prior work and make a plausibility case for it, using some new data and reconsidering extant data. Ultimately, years of analysis will be necessary to see how well these ideas pan out, and to work out the details of a full-blown theory.

By way of preliminaries, let me define problem solving for the purposes of this paper. I noted above that in my earlier work on mathematical problem solving (see, e.g., Schoenfeld, 1985, 1992) I focused on *non-routine* problem solving – on what Hatano (1982) would call “adaptive expertise.” The goal in that work was to describe the kinds of mathematical understandings possessed by people who were good at solving problems that they did not, *a priori*, know how to solve – that is, they did not have a solution method accessible to them when they began to work on the problems. That work showed the importance of heuristic strategies (tools for making progress on difficult problems), metacognition (especially monitoring and self-regulation, for the effective use of the knowledge at one’s disposal), and beliefs. Here I wish to employ a much broader definition of problem solving. For purposes of this paper, a *problem* for an individual at any point in time is something that individual wants to achieve. To put this another way, *solving a problem* will be interpreted as *working toward achieving a high-priority personal goal*. Some of the things that qualify as “problems” in this categorization are: a neonate’s need to be fed; a mathematics task taken seriously by the individual trying to arrive at an answer to it; and, trying to teach a successful lesson on any particular topic. As will be seen below, this definition is deliberately broad – yet, it will allow for some very fine-grained studies of problem-solving behavior.

2. On rationality

This narrative begins with an assertion and a story, to suggest the scope of the ideas being discussed here. The assertion is that most human behavior is fundamentally *rational*, in the following sense: the actions that people take, at any moment, are designed to address problems that are (at that moment) of significant importance to them. As will be seen below, this form of rationality represents a particular form of *internal consistency* on the part of the problem solver; it does not necessarily produce behavior that appears “rational” to an outside observer. I note that this terminology, while problematic to some, does have a long lineage within the cognitive science community: In his first Presidential Address to the American Association for Artificial Intelligence, Allen Newell described a fundamental aspect of his approach to characterizing purposeful problem solving: “The behavior law [to explain an intelligent agent’s actions] is the *principle of rationality*: Actions are selected to attain the agent’s goals” (Newell, 1981, p. 6). What is important to understand in this context is that an individual’s goals are internal and established by that individual. Those goals may or may not be to solve the tasks given to them researchers or teachers, but rather to meet some other high priority (perhaps psychological or social) needs.

Here I will re-tell some stories from my 1985 book *Mathematical Problem Solving*, with an emphasis on the rationality of the behavior described.

One problem that I used in my problem solving research is this:

Problem 1

You are given two intersecting straight lines and a point P marked on one of them, as in Figure 1 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 one of them.




Figure 1: A construction Problem.

Typically I handed students a photocopied sheet of paper with this problem on it; the students made their construction on that sheet. When I gave this problem to one pair of students, however, they took out a blank sheet of paper and began, laboriously, to copy the figure onto that sheet of paper. They used the standard straightedge-and-compass construction for copying an angle, then measured off the distance from the vertex to P . From my perspective at the time (and from that of most people who have seen a videotape of the problem session), their actions seemed a complete waste of time. They certainly didn't help solve the problem – and, it should have been clear that I had more copies of the problem, so that they wouldn't be “spoiling” the problem sheet by writing on it. From that perspective, their behavior hardly makes sense.

Similarly, here is a “Fermi-type” or “back of the envelope” problem that requires little formal knowledge but a bit of ingenuity:

Problem 2

Estimate, as accurately as you can, how many cells might be in an average-sized adult human body. What is a reasonable upper estimate? A reasonable lower estimate? How much faith do you have in your figures?

Here is how I think about the problem. One needs to estimate the size of an “average” cell, the size of an average-sized adult human body, and divide the latter by the former. The estimate of human volume can be “quick and dirty.” If you guess that an average male adult weighs 75 kg, you are certainly within a factor of 1.5; for ease of computation make it 100 kg, and you're still within a factor of 2. That's certainly as good as you need. Estimates of cell size are more tricky. How

big is a cell? Cells were discovered by the use of early microscopes. They couldn't have been very powerful – perhaps between 10 and 100 power. The naked eye can resolve down to 1/10 mm, so a cell might be between 10^{-2} and 10^{-3} cm across. If one assumes, for the sake of simplicity, that an “average” cell is a cube, then there are between 10^6 and 10^9 cells per cm^3 of flesh. The volume of a 100 kg human is roughly 100 liters, so there are perhaps between 10^{11} and 10^{14} cells in that volume. Note that nearly all the error in volume estimation comes from the estimate of cell size – one can be cavalier about the estimate of human volume.

I gave this problem to a number of talented undergraduate mathematics majors, who worked the problem by themselves. Their behavior was oddly consistent. After reading the problem, the students would begin their work by making very detailed computations of body volume. They would approximate the head by a sphere, the arms and legs by either cylinders or sections of cones, and the torso by a circular or elliptical cylinder; they then made close guesses as to dimensions, and detailed computations of volume. Having spent perhaps ten minutes on those computations, they dispatched with cell volume in a matter of seconds – “say 1/100 of an inch” or “it's bigger than an angstrom unit, how about maybe 100 or 1000 angstroms?”

What makes these students' behavior all the more interesting is that once I started having students work the problem in pairs, I never again saw such behavior.

Here is the main point of these stories. From the experimenter's perspective, the students' behavior made little sense. There was a specific mathematical problem to be solved, and the students' attention to irrelevant mathematical detail made no sense. Indeed, from the experimenter's perspective, such behavior can be seen as *irrational*.

There is another point of view, however – that of the student. Consider the first problem, and imagine yourself as an undergraduate. Your professor has asked you to come into his laboratory, and to be videotaped as you work on a problem. In the first case, he hands you a geometry problem. You read it and your mind goes blank. You don't know how to solve it, and you guess that you won't be able to solve it. You don't want to look stupid. You also don't want to really try to solve the problem: if you do try, you'll have to acknowledge to yourself that you failed to solve it, whereas if you don't try very hard, then you can explain away your failure by telling yourself that you never really tried. So: your goal is *not* to solve the problem. Your goal is to exit gracefully from this situation, with your ego and your professor's judgment of you both intact. What can you do? The problem involves geometric constructions. Can you demonstrate geometric knowledge, showing that you *do* know some geometry, while not really trying hard to solve the problem itself? Aha! What if you copy the figure in the problem, using correct geometric procedures? That way you demonstrate some relevant mathematical knowledge,

and you stall for time. Perhaps you will have an inspiration, in which case you can solve the problem. And if you don't, you'll have spent so much time on the construction that, of course, you couldn't possibly have had time to solve it. Hence your ego emerges intact either way.¹

Now consider the "cells" problem. This problem seems to come out of nowhere – one can imagine a student saying (to him-or herself), "I don't know a thing about biology. What in the world can I do?" For the student, the goal may *not* be to engage fully with the problem – this can just reveal his or her ignorance! Rather, the student has a goal similar to the one above – to engage in behavior that appears to be mathematical on the surface, and escape with your ego intact. With this as a goal, what can one do that is mathematically relevant? The problem involves computing either masses or volumes. Aha! You know how to compute the volumes of geometric solids. This involves estimation and the use of mathematical formulas – good mathematical behavior. So, you engage in the careful estimation of human body volume. When you get to the part of the problem that deals with cell volume, you zip through it as rapidly as possible. As a result, you spend 90% of your time "being mathematical." You emerge from the problem session having produced some legitimate mathematics for the professor, and (by virtue of having done something relevant) with your ego intact. In the words of Warren Hinckle (1990), the students succeeded at the following task: "if you have a lemon, make lemonade."

Looked at from this point of view, the students' actions in working both mathematical problems were absolutely and perfectly rational. In both cases, the students entered the laboratory context with a certain set of beliefs: this is who I am, this is what I know and can do, *etc.* In both cases, they were confronted with a task, but in a larger context: a mathematics professor was going to judge their behavior as they worked on the task. In both cases they established goals for themselves, as a function of beliefs and context. The goal-setting depended on their perceptions of the difficulty of the task, their ability to solve it, and the likely reaction of the professor (and themselves) to their efforts. In both cases the primary goal turned out not to be mathematical (i.e., solve the mathematics problem). Instead, the goal was to find a comfortable exit strategy from an uncomfortable situation – to display mathematical behavior, and to leave with one's ego intact. With this goal established, the students searched their knowledge bases. In each case (though with different mathematics, of course) the students found some mathematical behavior in which they could engage – behavior that would have them acting mathematical, displaying some knowledge, and not be seen flailing.

¹ One of the students discussed here later became my research assistant. She told me that this is what she had done.

This was an excellent choice of tactics, given their goals! In sum, their behavior could be seen as *totally rational*, once one understands what their goals (in that particular context, at that particular moment) actually were!

For purposes of this paper, I will stipulate that most human behavior is of that type. That is, if you know what problem an individual is *really* trying to solve (which may or may not be the problem that the person is “officially” trying to solve), then the person’s actions are likely to be rational in the following sense: given what the individual *knows*, the choice of that action represents a plausible mechanism for achieving the person’s *real* goals. The challenge, then, is to understand what problems people are really trying to solve.

(I should note the following. I mentioned this idea in a recent presentation. I was joined after the talk by a psychotherapist, who said that he had much the same idea in his psychotherapy practice. He often worked with patients who had difficulty changing what was clearly dysfunctional behavior. On the surface, the behavior seemed to make no sense. But, he would ask, “what are you getting out of this?” Often, for example, an obviously dysfunctional behavior (e.g., alcoholism) would result in the individual’s getting a great deal of attention from family members. *That* was the major problem being “solved” by the individual – and once that was understood, things fell into place.)

My first premise, then, is rationality of the kind described here.

My second premise is that rational behavior of the type described here can be modeled, in very fine-grained detail – if one has a good sense of the knowledge, goals, and beliefs of the individual whose behavior is being modeled. More specifically, I will argue that the “architecture” of the model is an abstraction of the stories told above. That is:

- an individual enters into a particular context with a particular body of knowledge, goals, and beliefs;
- as events take place in that context, the individual prioritizes goals in response to those events. Thus, for example, a student is given a problem to work and establishes a set of top-priority goals to work toward. (In the cases described above, the top-priority goals involved aspects of self-preservation, and students acted accordingly. In the vast majority of problem-solving sessions I have recorded, the highest priority goal of the problem solver *is* to solve the given problem, and he or she acts accordingly);
- what is considered to be relevant and appropriate knowledge to employ toward achieving those goals is shaped by the individual’s beliefs. For example, a student may react to a question from a teacher one way,

because he or she expects the teacher to expect a formal mathematical argument in response to the question. The same question from a peer might trigger a different (and equally mathematical but informal) response;

- the individual pursues a path toward achieving the high priority goals by choosing and applying relevant and appropriate knowledge, as described in the previous bullet;
- as events unfold, goals can be re-prioritized and new knowledge can be used to achieve the new high-priority goals. Thus, for example, if a problem has been broken up into sub-problems, the individual will work on one or more of the sub-problems. If there is a perception of making progress, the individual may persevere by working on those sub-problems. If there is a perception that things are not going well, the individual may consider alternatives;
- the quality of the decision-making described in the previous bullet is very much a function of the individual's metacognitive skill;
- the process described here is recursive, in the sense that goal prioritization and knowledge selection occur at multiple levels.

I recognize that the preceding is a very abstract description. Thus I shall provide some worked-out examples, in another problem-solving domain – the domain of teaching.

3. Teaching “in the moment” as problem solving: A model

Teaching is problem solving, in the broad sense described above. On any given day, the teacher enters the class with an agenda – a set of (typically multiple) goals. The attempt to meet these goals is an act of problem solving.

My claim is that the teacher's problem solving can be seen as (indeed, modeled as) a function of the complex interaction of teacher's goals, beliefs, knowledge, and decision-making procedures. I begin by elaborating on each of these categories.

Goals

In the spirit of the previous section, I note that these may be quite varied. On the surface, the main goal of a lesson is usually to have students learn a body of subject matter. As Lampert (2001) notes, this is just one of many goals. In the fifth-grade class she discusses, her goal is also to help students learn to work collaboratively; to grow as human beings; to learn to study effectively; and more. Subject-matter goals may include mastery of the particular topic, developing a broad sense of mathematical inquiry, and more. Other teachers may have other, sometimes more


personal or idiosyncratic goals as well: maintaining discipline or avoiding classroom conflict, nurturing particular students, staying on safe ground for themselves with regard to content, or just getting through the day. In any particular situation, some of these goals will have highest priority.

Beliefs (and values)

A teacher's beliefs and values shape the prioritization both of goals and of the knowledge employed to work toward those goals. Does a teacher believe that students can learn from mistakes, for example, or that students should be given clear presentations of correct mathematics? Does the teacher see mathematics as a body of facts and procedures, or as a form of sense-making? Does the teacher believe that a particular student or group of students has the capacity to learn in particular ways? Does the teacher consider oral or written communication in mathematics to be important? Are there (perceptions of) external pressures, like the need for students to do well on exams? All of these beliefs and values serve to determine which goals have highest priority. When a teacher's work is modeled in fine-grained detail (see, e.g., Schoenfeld, in press), the beliefs that need to be delineated include: beliefs about the nature of learning and what supports it; beliefs about teaching; beliefs about students, both individually and collectively; beliefs about what is appropriate and inappropriate for classroom environments; and beliefs about the nature of mathematics, both in general and specifically with regard to the topic(s) currently being studied.

Knowledge

In broad-brush terms, I make the standard cognitive assumptions about knowledge and its organization: our knowledge is organized by way of schemata, which are "triggered" by particular contexts or associations. Readers of this paper, for example, when they see the following diagram,



are likely to think of the Pythagorean theorem and other related mathematics. A teacher's knowledge includes many categories of knowledge: of mathematics; of curriculum; of various pedagogical strategies; of specific student understandings and misunderstandings, and ways to deal with them (also known as *pedagogical*

content knowledge); of recent events in the classroom; of individual students, their (perceived) strengths and weaknesses; and more.

Beyond that, I need to point out that in any given context, some of an individual's knowledge is likely to be more accessible than other knowledge – if you are working geometry problems, much of your geometric knowledge is likely to be “activated,” whereas in another context (say coming across a geometric clue in a crossword puzzle) the activation level of that knowledge is lower, and it may be more difficult to bring that knowledge to conscious awareness. There is a large psychological literature relevant to this point, which I will take as a given for purposes of this paper.

Decision-making

Here is the basic mechanism by which the model works. At any given time, the teacher has a particular set of high priority goals and, most likely, a larger agenda within which those goals are situated. To make the situation concrete, imagine a teacher about to begin a lesson. The teacher expects to conduct routine introductory business, go through the day's assigned homework, and then turn to new material. The teacher begins the class session, as intended, by taking roll, and then asking if the students have any questions. The question-and-answer session goes on for a few minutes, with the teacher responding to questions according to a determination of the importance of the issues raised, and the time it will take to answer them. In some cases, the teacher may defer answering a question (“see me after class”); in others, a question may lead to a long discussion. The choices are made on the basis of the teacher's beliefs and values. (This issue is important, this one is not; how much time can I spare; do I want to answer this one now; do I want to answer this one now in public; *etc.*) When students run out of questions or the teacher decides that enough time has been spent on this activity, the activity is brought to a close and the class proceeds to the next part of the intended agenda – reviewing homework.

Suppose the teacher has assigned a collection of problems that the students are supposed to have worked. The teacher has many choices about how to review these. These include a series of possible classroom routines, all of which the teacher could implement:

- having students volunteer or be called on, and work through some or all the problems in sequence;
- leading a “Socratic” discussion of some or all of the problems
- presenting solutions at the board, and asking students for comments or questions;

- asking students to identify problems that caused them difficulty, and focusing only on them;

and more. The teacher's choice may be made on the spot, as a function of the time remaining in the class, or of other things the teacher wishes to accomplish that day. Once that choice is made, the top-level goal has been established: to go through the homework in the fashion chosen. The first subgoal is to work through the first problem. If nothing unusual happens, this task is accomplished and the class moves on to the next problem. At any given time, however, something can happen to change plans. Here are some examples.

a. The teacher has chosen for students to work through the problems, but a particular student's explanation seems incoherent and is confusing the class. With a particular set of beliefs, the teacher might decide to step in and demonstrate a correct solution. With a different set of beliefs, the teacher might lead the student through a solution. With yet a different set of beliefs, the teacher might air the student's (mis)understandings, and use them as a vehicle for addressing such issues with the whole class. Note that there are various ways to address these issues, as discussed in the list above. This is a matter of knowledge and choice – knowledge and ability to implement the options, and choice, with regard to values and beliefs, subject to the constraints of time, *etc.*

b. In the middle of a routine problem solution, the student makes an error indicating a fundamental misconception – one that may be shared by other members of the class. For example, the student may write

$$(a + b)^2 = a^2 + b^2.$$

The teacher then faces the same kinds of decisions as discussed in (a).

c. The class may become restive, at which point the teacher may decide either to persevere or to embark on a new activity.

d. The middle of a discussion, a student may make a comment that contains the seed of an interesting mathematical idea – but one that would take 10-15 minutes to work through with the class, causing a significant disruption in the teacher's planned agenda.

Note that the teacher has many choices, among them:

- i. The teacher may say "That's interesting, I'll talk to you about it after class";

- ii. The teacher may say “We’ll discuss this in tomorrow’s class” and plan to do so;
- iii. The teacher may say “Your question raises an interesting issue. Let me explain it to you” and take 2 minutes to do so;
- iv. The teacher may invite the class to work through the issue, taking 10-15 minutes to do so.

Which of these choices the teacher makes will depend on the teacher’s beliefs about what is important, how comfortable the teacher is about implementing any of these choices (a function of what knowledge is available at that moment), and his or her perception of the value of that choice compared to the cost (in time and disruption of the lesson agenda).

I note that this is not a hypothetical – we will work through such an example in the next section.

4. Using the model to characterize the actions of specific teachers – a summary and one brief worked-out case

Before proceeding with specific examples, I want to provide some context for what follows.

First, as noted above, I have argued that most human behavior is *rational* in the sense that people act in ways designed to meet goals that are important to them. Indeed, in example (d) above, I suggested that their choice of strategies or knowledge may often be made on the basis of what will “yield” the best results at least cost. This characterization may make it sound as though humans are acting like computers, mechanically establishing goals and making deliberate choices. I do NOT mean this! Most of the time, people make instantaneous decisions based on what “feels right.” What I am suggesting is that a post hoc analysis will reveal that their decisions are *consistent with* rational choices – and that such consistency can be modeled.

Second, I want to stress that the model of problem solving described here is a model of decision-making *in action*. There is much more to teaching than is described here – there is planning, for example; there is a teacher’s learning, over the course of his or her career. The model does not address those things directly. What it does address is what the individual is doing at the moment, while teaching.

Third, my descriptions have been very general up to this point. They have been deliberately so – I am conjecturing, after all, that the model applies to all problem solving! In the examples that follow, however, I will indicate that the model can be

made very specific and fine-grained – so that it applies to explain, on a line-by-line basis, the decisions that a teacher makes in the middle of teaching.

Cases studied, and partially worked-out examples

Over the past decade my research group has studied a number of cases of teaching, in very fine-grained detail. The goal, as discussed above, has been to explain every action the teacher makes while teaching, as a function of the teacher's knowledge, goals, beliefs, and decision-making. Because of the level of detail involved, the typical papers are very long: the two main papers, describing lessons taught by Jim Minstrell and Deborah Ball, are over 100 pages long each. Here I shall simply summarize some of the main points, and give some brief examples. The detail can be found in Schoenfeld (1998, 1999, 2000 in press) and Schoenfeld, Minstrell, and van Zee (2000).

Case 1, a beginning teacher with traditional content

The first case we modeled was of a beginning teacher teaching a rather traditional lesson. What we saw was that the teacher reached a “roadblock” in his lesson, where he was simply unable to continue as planned. The key aspects of the situation studied were this:

- A. The teacher had a plan for the lesson that was somewhat under-specified. The plan was to have students work through the algebraic simplification of expressions such as (x^5y^3/x^3y^2) , and, on the basis of that experience, to have them conclude that

$$x^0 = x^{5-5} = x^5/x^5 = 1.$$

His plan was to call upon students who had obtained the right answer, have them explain how they arrived at that answer, and elaborate on the answer for the class.

- B. Because he was inexperienced, he did not know to anticipate a substantial degree of confusion when students “cancelled” x's in the expression

$$\frac{x \ x \ x \ x \ x}{x \ x \ x \ x \ x}$$

and saw “nothing” when they were done:

$$\frac{x| \cdot x| \ x| \ x| \ x|}{x| \ x| \ x| \ x| \ x|}$$

He had no back-up plan for having the students deal with the subject matter.

- C. He believed (as many beginning teachers in the U.S. believe) that it is inappropriate to simply “tell” students the correct answer – that a “constructivist” teacher must work with ideas generated by the students.

The factors A, B, and C combined to create the problem for the teacher. He wrote the problem x^5/x^5 on the blackboard for the students to work, expecting at least one student to get the right answer. When none of the students called out the right answer, he could not implement his intended strategy of elaborating on a student explanation. He searched his knowledge base for an alternative strategy, but none was available. Since he believed that he should not “tell” the students the answer, he was stuck. (If you see the videotape of the class, you can see him slump at the blackboard at this point – he is unable to go further.)

Note that if the teacher had a different belief system, which allowed him simply to lecture the content to the students, he would have been able to proceed perfectly well. He had the relevant knowledge, but his beliefs about how to teach prevented him from using it.

Case 2, an experienced teacher with novel content

In Schoenfeld (1998) I present the very detailed examination of a full lesson taught by Jim Minstrell, a high school teacher-researcher widely recognized for his skill. Minstrell had created a lesson to help his students understand the issues one confronts when gathering and analyzing data. The issues in the lesson are, which data does one use (e.g., does one include or exclude “outliers”), and how does one compute the “best value” (e.g., would mean, median, or mode be a better choice as a measure of central tendency) for the situation?

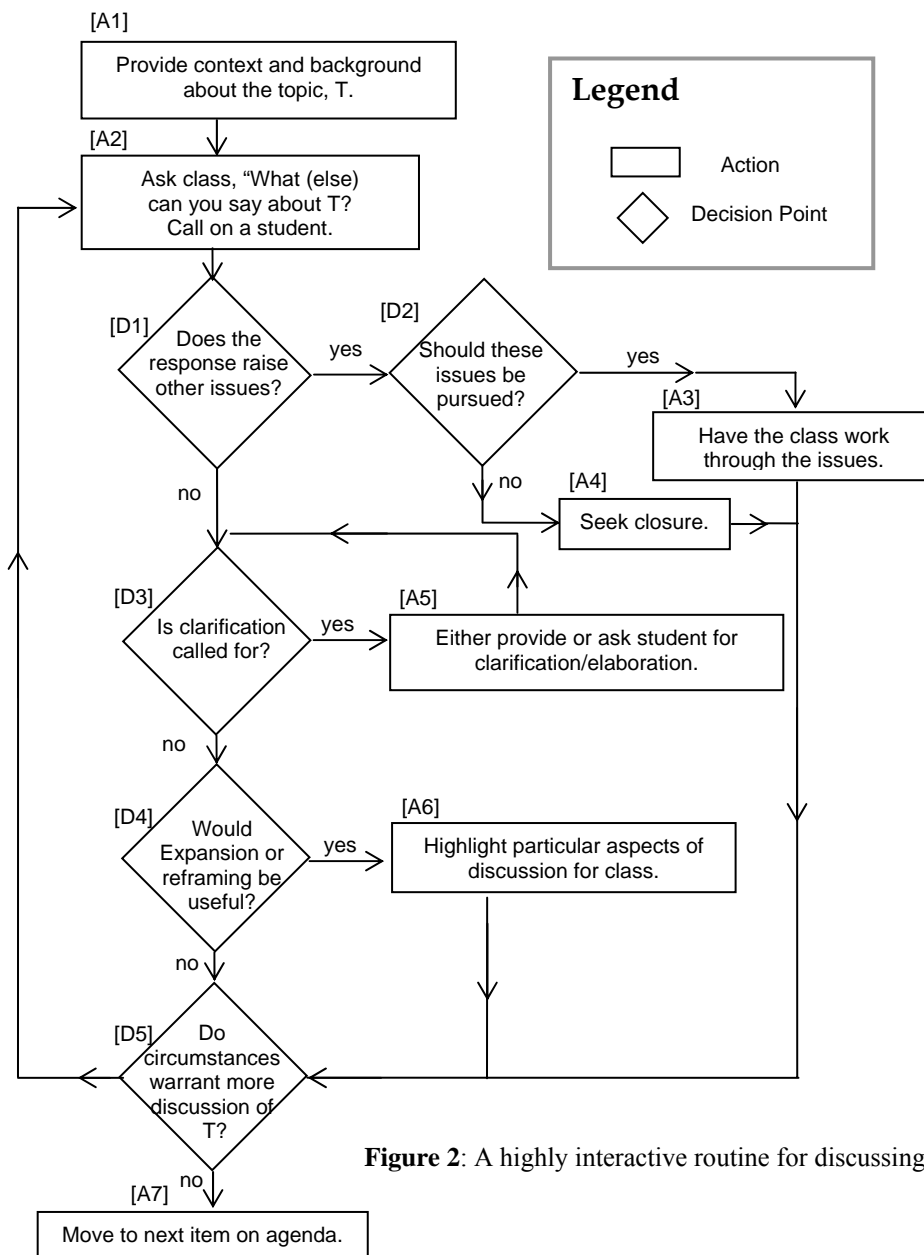


Figure 2: A highly interactive routine for discussing a topic.

Minstrell has developed a unique interactive style, in which he rarely makes declarative statements, but instead asks questions and works with the answers

(correct or not) provided by the students. He also has a particular kind of classroom routine that he employs for soliciting comments from students – one that, as it turns out, is also employed by teachers such as Deborah Ball and myself. That routine is given in Figure 2. (See next page.)

The previous day Minstrell had asked a group of 8 students to measure the width of a table. They had obtained the following values:

106.8; 107.0; 107.0; 107.5; 107.0; 107.0; 106.5; 106.0.

The question before the class was, “what is the best value to use for the width of the table?”

Minstrell’s plan was to start out the class with “routine business” (any questions the students might have about class organization, *etc.*), and then engage the class in the issue of “best value.” That discussion would have three parts: the consideration of which data to include (both in general and in this case), the consideration of how to get the “best value” for the numbers they had, and a discussion of “precision” – how to describe the magnitude of the possible error regarding the “best value.” As always, he planned to interact with students using his questioning strategy. He also had some very high priority goals for the class, among them:

- to foster the students’ understanding of science as a sense-making activity;
- to have students become comfortable asking questions;
- to model the process of inquiry.

The lesson begins with “routine business.” Minstrell asks students if they have any questions regarding the conduct of the course, grading policy, *etc.* After these preliminaries have been taken care of, he turns to the question of analyzing data.

Minstrell asks the students why, in general, they might consider some or all of the data. A student responds by saying “eliminate [the] highest and lowest,” which Minstrell pursues by asking if the students can explain where one might do that. Among the responses are that high and low scores are sometimes dropped in sporting events.

Minstrell clarifies this, and then asks the question again: “OK. What’s another way at going at taking some of the numbers and not all of them?” This time a student answers in terms of “extreme values,” and Minstrell pursues this – calling atypical values “outliers” and explaining that one might be suspicious of such numbers. He continues, “OK? What was another one?” and spends some time discussing the credentials of those who took the measurements – some numbers might be seen as more “trustworthy” than others because they were gathered by people with greater expertise. He tries again: “OK? Any other reasons you can think of to only take

some of the numbers?” There is no response, and he has no reason to bring other suggestions into the conversation, so he brings this part of the lesson to a close. (In the terms of the model, he has met the goal of working through “which numbers should we consider; thus the next goal, “how should we combine them” becomes highest priority.) In doing so that he has made consistent use of his knowledge base, selecting and implementing the iterative strategy described in Figure 2. The routine described in Figure 2 corresponds, on a line-by-line basis, to his actions over this part of the lesson.

Minstrell now begins the second part of the data analysis discussion. He asks,

“So now we've got some numbers there, what are we going to do with those numbers? What's one thing that we might do with the numbers?”

A student says “Average them” and Minstrell, consistent with his questioning style, asks “Now what do you mean by ‘average’ here?” He elaborates on the definition of *mean*, and then returns to the question:

“Any other suggestions for what we might do? So we can average them.
[8 second pause]
Any other suggestions there for what we might do to get a best value?”

A student says “You've got a bunch of numbers that are the same number,” a statement Minstrell pursues with his questioning style. The result is a clarification of the *mode*.

Consistent with the implementation of the routine in Figure 2, Minstrell returns to the top-level question: “Anybody think of another way of giving a best value?” A student provides an unexpected response:

“This is a little complicated but I mean it might work. If you see that 107 shows up 4 times, you give it a coefficient of 4, and then 107.5 only shows up one time, you give it a coefficient of one, you add all those up and then you divide by the number of coefficients you have.”

Up to this point, Minstrell has been implementing the routine in Figure 2 smoothly.² One can think of each iteration of the routine as corresponding to the establishment of a subgoal – “let’s hear if the students have another idea, and work through it” – and the completion as meeting that subgoal. This new suggestion

² Note, again, that I am not claiming that Minstrell is following this routine consciously – I am merely claiming that his behavior is consistent with this routine. I uncovered this specific routine when doing an analysis of Deborah Ball’s teaching (Schoenfeld, 2002; in press). When I described it to Deborah, her reaction was “That’s interesting. I wasn’t doing it consciously, but now I can see that I use that routine quite a lot.”

changes things, however. It is “outside the space” of things Minstrell has set up. The issue in terms of modeling: Can we say, in a principled way, what he will do? (Note that we have presaged this situation in the general discussion. A priori, Minstrell could respond in any of a number of ways, from (i) telling the student he’ll talk to her about this idea after class, to (iv) inviting the class to work through the issue, perhaps taking as long as 10-15 minutes to do so.

The model works as follows. In terms of Figure 2, Minstrell has asked question [A2] and the student’s response, [D1], does indeed raise other issues. Hence Minstrell must decide (in [D2]) whether and how to respond. The student’s comment is relevant in terms of subject matter, and represents a legitimate attempt at sense-making on her part. Recall that Minstrell has certain top-level goals for the class, among them that the students see science as a form of sense-making. He wants the students to feel free to raise relevant issues (that is, he wants the environment to be “risk-free” when students make conjectures or inquiries). He knows that he can work to create the right kind of environment (and encourage other students to take the same risks) by responding positively to the student’s question. Hence, Minstrell will choose option (iv) – he will invite the class to work through the issue, even though the cost will be a temporary deflection of his agenda for the lesson. Once this decision is made, the next question is how he will pursue the issue. This too is a matter of beliefs, values, and knowledge. Minstrell has more than enough content knowledge to pursue the issue. He favors the questioning strategy that invites student input, rather than telling; he also wants to make sure all the students understand the issue before pursuing it. Thus, (the model of) Minstrell will ask the student to repeat or clarify what she has said, and then ask the class for ideas or suggestions. This is, in fact, what the real Minstrell did.

Working through the suggestion with the class – showing that one interpretation of what the student said led to an inappropriate formula, but another interpretation led to the formula for what we call the weighted average – did indeed take some time. When the discussion was concluded, Minstrell returned to his original agenda. He had the students discuss the median (the third measure of central tendency, which had not yet been raised). The discussion of median finished the second main chunk of the lesson (how to choose the “best value”), at which point he could turn to a (condensed) discussion of “precision.”

In sum, Minstrell’s behavior, even in unexpected circumstances, was entirely “rational” (consistent with his goals and values), and his decision-making during the full hour of class could be explained on a line-by-line basis. The model works exactly as described in general. At any moment Minstrell has certain goals, and he sorts through his knowledge base to find an approach that is consistent with those goals. As events proceed, some goals are met, or new goals emerge because of

contingencies. When this happens, goals are re-prioritized, and actions consistent with the new high priority goals are taken.

Case 3, An experienced teacher and an “emergent” lesson

The third body of instruction discussed here is a segment of a lesson taught by Deborah Ball, which has become rather famous as the “Shea number” tape. (See Schoenfeld, in press, for detail.) Ball had been teaching a third-grade class. The previous day she had had the class meet with a group of fourth graders (her previous year’s class) to discuss some of the mathematical issues that had emerged in both classes as they considered the properties of even and odd numbers. Some troubling issues had arisen for the students – for example, is the number zero even, or odd, or “special” (not fitting into either category). Some students argued that zero is even, some that it is special. The issue had not been resolved.

Ball starts the class by asking her students reflect on their experience the previous day. Her intention, at least in part, is to have them “go meta” – to reflect on how the meeting shaped their thinking. When a first student comments, Ball interacts with her and sums up: “so you thought about something that came up in the meeting that you hadn’t thought about before.” After another interaction between two students, she points out that some issues (e.g., whether zero is even or odd) take a long time to figure out – that even the fourth graders hadn’t resolved it yet! Then, when a student makes the following comment:

“Um, first I said that um, zero was even but then I guess I revised so that zero, I think, is special because um, I– um, even numbers, like they they *make* even numbers; like two, um, two makes four, and four is an even number; and four makes eight; eight is an even number; and um, like that. And, and go on like that and like one plus one and go on adding the same numbers with the same numbers. And so I, I think zero's special”,

Ball makes a rather unusual move:

“Can I ask you a question about what you just said? And then I'll ask people for more comments about the meeting. Were you saying that when you put even numbers together, you get another even number, or were you saying that all even numbers are made up of even numbers?”

This is a striking intervention, in that it derails Ball’s announced reflective agenda. The class spends a substantial amount of time discussing the issue, and it is not easy to return to reflections. People who have seen the tape have been very surprised at Ball’s move, arguing that it makes no sense.

The issue here: does it make sense? Ball is a highly accomplished teacher. Why would she do such a thing?

The answer depends on knowing the history of the classroom discussions that took place prior to the meeting of the third and fourth grade classes, and on knowing Ball's agenda. Ball had planned for the reflections on the previous day's meeting to take a few minutes at the beginning of class. After that, she planned to return to the main line of work the class had been pursuing – discussions about the properties of even and odd numbers. Students had noticed that doubling gave rise to even numbers. (The class's working definition was that a number was even if you could divide it into two equal piles without leaving anything over. Hence doubling produced even numbers.) Some students had conjectured that the sum of any two even numbers was even. Another conjecture that had been aired was that any even number was the double of an even number – that is, that all even numbers “come from” other even numbers, in the same way 4 “comes from” 2 and 8 “comes from” 4.

For Ball, understanding what her students think is always a key to a successful lesson – so the success of the latter part of her planned lesson, the exploration of student conjectures, depended in part on understanding what her students believed about combinations of even numbers. Did this student (and others) believe that every even number can be written as the double of another even number, or only that the double of every even number is also even? If they believed the former, the lesson would evolve differently – and this was important. Hence Ball decided to make a brief, announced *detour*, which she signaled by saying she was going to ask about what the student had just said, and then return to her agenda by “ask[ing] people for more comments about the meeting.” Whatever one's judgment about the appropriateness or wisdom of this move, the fact is that it is consistent with Ball's beliefs about what is important (understanding her students' understandings) and with her goals and agenda (spending most of the class period facilitating a conversation about the students' conjectures about the properties of even and odd numbers). Moreover, it is clear she expected the exchange to be over quickly, so that the cost of the detour would be small. Under these circumstances, her decision can be seen as fundamentally rational in the sense that I have discussed. Moreover, a cost-benefit analysis of the cost of asking the question (a brief disruption to the flow of the argument) versus the benefits (setting the planned discussion on a more stable base) shows it to be a reasonable, though not obvious, choice. This a model of Ball's knowledge, goals, and beliefs reveals this choice to be both rational and within the realm of possibility.

In sum, this kind of analytic model produces behavior that is entirely consistent with a broad range of teaching – all of which can be seen as problem-solving behavior.

5. Using the model as a model of mathematical problem solving

I will now argue that the theory outlined above also serves as a theory of mathematical “problem solving in the moment” – a theory that serves to explain how and why people do what they do as they are engaged in solving mathematics problems.

A prefatory comment is appropriate here. I am about to revisit some of the data from my 1985 book *Mathematical Problem Solving*. The question is, how do things differ in his interpretation?

In that book, I offered what I called a *framework* for the analysis of mathematical problem solving behavior. I described four categories of mathematical knowledge and behavior:

- resources (the knowledge base);
- heuristic (problem-solving) strategies;
- “control” (Monitoring and self-regulation, aspects of metacognition);
- beliefs.

(These were joined in 1992 by the category of Practices, the consistent activity patterns of a particular intellectual or other community).

My argument was that if you wanted to understand someone’s success or failure in a problem solving attempt, you needed to examine all of these categories. That is, any one of these (the presence or absence of particular knowledge; access or lack of access to heuristic strategies; effective or ineffective metacognitive decision-making; productive or counter-productive beliefs and practices) could provide the reason for an individual’s success or failure as he or she tried to solve a problem. Moreover, I argued that these categories were sufficient for explanations – that success or failure could be explained in these terms.

What was *missing* in this approach was a sense of how all these things fit together – a description of *mechanism*. How did the categories interact with each other? Why did people do what they did when they were in the midst of a problem solving attempt? There were suggestions of the interactions, specifically in the ways that beliefs served to prioritize knowledge. For example, I argued that students who believe that “proof has nothing to do with discovery or invention” would fail to access some relevant proof-related knowledge when they were working construction (“discovery”) problems, even though they could clearly be shown to have that knowledge. But, a theory is more than that.

I suggest that the description of people's decision-making in the act of problem solving described here now has the potential to be a theory of problem-solving-in-action. To recap, the key elements of the theory are

- knowledge;
- goals;
- beliefs;
- decision-Making³,

The basic idea is that an individual enters *any* problem solving situation with particular knowledge, goals, and beliefs. The individual may be given a problem to solve – but as we saw early in this paper, it is not necessarily the case that solving that problem will become *the* problem the individual sets out to solve! Thus, what happens is that the individual establishes a goal or set of goals – these being the problems the individual sets out to solve. The individual's beliefs serve both to shape the choice of goals and to activate the individual's knowledge – with some knowledge seeming more relevant, appropriate, or likely to lead to success. The individual makes a plan (often establishing subgoals, *etc.*) and begins to implement it. As he or she does, the context changes: with progress, some goals are met and other take their place. With lack of progress, a review may suggest a re-examination of the plan and/or re-prioritization of goals. When unexpected events happen (e.g., new information becomes available), a re-prioritization also occurs. This cycle continues until there is (perceived) success, or the problem solving attempt is abandoned or called to a halt.

In what follows I am going to re-visit a problem solving session described in my 1985 book, and re-interpret what happened. Time and space do not permit me to do the kind of exhaustive analysis for this paper that I did in the teacher-model work (e.g., Schoenfeld, 1998; in press), so this analysis is still on the speculative side. I will assert at this point that I am confident that with enough time, I could do a much more detailed analysis.

The problem-solving episode in question, which is given in full in the Appendix, was discussed in Chapter 9 of *Mathematical Problem Solving* (Schoenfeld, 1985). In the discussion I focused largely on aspects of metacognition. The key point of the analysis was that the problem solver managed to terminate a number of fruitless

³ The relationships between the old and new categories are straightforward. Resources and strategies (and some practices) are part of the knowledge base, beliefs remain much as they were (but prioritize both goals and knowledge), and metacognition becomes part of decision-making, which includes the prioritization of goals. Success or failure will still depend on the efficacy of the knowledge base, the appropriateness of the individual's beliefs, and the quality of decision-making.

attempts to solve the problem, thereby giving himself enough time to find a correct solution. I wrote that his solution attempt was “an illustration of the way that executive [metacognitive] skills can make a positive contribution to problem solving performance.”

I suggest that the reader read through the Appendix, then return here for a narrative description.

In line 1 (and beyond) it is clear that GP (the problem solver) does establish finding a solution to the given problem as his major goal. He acknowledges (line 2) not knowing where to start on the problem, and then explicitly employs a heuristic strategy (drawing a diagram that looks close to correct, in the hope of gaining insight) in line 5. The figure does indeed suggest an approach – in line 6 he notes that the two triangles are similar, and that it may be possible to determine the answer analytically. This triggers more knowledge and the establishment of a subgoal – solve analytically for the size of the altitude of the smaller triangle. He does so in line 9. Then, he has a second sub-problem – how to construct a line that has the value he has found analytically, $A/\sqrt{2}$. This triggers an explicit memory search (line 12), which is partly successful – he remembers how to construct $\sqrt{2}$ (line 16), then $\sqrt{2}/2$ (line 17), and ultimately, in lines 18-19, he constructs $A\sqrt{2}/2 = A/\sqrt{2}$. He gets there by doing a series of successively more complex constructions, each one a new subgoal established after the preceding one has been met.

In line 22 he turns to the second (and more difficult) part of the problem. He spends lines 24-27 exploring the problem (again, a good heuristic strategy – as Pólya says, “first, you have to understand the problem”). In line 28 he poses a possible misdirection, trying to apply his solution to the first part of the problem inductively. Here metacognition and decision-making kick in: he decides (line 33) that the approach is not profitable. This calls for re-selecting a top-priority subgoal. He begins working on the problem of constructing the top triangle, with area $1/5$ that of the original triangle T . Using the same approach that he used to solve the first part of the problem, he makes some progress, and in line 40 determines an algebraic expression that he needs to construct: $\sqrt{3}/\sqrt{5}$. Since an expression with two roots is too complex (again, a goal-directed heuristic), he re-expresses this as $\sqrt{15}/5$. This raises another issue, whether $\sqrt{15}$ or $\sqrt{15}/5$ is constructible (line 42). It calls for another knowledge search (lines 43-48). He realizes firmly that division by 5 is not a problem, so the solution to the problem hinges on his ability to construct $\sqrt{15}$ (line 48). He engages in some more conscious memory search (lines 49-54) and finds an appropriate approach, passing by some unprofitable ideas (line 53) and ultimately settling in on a correct approach (lines 55 and 56).

This brief narrative suggests the way in which the theory works, providing a potentially complete description of the problem solving session at a level of mechanism – saying how and why the problem solver did what he did, and to what effect. GP’s beliefs about himself and about mathematics are clearly relevant: he starts out with the assumption that he can tackle problems like this, and work his way through them. In numerous places, his decision-making facilitates his solution – in the establishment of goals and subgoals, in effective monitoring and self-regulation (lines 33, 54), in goal-directed memory search (e.g., lines 46-47) and in the selection of appropriate heuristic strategies (lines 5, 31). Moreover, we see the ways in which things interact: his use of the heuristic “draw a diagram (of the goal state)” triggers the recognition of similar triangles, which then suggests a solution path that had not been apparent beforehand.

In short, this kind of approach and interpretation suggest that if one knew enough about GP’s knowledge, goals, beliefs, and decision-making, one could model this solution down to a very fine level of detail. As such, this would be a model in substantiation of my broad theoretical claim – that nearly all problem solving (in the moment) can (a) be seen as rational, and (b) can be modeled as a function of individuals’ knowledge, goals, beliefs, and decision-making. I am confident that all of the problem solving protocols discussed in *Mathematical Problem Solving* can be re-analyzed this way.

6. A tongue-in-cheek example from the cradle

I hypothesize that this theoretical perspective can be applied to characterize the problem-solving activities of the very young. I can not resist a *reductio ad absurdum* here, but I think there is some truth to it.

Consider a hungry week-old infant. That child has one overriding goal: food! And, that child has one strategy in its knowledge base: cry! Typically, the strategy works, although sometimes with delay – the infant’s mother may not immediately identify the cause of the child’s discomfort. As the child gets older, its collection of strategies gets larger – eventually, for example the child can say “mama” and “papa.” Now when it is hungry, it may cry – but it may also call a specific parent! (This is knowledge at work). As its vocabulary grows, it may know how to identify the sources of its discomfort, *and* who is likely alleviate them – hence choosing one parent over another, and asking specifically for food or drink. Hypothetically, a rather simple model could describe a young baby’s actions; as the child developed, increasingly complex knowledge, goal-setting, and decision-making (shaped by the child’s evolving beliefs) could characterize the development of the child’s problem solving skills. Hence it might be possible, at least theoretically, to characterize problem solving from cradle to grave.

7. Brief Discussion


As indicated in the introduction, this paper is a theoretical manifesto. I claim that it is possible to unify the earlier work I conducted on the solution of non-routine mathematical problems (Schoenfeld, 1985, 1992) with more recent work modeling teachers' decision-making (Schoenfeld, 1998, 1999, 2000, in press; Schoenfeld, Minstrell, and van Zee, 2000). If this effort is successful, it will provide a theoretical mechanism for characterizing a very large part of human goal-directed activities⁴. Time will tell whether this attempt will be successful. But, I hope to have provided enough evidence to convince the reader that the attempt is plausible and worth undertaking.

⁴ Some percentage of human behavior is random, of course. But, especially when one is acting in familiar contexts, much behavior is rational in the sense that I have described.


Appendix

The full text of a mathematics faculty member's attempt to solve a problem. Taken (with my permission) from Schoenfeld, 1985.

1. (Reads problem): You are given a fixed triangle T with Base B . Show it is always possible to construct, with ruler and compass, a straight line parallel to B that divides T into two parts of equal area. Can you similarly divide T into five parts of equal area?




2. Hmm. I don't know exactly where to start.
3. Well, I know that the ... there's a line in there somewhere. Let me see how I'm going to do it. It's just a fixed triangle. Got to be some information missing here. T with base B . Got to do a parallel line. Hmmm.




4. It said the line divides T into two parts of equal area. Hmmm. Well, I guess I have to get a handle on area measurement here. So what I want to do ... is construct a line ... so that I know the relationship of the base ... of the little triangle to the big one.
5. Now let's see. Let's assume I draw a parallel line that looks about right, and it will have base little b .

6. Now, those triangles are *similar*.




7. Yeah, all right then, I have an altitude for the big triangle and an altitude for the little triangle so I have little a is to big A as little b is to big B . So what I want to have happen is $\frac{1}{2}ba = \frac{1}{2}AB - \frac{1}{2}ba$. Isn't that what I want?

8. Right! In other words I want $ab = \frac{1}{2}AB$. Which is $\frac{1}{4}$ of A times [mumbles; confused] $(1/\sqrt{2}) \times A \times (1/\sqrt{2}) \times B$.
9. So if I can construct the $\sqrt{2}$, which I can! Then I should be able to draw this line ... through a point which intersects an altitude dropped from the vertex. That's little $a = A/\sqrt{2}$, or $A = a\sqrt{2}$, either way.
10. And I think I can do things like that because if I remember, I take these 45-degree angle things, and I go 1, 1, $\sqrt{2}$.
11. And if I want to have $a \times \sqrt{2}$... then I do that ... mmm. Wait a minute ... I can try to figure out how to construct $1/\sqrt{2}$.
12. OK. So I just gotta remember how to make this construction. So I want to draw this line through this point and I want this animal to be - $(1/\sqrt{2}) \times A$. I know what A is, that's given, so all I gotta do is figure out how to multiply $1/\sqrt{2}$ times it.
13. Let me think of it. Ah huh! Ah huh! $1/\sqrt{2}$... let me see here ... ummm. That's $\frac{1}{2}$ plus $\frac{1}{2}$ is 1.




14. So of course if I have a hypotenuse of 1 ...
15. Wait a minute ... $(1/\sqrt{2}) \times (\sqrt{2}/\sqrt{2}) = (\sqrt{2}/2)$... that's dumb!
16. Yeah, so I construct $\sqrt{2}$ from a 45, 45, 90. OK, so that's an easier way. Right?
17. I bisect it. That gives me $\sqrt{2}/2$. I multiply it by A ... now how did I used to do that?
18. Oh heavens! How did we used to multiply times A ? That ... the best way to do that is to construct A ... A ... then we get $\sqrt{2}$ times A , and then we just bisect that and we get A times $\sqrt{2}/2$. OK.


19. That will be ... what! ... mmm ... that will be the length. Now I drop a perpendicular from here to here. OK, and that will be ... ta, ta ... little a .




20. So that I will mark off little a as being $A\sqrt{2}/2$. And automatically when I draw a line through that point ... I'd better get $\sqrt{2}/2$ times big B . OK
21. And when I multiply those guys together I get $(2/4)AB$. So I get half the area ... what? ... yeah ... times $1/2$ - so I get exactly half the area in the top triangle, so I better have half the area left in the bottom one. OK.




22. OK, now can I do it with 5 parts?
23. Assuming 4 lines.
24. Now this is going to be interesting because these lines have to be graduated ... that ...
25. I think, I think, rather than get a whole lot of triangles here, I think the idea, the essential question is can I slice off ... $1/5$ of the area ... hmmm ...
26. Now wait a minute! This is interesting. Let's get a ... How about 4 lines instead of ...
27. I want these to be ... all equal areas. Right? A_1, A_2, A_3, A_4, A_5 , right?
- 28.




29. Sneak! I can ... I can do it for a power of 2. That's easy because I can just do what I did at the beginning and keep slicing it all the time.
30. Now can I use that kind of induction thought?
31. I want that to be $2/5$. And I want that to be $3/5$. (pointing to relevant regions)
32. So let's make a little simpler one here.



33. If you could do that then you can construct $\sqrt{5}$. But I can construct $\sqrt{5}$ to 1 ... square root of 5, right?
34. So I can construct ... OK. So that certainly isn't going to do it. No contradiction ...
35. Now, I do want to see, therefore, what I have here.
36. I'm essentially saying it is possible for me to construct it in such a way that it is 1,2,3,4,5, $1/5$ the area ... OK.
37. So little a times little b has got to equal $1/5 AB$. So I can certainly chop the top piece off the area and have it be $1/5$. Right? Right?




38. Now the first part of the problem, I know the ratio of the next base to draw ... because it is going to be $\sqrt{2}$ times this base. So I can certainly chop off the top $2/5$.



39. Now from the first part of the problem I know the ratio of the top ... uh, OK, now this is $2/5$ here, so top $4/5$. OK. All right. So all I gotta be able to do is chop off the top $3/5$ and I'm done.

40. It would seem now that it seems more possible ... let's see ...



41. We want to make a base here such that little a times little b is equal to ... the area of this thing is going to be $3/5$... $3/5 AB$... in areas, right! And that means little a times little b is $[(\sqrt{3}/\sqrt{5})A][(\sqrt{3}/\sqrt{5})B]$. OK, then can I construct $\sqrt{3}/\sqrt{5}$? If so then this can be done in one shot.

42. Well let's see. Can I construct $\sqrt{3}/\sqrt{5}$? That's the question. $\sqrt{3}/\sqrt{5} \times \sqrt{5}/\sqrt{5} = \sqrt{15}/5$.

43. $\sqrt{15}, \sqrt{15}$. Wait a minute. $\sqrt{15}/5$. Is $\sqrt{15}$ constructible? $\sqrt{15}$ is ...


44. It is $\sqrt{16-1}$. But I don't like that. It doesn't seem the way to go.

45. $16^2 - 1^2$ equals ... [expletive deleted]

46. Somehow it rests on that.


47. [Expletive] If I can do $\sqrt{15}$. Can I divide things and get this?

48. Yeah, there is a trick! What you do is lay off five things. One, two, three, four, five. And then you draw these parallel lines by dividing them into fifths. So I can divide things into fifths so that's not a problem.



49. So it's just constructing $\sqrt{15}$, then I can answer the whole problem.

50. I got to think of a better way to construct $\sqrt{15}$ than what I'm thinking of ... or I got to think of a way to convince myself that I can't ... ummm ... x^2 ... 15.
51. Trying to remember my algebra to knock this off with a sledgehammer.
52. It's been so many years since I taught that course. It's 5 years. I can't remember it.
53. Wait a minute! Wait a minute!
54. I seem to have in my head somewhere a memory about quadratic extension.
55. Try it differently here. mmm...
56. So if I take a line of length 1 and a line of length ... And I erect a perpendicular and swing a 16 [he means a $\sqrt{16}$, or 4] here. Then I'll get $\sqrt{15}$ here, won't I?



57. I'll have to, so that I can construct $\sqrt{15}$ times anything because I'll just multiply this by A and this by A and this gets multiplied by A divided by 5 using that trick. Which means that I should be able to construct this length $[A\sqrt{3}/\sqrt{5}]$ and if I can construct this length then I can mark it off on here [the altitude to from the top vertex to B] and I can draw this line [the parallel to the base] and so I will answer the question as YES!!

Bibliography

DECORTE E., GREER B. & VERSCHAFFEL L. (1996) Mathematics teaching and learning, In D. Berliner & R. Calfee (Eds.), *Handbook of Educational Psychology*, 491-549, MacMillan, New York.

HATANO G. (1982) Cognitive consequences of practice in culture specific procedural skills, *The quarterly newspaper of the Laboratory of Comparative Human Cognition*, **4**, 15-18.

HINCKLE W. (1990) *If you have a lemon, make lemonade (reissue edition)*, W. W. Norton & Co, New York, NY.

LAMPERT M. (2001) *Teaching problems and the problems of teaching*, Yale University Press, New haven.

NEWELL A. (1981) The knowledge level, *AI Magazine*, Summer 1981, 1-33.

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

SCHOENFELD A.H. (1992) Learning to think mathematically: Problem solving, metacognition, and sense-making in mathematics, in D. Grouws (Ed.), *Handbook for Research on Mathematics Teaching and Learning*, 334-370, MacMillan, New York.

SCHOENFELD A.H. (1998) Toward a theory of teaching-in-context, *Issues in Education*, **4(1)**, 1-94.

SCHOENFELD A.H. (1999) (Special Issue Editor) *Examining the Complexity of Teaching*. Special issue of the *Journal of Mathematical Behavior*, **18(3)**.

SCHOENFELD A.H. (2002) A highly interactive discourse structure, in J. Brophy (Ed.), *Social Constructivist Teaching: Its Affordances and Constraints* (Volume 9 of the series *Advances in Research on Teaching*), 131-170, Elsevier, New York.

SCHOENFELD A.H. (In press) On modeling teachers' in-the-moment decision-making. *To appear in Alan H. Schoenfeld* (Ed), *A study of teaching: Multiple lenses, multiple views, a volume in the Journal for Research in Mathematics Education* monograph series.

SCHOENFELD A.H., MINSTRELL J., & VAN ZEE E. (2000) The detailed analysis of an established teacher carrying out a non-traditional lesson, *Journal of Mathematical Behavior*, **18 (3)**, 281-325.

ALAN H. SCHOENFELD

Elizabeth and Edward Conner Professor of Education

Graduate School of Education

University of California

Berkeley, CA 94720-1670, USA

alans@berkeley.edu

