The Why? What? When? and How? of Tutoring: The Development of Helping and Tutoring Skills in Children

David Wood
Centre for Research in Development, Instruction, and Training (CREDIT), University of Nottingham

This article was originally presented as a keynote address by the author at the 2002 Teacher Leader Institute in Boston, Massachusetts.

Literacy Teaching and Learning
Volume 7, Numbers 1 & 2

pages 1–30

I'm sure that, like me, most of you share a view of the learning process and the design of effective learning environments in which the learner is active, selective, and constructive. For me, learning is also an inherently social process. The learner may want to dispute, to contest, to discuss, or to coordinate and reconcile different points of view. When in difficulty, it's often wise to look for help, and in a good learning environment, that help is available.

My topic for today is tutoring and learning. I will put most of my effort in this first presentation into discussing the analytical and technical questions about the what, when, and how of tutoring, but I would like to concentrate first on the why. Why tutor? Like all why questions, it essentially forces you back to your values: it's not about evidence so much as ethics, values, and moral stance. So, I'm going to start by saying a little bit about why I am involved, interested, and fascinated by the very complex human activity we call tutoring.

THE WHY OF TUTORING

In a nutshell, the argument I'm going to try to persuade you to accept runs something as follows: tutoring is an outgrowth of helping. Helping is an innate human propensity—we're born to help. Even though helping is a natural part of human nature, it is a much neglected topic, particularly in research. We often have a lot to say about our psychological problems rather than our strengths. One reason I'm interested in helping is that I want to try to do a little bit to restore the balance. What I'm going to try to persuade you about in looking at the development of helping in children is that there are some fascinating phenomena about childhood about which we know very little.

My own interest in tutoring did not start, in fact, from the moral issue of needing to understand human helping. It really stemmed from a series of experiences when I was working at Harvard in '69 through '72. A colleague of mine, Ken Kaye, was doing a study looking at cooperation between pairs of toddlers of about 18 months to 2 years of age. He'd shown one toddler how to solve a very simple problem. The task was this: you have a transparent tube, and inside the tube is a pile of cookies. The problem is such that if the child learned to press a little lever on the bottom of the tube, a little cap lifted up (rather like one of those flip-top trash cans), and the child could pick out the top cookie. You can train 18-month- to 2-year-olds to solve that problem quite easily, but the question Ken was asking was what happens when, having got one toddler sated with cookies, you introduce another toddler who does not know how to solve the problem into the situation. Well as you might imagine, within a short period of time, what the new toddler does is to go up to the transparent tube and tries unsuccessfully to get a cookie out. He or she usually starts to get increasingly frustrated in the process. The question now is what does the first toddler do? By 18 months to 2 years, children would spontaneously help each other.

It struck me then that what was going on here was that when people (even very young ones) see somebody else doing something that they themselves can do, and see them getting frustrated by not being able to do it, then there's a perceptual invitation to get involved. Some of us feel the invitation very strongly, others not quite so strongly. What I'm asking you to believe is that there is a natural invitational quality in the sight of people failing to do things that you can do and that this is the basis of helping.

My second experience was quite a different one. It came from an investigation that Jerry Bruner and Barbara Koslowski (1972) were doing on the development of reaching in babies. Bruner's argument was that you didn't need reinforcement to explain the emergence of reaching behaviors (this was in the 1960s, don't forget, we were still fighting the reinforcement battle). Such learning, he argued, is driven by an intrinsic drive for mastery. Bruner and Koslowski argued that babies have the intention to reach before they successfully manage to reach. There are good physiological reasons why they can't reach early in life: their muscles are very underdeveloped. Human babies are born in a very immature state physiologically, so they don't actually have the anti-gravity muscle power in their shoulders and in their arms. They can't actually reach because they don't have the power to do so.

What Jerry and Barbara did was to film babies in their pre-reaching phase. The kinds of things they observed were really quite neat. For instance, if you dangle a small object in front of babies' eyes, they would make small grasping movements, like grabbing their clothes. If you put a large object in front of their eyes, they might lean forward, bending from the waist, and throw open their arms. The argument continues that in the pre-reaching stage you see all the pre-adaptive components, which show that babies already know what they have to do: a big reaching movement for a big object and small reaching movements for a small one. All the components were there. The baby's mouth was meanwhile going like mad, drooling away. It was quite obvious where the object was going to finish up when the baby eventually got it.

I was present when Jerry presented this film to an undergraduate psychology class. What struck me as I was sitting offside and to the front was the reaction of the students. In the film, you saw all this really strong intentional behavior: the infants weren't reaching but trying to reach, and there was no sense of consummation at the end (you never got to see a baby get an object!). The students were on the edge of their seats, the tension in the room palpable. Again, it struck me that here was the same kind of phenomenon taking place as I'd seen with Ken's toddlers. When you see someone trying to do something that you can do and they're struggling, there is a real desire to get in there and help. The students in Bruner's class wanted to get in there and give the object to that baby. And babies know it, right? Do babies look for help?

We know that seeking help certainly occurs well in advance of language. Suppose you look at babies in the first 4 months of life. A good way to charac-

terize what goes on is that they're experiencing a love affair with the human face—they will remain locked on a face. At about 4 months, when they usually start to reach, they become more locked into the world of objects and begin to scan and look at things, particularly bright, moving ones. Then, at somewhere about 8 to 10 months—it varies from child to child—babies start to coordinate looking at objects and looking at people, often looking from one to the other. And when they learn to reach but an object lies out of reach, they will look at a nearby person as they stretch out (asking for help?).

So my argument here is that babies are innately endowed with some knowledge of the social world, and they expect (though not, I suspect, in a conscious sense) others to be able to help them to achieve intentions that they can't yet achieve themselves. They're natural help seekers. That's the theme I'm going to take up later. I am going to be looking at help seeking and the role of help seeking and self-helping in development as sources of individual differences that contribute to learning.

One final speculation about the why of tutoring concerns the dominant view of human evolution as a theory of sociobiology and the selfish gene. The basic idea is that genes are driven, as it were, to perpetuate themselves. The explanatory unit of evolution is the gene, not the species. You can go a long way in understanding huge sways of physical evolution and mental evolution on the basis of the assumption that selfish genes drive the process.

I was delighted to hear recently one of the main architects of this theory, Richard Dawkins, admit that there is a real problem with sociobiology theory. The problem is they cannot understand deep altruism. Surface altruism, or reciprocal altruism as they tend to call it, is easy. That's where you help people who help you, particularly members of your own family. There's no problem, for example, understanding why aunts help nephews since they share familial genes; the genes drive the aunt to help the nephew because in that way there's a greater probability that some them will get perpetuated. What Dawkins (1989) says sociobiologists can't understand is a phenomenon like adoption. They cannot understand how putting an investment into an organism that is unlikely to share any of your family genes has any possible evolutionary value for the genes. Yet that's what teachers do all the time. They help children to learn so they can adapt better to the environment, yet few of those kids are going to be genetically related to the teacher. So my challenge to social biologists is not only to understand phenomena like adoption, but also to understand helping and explain teaching.

Those, then, are some of the reasons why I think we should be interested in tutoring, not just because of its obvious practical significance for educational purposes, but because it is fundamentally important for understanding ourselves. Teachers are an enigma; they should not exist on a psychobiological account, and yet I know all too well that teachers do indeed exist.

THE HOW, WHAT, AND WHEN OF TUTORING

I want to talk now about the more analytical and the technical aspects of the how, what, and when of tutoring. I'll use an example that has to do with the development of concepts of number and counting in young children. I think it's absolutely fascinating work, and I'm going to focus you on the age of 6, when children are just beginning to understand what my colleagues in math call *base value ten.* These are students who already know that 15 is more than 12 and understand that if you've only got 12 p and you want to buy something that costs 13p or 14p or 15p, then you haven't got enough money. They can count well beyond 15 or 18, or whatever the numbers are that you are dealing with, but there's a very important watershed that these children have yet to go through. They can't yet count on from 10.

The best way to describe this is through a study invented by Nunes and Bryant (1996). Imagine that you've got a wallet, and in full view of a 6-year-old child, you put 7 pennies inside the wallet. You put the wallet down, and the child knows full well and can tell you that there are 7 pennies in it. Then you give them some more pennies outside of the wallet and ask them to count up to 11 pennies. There aren't 11 pennies outside the wallet, so the child needs to remember the fact that there are 7 pennies inside the wallet to get to 11. Got the situation in mind?

Children we face with this task fall into three groups:

- One group will count the wallet as one. So they count "one" and then count the pennies outside the wallet. They remember the pennies inside when asked, but they don't seem to know that they can be used as a set of 7.
- Then there is an intermediate group that will count "1, 2, 3, 4, 5, 6, 7" (pointing to the wallet), and then count "8, 9, 10, 11" over the pennies outside.
- The third group, the count-on group, will just count "7" (for the pennies in the wallet) and then "8, 9, 10, 11."

It's around the same time that the groups of children who can count on begin to discover that what we call a 10p, or in the United States 10 cents, represents a set of 10 units and that it too can be counted on from. The mastery of such number concepts is really quite complex (Piaget came up with a good explanation of what's going on, but we won't go over that here).

Now, this is the question. We know that this understanding about counting on is emerging round about the age of 6. We have used Bryant and Nunes' shopping task to do some roaming around the known, in Reading Recovery terminology. The question we asked is what happens when you put two children together, one who does not yet have the idea about counting on and one who does, and give them a shared task. You give the children a couple of 10-

penny pieces and a bunch of pennies, and you ask them to go shopping and to buy something. If they count up all the objects (counting the 10-pence pieces as "ones"), they haven't got enough pennies to buy this thing, but if they count on from one of the 10-pence pieces they will have enough.

So what do the pairs of children do with the shopping task? We tell the children we want them to agree on whether there's enough money provided to buy objects (that they can only afford if they count on from a 10-pence piece). We ask them to take turns doing the counting. First one of the children does it, then the other child. Each time we ask, "Do you agree with the answer?" Here's an episode from one pair.

Two children sitting at a table with an adult.
Adult: This costs 15 pence. Do you have enough?
Girl: (counting) 10, 11, 12, 13, 14, 15. Do you agree?
Boy: Uhmm…
Girl: Let's do it again. 10, 11, 12, 13, 14, 15. Do you agree?
Boy: Yes.
Girl: Okay.
Adult: This time, Lee counts first, and we have something else this time. This new eraser costs 18 pence. Do you want to see if you have 18 pence?
Boy: 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14 (counting a 10 p as 1). No, we don't have enough.
Girl: This is 10 (picking up the 10p). 10, (slowly) 11, see these are pennies so we count them, 10, 11, 12, 13, (slowly) 14, 15, 16, 17, 18. Because these are pennies, okay?
Boy: I agree.
Adult: He agrees with you. Have you got 18 p here? (Girl counts under her breath and nods in agreement). Would you like to count and see if you agree with that, Lee? Count and see if you agree.
Boy: 10, 11, 12, 13, 14, 15, 16, 17, 18.
Adult: Okay, we can move on then if you agree.

We came back the next day, and we asked the boy to do a whole variety of counting-on tasks on his own. We found that he had mastered the idea. The girl is no doubt a joy to her mother! I bet she's a joy to her teacher, and I suspect she's a joy to her friends. For me, the study really gives rise to the question, "Are teachers born or are they made?"

Another thing that is significant on the video clip is the expression on that little boy's face right at the beginning when he first tries the task. I think his thoughtful look suggested that he was experiencing an absence of a strong feeling of knowing and that he knew there was something he didn't quite understand about the task. I want to come back to that later when I discuss self-correction and help seeking.

In this study, we deliberately chose children who were just on the cusp of discovery about counting-on to work with their more knowing peers. So I'm not suggesting that what we have seen would be a common happening in daily life. But the example demonstrates that, although no one asked the girl to help and no one asked her to teach, she did so naturally. She was not the only child to do so, though most were not nearly such competent helpers.

I want you to note that when she handed over the 10p to the other child, she started the counting off but then stepped back. Another thing I want you to notice is the lovely, gentle pace of the interaction. These situations, when they work like the one you have just seen, are like a ballet. They do happen in other contexts. We know, for example, that children as young as 4 years will adapt the speed and complexity of their talk when they speak with younger children, adapting their communication to fit their interlocutor, the other person. Those abilities are being brought to bear (they are abilities, not really skills because skills are acquired through training and learning) in the example we have just shared.

TUTOR CHALLENGES

What I'm urging you to accept is that the example of the 6-year-olds represents a complex bringing together and integration of a range of competencies and skills in order to tutor others. These are competencies and skills that experienced and effective teachers have developed and honed to a fine art. But there are still difficulties and challenges even for the experienced teacher—certainly there are for me when I try to tutor.

Tutor Challenges

- Knowledge of the task
- Relating knowledge to performance
- Perspective taking
- Self-inhibition: from doing, to guiding, to fading
- Communicative competence
- Timing

Knowledge of the Task

One reasonably self-evident problem is that having knowledge of the task is not enough to guarantee effective tutoring. There were children in the study I have just illustrated who knew perfectly well how to do the task but seemed incapable of helping the other child in a contingent fashion. One outcome of our research is the finding that the chances that learners will be able to go on to do a whole range of tasks that they couldn't do before they were tutored depends on the contingency of their tutoring experience.

Relating Knowledge to Performance

What these children have to do in order to provide contingent support for learning is to relate their knowledge of the task to the ongoing, dynamic performance of their peer. This is a really hard and difficult problem. It demands what Jerry Bruner years ago called the "pedagogy of subject matter knowledge." In other words, it's not enough to know about the knowledge and skills that go into competent task performance, you've also got to know how to interpret and react to the various difficulties or sequences that learners are likely to go through as they themselves develop that knowledge and master those skills.

Perspective Taking

Some children as young as 3 years are quite able to put themselves in the perspective of another child and provide help that is contingent on that perspective (Wood, Wood, Ainsworth, & O'Malley, 1995). If you can't or will not try to see a situation from the learner's point of view, you are very unlikely to provide assistance that proves helpful.

Self-Inhibition

Another challenge is self-inhibition: leaving enough space for the learner to demonstrate whether he or she is able to carry out the task. George Herbert Mead (1950) pointed out that when we ask somebody to do something, politely, and they don't do it, we tend to get annoyed, especially if we have to ask two or three times. If they still don't do it, you tend to do what you asked from yourself. Mead argued that language sets up a tension and an urge to action in the speaker. I think that such tensions come into play when, in tutoring, you suggest actions to the learner but refrain from doing what you yourself suggest.

I think, in fact, that self-inhibition is a fundamental problem in teaching, and that's one reason why what the 6-year-old girl was doing was so special. She was able to set the other child going and then seemed naturally to inhibit herself from further action in order to leave space for the other child to demonstrate whether he was able to carry on with the task. I also think the need for constant self-inhibition in tutoring is why teachers of young children are often wrung out at the end of the teaching day. It's very, very stressful to inhibit yourself and to deny yourself the right to act.

Communicative Competence

Self-inhibition is one foundation for communicative competence in tutoring. The tutor also can't say too much in the course of the interaction without the

risk of losing or boring the learner. They must also be prepared to fade their role in the interaction, ultimately remaining mute and inactive. This all needs to be done in such a way that what is formulated by way of hints or suggestions happens in a way that's likely to be understandable to the learner. The child in the example provided was a star. Hers was quite a minimal style of teaching if you think about it: everything said was precise and to the point.

Timing

Timing in tutoring is, I think, another real challenge. In Reading Recovery I am struck by how you use time. You've got 30 minutes or so and so much to do. What I'm going to try to persuade you of is that it is worth thinking about how you use that time in some detail. Consider that there are times where and when the child needs to move on and work relatively fast. How do you separate these occasions from the times when the child needs to take time for thought—time to try for self-correction being one of those? How on earth you regulate your own use of time in real-time tutoring, I don't know. Some of you obviously do it, but I think that it is extremely difficult, so I want to spend a lot of time later talking about the nature and role of timing.

ESTABLISHING AND MAINTAINING THE TUTORIAL RELATIONSHIP

This goes right back to the 1970s, to the very first paper that Jerry and I wrote on what's going on as one person helps another to solve a problem that, left alone, they can't solve on their own (Wood, Bruner, & Ross, 1976). Jerry's got a much nicer way of putting this, he calls it the "loan of consciousness." That's one of his more poetic turns of phrases. I quite like that.

So, what we're suggesting is that there are a number of fundamental activities involved in establishing and maintaining the tutorial relationship.

Learning and Tutoring: Scaffolding Functions

- Task induction
- Highlighting and salience
- Removing distractions
- Reducing degrees of freedom
- Reminding
- State maintenance
- Modeling

Task Induction

This is really, I think, what you do when you're making use of running records to design lessons: you're trying to find a task that is challenging but manageable

and into which you can induce the learner. I think there are some really deep things underpinning this, not the least of which is the role of an absence of a strong feeling of knowing on the learner's part. By a feeling of not knowing, I mean the sense that there's more to be known than what you currently know, but you don't know what that knowledge is yet. Setting manageable problems for the learner really represents an attempt to solve the problem of what is it that the learner can just begin to recognize but not yet produce, I will argue later. That's the challenge of manageable problems, and it's where we grossly go wrong in schools in my opinion; I think it is one of the major challenges facing schools.

Highlighting and Salience

That little girl in the example did this beautifully: "This is 10…10 (slowly)." I think that is one of the most important scaffolding functions, where you draw the learner's attention to something that the learner has not yet taken full cognizance of. My suspicion is, and I've got no real evidence of it, is that this is the way tutoring often works in everyday life. Children observe other people paying attention to things that they haven't yet thought to pay attention to for themselves. In other words, the tutor makes new parts of the world salient for the learner.

I think a lot of what I saw in Reading Recovery (in fact Marie Clay has written a good deal about it) concerns the nature of the strategy that the teacher uses to help the child overcome problems. The focus of the teacher's attention at this point is vital because what the teacher highlights is what the learner is likely to attend to. So, for instance, if you provide learners with a way of solving a problem that only relies on local situational cues, you may limit their learning. If, on the other hand, you draw their attention to more strategic clues which they can then use themselves on later occasions to solve the same kind of problem, you can help them to learn generic strategies rather than just support local learning that has limited power of generalization. I think this is another really crucial issue.

Removing Distractions and Reducing Degrees of Freedom

In the pyramid-building tasks that we used in our early studies, we often saw mothers resting their hand over a block, serving to hide it so that their child would not get monopolized by it. They essentially freed the child to concentrate on other potentially more timely things by reducing the scope of the task. We often do achieve the same function by verbal means—asking the child to ignore some features or to pay attention to others. By such means, the tutor can progressively simplify the task to a point where it comes more within the learner's grasp.

Reminding

You do a lot of this in Reading Recovery lessons, where, for example, you remind the child of a word, letter, sound, or written symbol that you know he or she has shown some mastery of previously and may be able to use to meet a new demand. In this way, you may help children to develop a strategy of reasoning by analogy from past experience and to appreciate that they have relevant knowledge that, if they thought to bring it to bear on the current problem, they could use as a frame to help solve it. In this way you may also enable them to achieve joint success on a problem that they can't currently bring off by themselves. Of course, an important issue is how you highlight the task to promote strategy development, and not just local learning. Reminding is also a very powerful scaffolding tactic for state maintenance.

State Maintenance

I think you wisely keep your lessons to 30 minutes in Reading Recovery. There's a lot of evidence on your side. It's the same in lectures where, after about 20 minutes, you find that the retrieval rate for what is presented starts to fall. The importance of being in the right state or level of arousal for the learner to learn is crucial, and if you go on for too long, you risk losing the child. I keep asking my neurophysiological colleagues what biological processes underpin this phenomenon. We seem to be able to play games forever, and we can watch a film for a long time; yet somehow, when we are doing intense intellectual work, which is right in the limit of our zone of proximal development, we can concentrate for about 20 minutes, and then we're shattered. Why? My colleagues don't seem able to answer me.

Modeling or Demonstration

I'm not going along with the early behaviorists who said that modeling and imitation are primary vehicles of learning and knowledge transmission. I think the processes involved are much more subtle. The child who tries to emulate what he or she sees other people doing and achieving is clearly exploiting a remarkably powerful way of learning, especially learning in everyday life. But modeling and trying to teach through demonstration and imitation can be seriously overused and abused in tutoring, and there are limits on its efficacy as a tutoring strategy.

CONTINGENT SUPPORT FOR LEARNING

Imagine we see a child or a learner who we think is in difficulty, and we decide that the time is right to intervene. We have decided *when* to tutor, and the

question now becomes *how:* What kind of support for learning are we going to use?

Contingent Support for Learning

- Level 1: General verbal intervention
- Level 2: Specific verbal intervention
- Level 3: Specific verbal intervention plus nonverbal indicators
- Level 4: Prepares for next action
- Level 5: Demonstrates action

A good tutor will treat each of these tutorial moves as a hypothesis. They're really hypotheses about how much help the learner needs to do what it is they're trying to do. They're not one-size-fits-all moves.

Level 1: General Verbal Intervention

A general verbal intervention might be something like "It could be" or "You have a go" or "It's your turn now"; it could be a note of warning: "I'm not sure about that." A general verbal intervention is really signaling the current state of activity, but it's not trying to provide a distinct goal or objective to the child; it's not reducing degrees of freedom by very much. It's very important because it's signaling the fact that you are there and you're monitoring what's going on. It may give general feedback about efficacy ("Oh, that's great" or "Not too sure"). It's providing an external evaluation in general terms of what is happening.

Level 2: Specific Verbal Intervention

The hallmark of a specific verbal intervention is that you start to specify some action or something to pay attention to be searched for. Often, this may be expressed as a question. The task of taking that question or taking the utterance and translating into the next step is then left to the learner. In the construction task[1] you saw illustrated on video, next steps suggested included, "Why don't you find the next biggest blocks?" or "I don't think that one looks right, I think it's too small to go on top there." You're giving a specific verbal specification of the next step or of some specific feature that needs to be put right.

---

[1] The task referred to is a construction toy that invites children to put together 21 blocks to create a pyramid. Blocks can be assembled by fitting pegs into holes in order to create different-sized levels of the pyramid. Children also have to pile the levels from largest to smallest in order to create the final structure.

Level 3: Specific Intervention and Nonverbal Information

This is where the tutor adds some nonverbal intervention to what he or she says to the learner. It is the kind of situation where you point to highlight what you might be referring to verbally or where you provide a frame in which you're going to concentrate the child's attention by nonverbal cues (such as looking or pointing). So here you're beginning to solve the search problem for the child. With only a verbal instruction (Level 2), you leave the child to do all the searching in the situation. With Level 3, you're providing clues that help the child to solve that search problem.

Level 4: Prepares for Next Action

Level 4 is essentially the same as a closed question in verbal interaction. "Is it A or is it B?" or "Does that sound like s-or-t?" Here you are essentially offering the learner two alternatives to choose from: either an alternative for attention or action, or a choice between acting or of not acting at this point. Now you're really exerting very strong control over the next action.

Level 5: Demonstrates Action

You might take complete control over the next action by demonstrating or modeling what it is that should be done next in order to achieve success.

So, being contingent means

- When the learner is in trouble, offer help immediately. If you've already offered help and the learner is still in trouble (for example, you thought that giving the child a specific verbal hint or asking a specific question might be sufficient, yet the child appears not to be able to understand), you immediately offer more help.
- When the learner succeeds with help, offer less help if you intervene again. Obviously, if you're already just providing very general feedback, that's an invitation to just get out of it and leave the learner to get on with it. But if, for example, you use a Level 3 instruction (where you've not only told the child but pointed to something), then the next time you meet that problem or demand, you are going to want to step back and maybe try Level 1 or Level 2.

All of this sounds so easy, obvious, and typically what happens in tutoring. But it is not easy and does not occur typically. Often, what you find in tutoring situations, for example, is that tutors will insist on repeating what they've already said, often adding more detail, thus increasingly obscuring the message from the learner's point of view.

- Keep to the point and be succinct.
- Always succeed.

These admonitions also sound so easy to follow, but they are also so difficult to put into practice. For example, suppose you give a Level 2 instruction to a child—you suggested a goal or activity. Suppose the child actually does something else entirely, but that something else is relevant to the problem. What do you do? Should you go to Level 3 on the basis that the child failed to understand Level 2, or do you change your tutorial goal to help the child achieve what you think he or she is trying to achieve? Obviously it's a privilege of occurrence issue: if the learner has to know something before mastering something else, then you wouldn't follow the learner. But if there's no privilege of occurrence (if you could be doing A but an opportunity to do B comes up, and B is equally important, and the learner seems to want to go there), do you follow the learner? I say you do, but it's actually very difficult to switch your objective in mid-interaction.

CONTINGENT TUTORING

The idea is that working collaboratively with the learner, you will have found the appropriate level at which that child can succeed, maybe only minimally, in completing the next step of the task. We now have to expand the theory and articulate it much more carefully to describe the dimensions of contingent tutoring.

Three Dimensions of Contingent Tutoring

1. Instructional contingency – **how** to support activity
2. Domain contingency – **what** to focus on next
3. Temporal contingency – **if and when** to intervene

Instructional Contingency

I've just been illustrating one of the dimensions of contingency, instructional contingency. This focuses on how the tutor adjusts the amount of help offered, not only on the basis of the child's unaided task performance, but also on the basis of how the child responded to the tutor's previous attempts to help. Thus, the tutor has to remember how much help was provided before: "Was this child successful? If so, I must remember to fade this time." This is very, very difficult to do, but that's only one set of demands that is implicitly fulfilled when tutoring contingently.

Domain Contingency

Domain contingency concerns the issue of what to focus on next in the time

course of teaching. Now, this is a multilayered requirement. It can mean what will be focused on the current task or in the next lesson; it could involve the choice of a book. How do you decide what to choose to work on with learners next? It's a case of the amount of uncertainty you think they can handle at one time. If learners have got too many uncertainties, too many ways of being tripped up, or of making an error, they may get as much feedback as you like, but it's no use if they are so overwhelmed that they don't know what the feedback is supposed to feed back to! Learners must be provided with a really solid base of success on the next learning task so that, if and when they do get into difficulty, they have a good chance of understanding what feedback is meant to refer to in their activity. It's a way of actually managing the intellectual challenge of the task. Maintaining a high success rate for all learners is a key demand constraining the definition of domain contingency.

Even when you've decided what task you're going to induce the learner into, you've got a secondary problem, which is that almost always, there is more than one way to solve any problem. Heather Wood and I have worked with the construction task I illustrated with over 600 3- to 5-year-olds. Yet, when we see new children working with the task, we still get surprises. Children will do things—some build from the corners outwards, for example, coming up with novel solutions to problems that I certainly never envisaged. And I designed the task! That means that you cannot have too fixed an agenda as a tutor. You've got to know where you're going, but it's really a process that we call *leading by following*. You've got to know where you're going but always maintain an element of flexibility, and you've got to be open to surprise—always. Even if you can't respond to surprise in real time, you may respond to it later and, perhaps, make use of what you learn in future tutoring. How many times have you, like me, experienced times when you had tried to explain something to a student, and you know you have made a mess of it? Later you think, "Oh, of course, the child thought such and such. Well that's fine because next time that kind of thing happens, I'll do a better job of explanation."

So there's the possibility of constant reflection and development of knowledge of the learning domain on a tutor's part—that's how the knowledge that goes into achieving domain contingency grows. That is also why learning to teach a discipline takes so long. You've got to build up a huge stock of really local, contextualized knowledge. That's why I couldn't write with authority about Reading Recovery, for instance. I simply do not have that knowledge. You've got it, I hope. There's no reason why I should be expected to have it, of course. Acquiring it requires real, practical experience over very extensive periods of time. The challenge we all face here at the Institute meeting is how can we support and accelerate this process; how can we promote the professional development of new teachers so that they can master high levels of domain knowledge more readily than their ancestors and tutors did?

Temporal Contingency (or One Reason Why We Decided to Study Computer-Based Tutoring)

We have discussed how to provide tutorial support, and now we return to the issue of if and when to intervene in the first place. I used to think this was a relatively straightforward issue, but not anymore.

When you saw the two children performing the shopping task, I drew your attention to the very gentle pace of interaction, and I think this was an important element in the success of their efforts. But how do we get to grips with the general question of how best to time events in tutorial interactions? Can you say ahead of time how long you should leave a particular learner to struggle before you provide any help, for example? No? Me neither. This has far-reaching consequences for our understanding of what it takes to become an effective tutor.

Marie Clay provides a detailed, elaborate, and well-grounded analysis of the knowledge and skills that go into the reading process. This analysis is absolutely key to understanding if, when, and how to offer tutorial help. She puts forward the image of the child as a problem solver who needs to access, use, and integrate multiple sources of information. That integration has to become smoother (certain aspects of it) and more automated over time. When you're thinking about your own domain and temporal contingency (whether it's based on knowledge you gain of a child by roaming around the known or from your running records or your actual tutoring), in order to make that vital selection of what kind of demand you're going to be putting on the learner (including if, when, and how to intervene), you have to have internalized Clay's analysis of the reading domain.

If you accept this argument, then it explains why it is impossible to give a general answer to questions such as "How much time should I leave before I decide to intervene?" But how can you study the importance of timing and temporal contingency if you think this is true? How do you get to grips with intuitions about the importance of temporal contingencies in tutoring and learning?

All of these questions came into sharp focus for me from trying to implement ideas about contingency in computer-based tutoring. Every commercially available tutoring system we found allows a fixed interval of time or a fixed number of attempts for learners before it stepped in if they didn't succeed. And then they are almost never offered contingent help. If, as a programmer, you can't decide ahead of time how long a period the machine should wait before the computer gives help, when it should hasten them on, or encourage them to slow down, you can't deliver contingent tutoring.

What Heather Wood and I decided to do in our computer-based studies was to leave decisions about if and when tutorial help is given to the learner. The computer was programmed to offer (instructionally contingent) help, but

only on a request from the learner. But then you face another set of very interesting questions. Do children vary in how long they wait before they themselves decide they need help? Does this impact on how well they learn? Are some children more aware of the fact that they need help in the first place? Does it matter in terms of learning outcomes? Do some children disadvantage themselves by not seeking help appropriately? Do some ask for too much help? Put another way, how do learners influence the construction of their own tutorial experience? Does it take two to tutor?

IT TAKES TWO TO TUTOR

The following list outlines the topics that we will focus on in an attempt to get to grips with questions about how individual differences across learners influence the process of tutoring and help to construct the contingencies we observe in learner-tutor interactions.

- How the child regulates their own and the tutor's activity
- How the child reacts to impasse or error
- Ability or readiness to look for help
- Apparent inattention to help offered
- Over-reliance on help
- Allocation of time on task; speeding up and slowing down in line with mastery of the task at hand

I thought it would be useful to start by saying a few words first about the kinds of predictions and expectations that motivated that work and how these grew out of scaffolding and contingency theory.

There are, believe it or not, people who are critical of contingency theory! I think this is very sad. But it's true. There are a number of criticisms, some of which are pretty vacuous and others of which have got more weight. It was partially in response to these criticisms that I was led on to the title for this session.

There have been about 20 or so different studies of scaffolding and contingent tutoring in a variety of contexts (see for example Elbers, 1996). These have provided pretty good confirming evidence. But one critique is that there is so much focus on the tutor (who, in my experience, is very neglected in work on learning) and that scaffolding and contingency portray the child into a passive role. As far as I'm concerned, nothing could be further from the truth. I suggest a thought experiment.

Imagine that you've got a perfectly contingent tutor who is interacting with two learners, one who, for whatever reason, has a high aptitude for learning in this domain that the tutor is working in and another child who is a struggler. Given that the tutor is perfectly contingent, you should be able to predict, just by observing the tutor, the relative difficulties facing the two children. The con-

tingent tutor, in a sense, wraps around the learner. The tutor wraps around by complementing what the learner needs to achieve success. So I believe that people are misguided in saying this casts the learner into a passive role. You cannot have contingent tutoring unless you recruit the child into collaborative interaction with you, and from then on, the contingent tutor's role is paced by the progress and the activity of the learner.

I think there is a deeper sense in which the critique has got some validity. Although we never looked in detail at individual differences among children in our early work (as we have pointed out in several of our publications), there are always some children who, even given far from optimum contingency of teaching, made extraordinary progress in learning. Some children learn with minimal help because they are blessed with high aptitude, whatever that means; other children struggle. This also relates to another criticism, which is that contingent tutoring isn't a very frequent experience in children's everyday life. I think that's a perfectly fair comment. Contingent tutoring is a description of an ideal that is almost impossible to achieve in practice because the complexities and the intellectual demands on the tutor are immense. Indeed, if I were to play God and I were designing a species that had to survive through collaboration, I don't think I would want to make that species so brittle that it could only learn under conditions which provide something like optimum tutorial support for the learning process! Contingency becomes a serious issue when we are dealing with learners who are struggling. Of course this applies to all of us some of the time. It applies to some of us most of the time. In fact, one of the studies I'm going to talk about later is based on a area of learning that most of us find difficult: polynomials in mathematics. We chose this domain because we know that the majority of otherwise bright, intelligent, easy-learning people find it a difficult concept area to learn.

So I accept that it is a fair criticism that you don't get to see a great deal of highly contingent tutoring in the world around us. It is really an ideal that becomes increasingly important as we deal with difficult learning situations.

One way to address the more general criticism about the impact of learners on the tutoring process is through the study of learner help seeking, about which I will have a good deal to say. People in various parts of the world have been looking at individual differences in children's help seeking—typically, seeking help from an adult tutor. You find a number of general findings coming out of the literature (for a brief overview, see Wood & Wood, 1999). In general, children who know less and who struggle to learn seem, on average, to be less skillful help seekers. They are less likely to seek help when their activity suggests that they are in trouble, and they also seem less able to make use of that help when it is provided. Now does that square with your intuitions? Anybody want to contest it? Is nature playing a cruel trick, then, because learners who need the most contingent tutoring seem less likely to signal to the tutor

when they need their help and are less able to benefit from help when it is given? Are learners who need most tutorial support the ones who are least able to meet the tutor halfway in trying to provide a contingent learning environment for them? Are they experiencing a double whammy? Is there anything, then, in the notion that there is sort of a secondary aspect in learning difficulties, which results from the way in which children set out to help the tutor to construct their own learning environment? These are questions that motivated me into the "It Takes Two to Tutor" title and to provide a way of trying to paint a fuller picture of the learner's role in the formation of a tutorial relationship.

In what other ways might tutees influence the way in which tutors tutor them? In the video of the two children in the shopping task, the child who did not initially know how to count on from the 10p seemed, to me, a joy to help. We saw an example of super peer tutoring, but in a context where there was a quite engaging and easy-to-engage tutee. One of the characteristics that make this kind of child relatively easy to teach is, I suggest, their temperament. There is quite a lot of evidence now that extremes of temperamental type do have some kind of biological basis (for a fuller description of this work, see Wood, 1998). There's an excellent study by Bell and Waldrop undertaken here in the United States, back in the 1970s and early 1980s, in which they identified two groups of infants who formed two extremes of a normal continuum (most of us are in the middle of this continuum). At one end were the slow-to-warm-up children. These children were very difficult to get engaged in sustained interaction. Maybe they were destined to become the very shy children of the future. The other group was of impulsive children. What I liked about the Bell and Waldrop work was the way in which they showed how an individual's position on this continuum was associated with the incidence of minor physical abnormalities. Now these are things like having one toe or a finger that is relatively longer than the others or some slight abnormality in the shape of the tongue. It is unlikely that any association with these physical features and temperament is created by other people's reactions to children because they are largely invisible. This work provides the most compelling evidence that I know of to show plausible proof of a relationship between children's temperament and biological predisposition. What nobody's done (and I'll throw open to the floor in case anybody knows of work that I've not come across) is to look at the impact of these characteristics on social interaction later in life, let alone on tutoring. My bet is that there would be some pretty strong relationships, giving us a way to explore the impact of children on the tutor.

When we're confronted by children who we find difficult to teach, is it going to be particularly difficult to help them to learn how to regulate their own learning, a form of self-teaching? In other words, do children who are hard to teach find it hard to teach themselves? I believe that to a nontrivial extent,

children do have to learn how to regulate their own learning. If so, any difficulty in learning through tutoring may have very general effects on future learning. I suspect, however, that it's not true of all children that those who are difficult to teach necessarily face problems in learning how to regulate their own learning. I don't think I was a very easy child to teach, for example (in fact, several of my ex-teachers have told me so). Yet, I think I do a reasonably good job of teaching myself. But I think there probably are some kinds of more specific associations between children who are difficult to teach and children who find it difficult to regulate their own problem solving, their own learning. In fact, one of the reasons I'm so interested in Reading Recovery is that you are explicitly trying to help children become more effective learners. Your work provides you with a natural context for finding out how far learning how to learn *can* be tutored.

So here is another set of questions that helped to motivate the design and analysis of the studies that I'm going to talk about. Again, I underline the fact that I'm going to be talking about how children learn from computer-based environments, and I recognize that it is hazardous to generalize what we might learn about help seeking, self-regulation, and learning there back into face-to-face tutoring. I haven't yet explored the relationships between how children regulate the learning environment when it's a mechanical tutor and how they regulate their learning with teachers. In fact, I think you will see, as we go through the details, how extraordinarily difficult it would be to study children in face-to-face interaction in the same depth that we can with a computer-based tutor that automatically registers and stores so much detail about the interaction. So, I may speculate about what computer-based studies imply for human tutoring, but I will welcome your views on whether you think any generalizations look sound.

LOOKING FOR HELP WHEN YOU'RE IN TROUBLE

One crucial area to look at is how learners react when facing an impasse, error, or a difficulty in their activity. This seems a reasonably self-evident focus, since if children actively seek help when they don't know how to proceed, they can help the tutor to construct a contingent learning environment for them. In the computer tutors that we've developed, we leave it to the child to decide when to seek help. So, temporal contingency (decisions about if and when to intervene) are put into the child's hands. Instructional contingency is in the tutor's gift. The child decides when he or she is going to receive help, and the tutor uses what can be seen as a simple running record of recent performance to decide how much help the child seems to need. On the basis of the child's response to help, the tutor then either offers more help or starts to fade. So we can start to get a grip on how different children seek help when they feel themselves to be at an impasse or in error.

When we first started designing this approach to temporal contingency, I was talking to John Anderson who, at Carnegie Mellon University, has done pioneering work on computer-based tutoring with his colleagues. He has urged caution about leaving help seeking to learners because he suspects they might abuse help by using it to avoid effort. In other words, they'd use the help to drive the tutor to give them answers rather than trying to work things out for themselves. Now, I've got a different and perhaps more benign view of learners. I believe that if learners need and are offered Level 3 help, they will naturally tend to work out and provide an answer. I don't think they will drive the tutor on to Level 5 and an answer (unless, perhaps, they want to confirm an answer that they have in mind). So, will children generally be help abusers? Would you expect some or most children to avoid effort by overusing help?

From my reading of the help-seeking literature, I suspected that more children would be help refusers (either refusing to seek help or not appreciating a need to seek it) and that most help refusers would be the children who needed that help most—a nice, testable hypothesis.

A second issue is whether or not some children are harder to tutor contingently because they cannot or will not understand offered help. So we ask, "Are all children equally likely to be able to make successful use of help when they do request it and it's provided?" There is always the possibility that the tutor gets the help itself wrong, of course, and this is a point I will take up later.

THE LEARNER'S USE OF TIME ON TASK

I think that computer-based tutoring may offer us a unique window from which to re-view an old topic: how learners use their time on task. As you know, there is plenty of evidence showing that one of the best predictors of learning outcomes is how long the learner spends on task. In fact, time on task bedevils empirical attempts to look at the efficacy of different tutoring regimes. Since any tutoring regimen that manages to keep the child on task longer is likely to enhance learning (provided it's quality time on task), it makes it difficult to establish any differential effects of different strategies of tutoring and to prove that specific details of tutoring strategies support learning. I will show how computer-based records of tutor and learner activity can be used to help us tackle these issues (though, as I said before, it might be hazardous to generalize any conclusions back to human tutoring).

The questions about how strategic the learner is in using time on task (e.g., working quickly when the learner has a strong feeling of knowing but taking time when unsure about what to do next) give us a window to study how the child regulates the tutoring environment in ways that relate closely to tutoring strategies in Reading Recovery.

I've been struck, on the recordings of Reading Recovery tutoring that I have been shown, by how you vary in your use of your lesson time with the

learner. Sometimes the tutor seems to be trying to speed up or accelerate the child's activity; at other times the tutor leaves much more time for the child. From what I have read and been told, I suspect your changes in pace reflect your assessment about the child's level of mastery. If you feel the child has grasped the basics, you start looking for greater fluency and toward greater *automaticity* or whatever term you want to use. But if you are working in the 10% or so of areas in which you expect the child to be meeting and mastering new things, you may provide more space and time. So, there is this whole issue of how your use of the learner's time on task reflects your assumptions about where the child's learning is at. The explicit theory about the use of time that underpins Reading Recovery lessons is one reason why I got excited about the agreement between my own view of what's important in tutoring and Reading Recovery theory.

Earlier, I used the metaphor of a ballet to describe interactions between tutors and learners. In Reading Recovery lessons, it's much more like a musical medley! How does this fit in with when we look at children and how they regulate their time on task? Do you see what question I'm trying to ask? We want to know if, for example, some children spontaneously use their time on task to speed up when things look familiar, but to slow down when they are unsure. Do they help to meet the tutor halfway in regulating the use of lesson time? Are other learners less effective in this way in helping the tutor to construct a contingent learning environment? I believe we have not been able to address this kind of question before because we've just never been able to get to this level of detail in recording children's activities. Computer-based studies give us a window on this too. Even more important, once we've got this information we can ask whether the child's use of time on task relates to learning outcomes. Well, to make it worthwhile listening to the details, I'll tell you the answer to that is yes, it does relate to learning outcomes. So it is worth going into detail.

A FEELING OF NOT KNOWING, ZONE OF PROXIMAL DEVELOPMENT, AND SEEKING HELP

Often (I'm tempted to say always, but I'm not sure) in effective tutorial interactions, we're dealing with the gap between what children are starting to recognize as appropriate to a task—what sounds right, what looks right, what feels right—and what they think they themselves can currently do and achieve. They can recognize what is appropriate before they can produce it for themselves. I think this may be a constraint on the child's zone of proximal development (and maybe this is bound up with acceleration), and in tutoring, it is about keeping the child at that point at which they can just start to recognize things they can't yet quite manage. Then children become critical players in the learning process because they can judge the results of their own efforts.

I suspect something like that's going on, and I think the data we're getting from our tutors provide, or at least suggest, that this is a plausible point of view. And if it's true, it's enormously important because it defines contingent tutoring—creating a space for learning. It is a basis for successful task induction and a shared view of the task by tutor and learner—crucial elements of scaffolding. If the instruction given by the tutor is then both temporally and instructionally contingent, we get close to an ideal process. I hope this helps to explain why the recognition production gap, the feeling of knowing, and the learner's use of time on task are so enormously important for me and, I think, for Reading Recovery practice.

What's the clever thing during your problem solving when you have an absence of a strong feeling of knowing about what it is you are doing? What strategies do you have available? One of them (not the only one) is to look for help. It might be that you look for help through reading, by surfing the Web, or searching databases; it might be asking a friend or a colleague or a tutor. You look for additional information from outside. Each of these social search processes is, I suggest, an outgrowth of help seeking. You use different strategies, but all have to do with recognizing that it might be useful to de-center and open out from your own perspective to incorporate insights, information, views, and advice from outside. I wonder why we don't encourage children to seek help more in schooling? Do we really develop help-seeking skills in children in schooling? That's, again, a general question. It's not meant to be challenging or rhetorical. But it is bound up with this whole issue of what do we advise children to do when they don't have a strong feeling of confidence or there is an absence of a strong feeling of knowing.

THE QUADRATIC TUTOR

I mentioned earlier that one of the computer-tutoring studies we have done involved learning about polynomial expressions in math, although I'm not going to go into detail here about the niceties of quadratic and cubic functions and how we sought to represent and teach these in our computer-based tutor. We chose the domain because we knew that the age range we were dealing with (mainly 11- to 13-year-olds) hadn't yet been formally introduced to these concepts in their curriculum. We also suspected that only around about 10–15% would actually develop a real understanding of it. (This projection is based on early work from the Chelsea Maths group in the United Kingdom; I don't think the situation changed much in recent years.) We wanted to look at contingent tutoring in an area that was likely to present a challenge for everyone. So we could stretch our own theory to the limit, we also wanted to deal with relatively high-flying, high-mathematically gifted learners and with children who were still struggling with the basics. We deliberately chose a wide spectrum

of performance to work with the computer-based tutor. To remind you, the tutor offers instructionally contingent instruction but only at the request from the learner. Decisions about if and when to seek help and most aspects of the timing and the pacing of the tutorial sessions are in the learner's hands. Success rates—how many problems they tackle, how long they take, and so on—are all due to individual differences from learner to learner.

If we've got the design of our tutoring roughly correct and predictions from contingency theory are borne out, it should follow, as I was saying earlier, that the contingent tutor should wrap around the learner. It should be the case that the different patterns of interaction between different learners and the tutor should reflect individual differences due to the learners. Just as if I was observing you with highfliers or strugglers and you were being contingent, I would expect only to have to observe you to be able to make some strong inferences about which kind of child you were working with.

Let me point out one crucial feature of QUADRATIC—this will be very important later. Although the tutor is designed to be instructionally contingent, it is not domain contingent at all. In other words, it works much like school teaching, achievement tests, and psychological experiments in that it confronts all children with the same sequence of problems most of the time. That, as you will see, places a very important constraint on how we should interpret the different performances of high- and low-achieving learners.

Before offering them tutoring, we tested all of the children on a good, reliable, and wide-ranging test of the precursors to the development of algebraic reasoning that was developed by Hart and her colleagues at Chelsea (see Wood & Wood, 1999). After the children had worked through the tutor, we did all the usual post-tests and looked for long-term retention of outcomes. This was a pretty classical pre-test, intervention, and post-test design.

If our expectations about contingent tutoring were to be confirmed, we should find that individual differences in the offline pre-test scores correlate highly with the online interaction measures with the tutor. They were. Taken together, the online measures produced a multiple correlation with test scores of over 0.8. This is almost as good as you would want from test-retest reliability. Indeed, I would argue that we're measuring the same individual differences as the test, but by different routes. It will not surprise you to hear that the learners who (on the offline test) knew more, achieved more, or had a higher aptitude for math[2] worked at a faster rate, made fewer incorrect moves, and sought help less often (but note that they also needed less help since they experienced fewer difficulties). They were also more likely to self-correct. Bear in mind that there

[2] I'm being completely neutral here with respect to those terms. We don't know why some learners were scoring high or low on the test. It could be opportunity, differences in motivation, native wit, or aptitude.

was no attempt to force help onto the learner. The tutor did say "yep" or "nope" to correct and incorrect actions, but learners were left with 100% opportunity to self-correct and as much time as they wished to do so. So the more children knew, the more likely they were to self-correct. They were also more likely to seek help. We found this out by looking at how many times each learner sought help after an error. Higher scorers made fewer errors and they asked for help less frequently overall, but the likelihood that they would seek help after an error was higher. Put another way, lower scorers were more likely to produce whole sequences of errors, despite the fact that the system had given them feedback about errors. The high scorers, then, could self-correct more, and if they did not self-correct, the chances were that they would ask for help, leading to much fewer episodes of error on error.

This seems to add up to really bad news for the children who are struggling to learn. I think that John Anderson is absolutely right in his conclusion that errors often impede learning. Our findings with DATA agree with others that have found a negative correlation between error rates and learning outcomes. Errors not only take time but they really do seem to detract from learning gains. Now, I'm not a Skinnerian and don't want to push this too far. Fred Skinner, if you remember, had an ambition to produce conditions for error-free learning and designed his early teaching machines with this goal in mind. But I don't think errors are all of a kind, and some may well be important en route to learning. Errors are, so to speak, in our mind and may not be in the mind of the child (who may simply want to try and see what happens, for instance, rather than solve our problem). So always put errors, when I use the term, in "scare" quotes. Errors can be a normal and important product of learning. For example, when a child is exploring and searching, he or she may make errors from a tutor's perspective, but what that child is finding out may be important to discovery of the nature of constraints in the situation. However, our evidence does, I think, show that persisting in errors in the face of feedback (saying what one is trying is not working out) inhibits learning. Some of our children produced as many as 30 unsuccessful attempts in a row, presumably guessing. Yet they were facing a really complex bit of math that I really don't think they could possibly get right alone.

One implication is that these children will be more difficult to tutor because they will be giving much less clear signals as to when they themselves feel that they need help. They are certainly going to help the tutor less by not self-correcting—which is one of the jewels in the crown of learning as far as I'm concerned. It acts as a clear sign that the child is recognizing the significance of the effects of his or her own actions—a form of metacognition or knowledge of one's own cognitive activity. Self-correction is a real sign of strategic self-regulation in progress.

So with QUADRATIC it seems that the children who were struggling most with the tutorial tasks were the ones most likely to score as lower achievers on

the initial tests. Are their apparently poor strategies for regulating the learning environment one of the reasons why they are low achievers in the first place? Are we uncovering here the process that explains why they are low scorers—because they aren't very good at engaging with and regulating their learning environments? Good question; I'll return to the answer later.

IS WORKING FAST A SIGN OF EFFECTIVE LEARNING?

Children who know more, as we have seen, work faster, work more accurately; they're more autonomous, and as we know from pre- to post-test gains, they're making faster gain rates on the specific things that they've been tutored in—all good, classical stuff. The more you know to start off with and given equal time to another child who knows less, the more you're likely to learn in comparison. So does it follow then, that *better learning* is fast, error free, and autonomous? Is that what we want children's behavior to look like? Certainly, from my experience of commercially available computer-based learning environments (and games), it seems that they set out to encourage and reward fast, error-free performance. Does it follow from our results that problem solving that is fast, error free, and autonomous can be taken as a sign as of better learning?

No, it does not. It is quite possible that we find an association between the magnitude of learning outcomes and performance with the tutor because higher achievers work faster and also learn more. It does not follow that working faster is what *causes* better learning. In fact, you might simply be measuring individual differences in prior knowledge with a different set of measures—speed, self-correction, autonomy, and so on. It's quite important, this. If I don't get this across, it makes what I want to say to end this talk impossible to make sense of. Let me try to explain the issue from another angle.

Imagine that there are two children with exactly the same test scores on a prior achievement test. These two children go on to work on the tutor, and one learns more than the other. Was this the child who worked faster or slower? In fact, it is more likely to be the slower of the two. Statistically, you can work this out for the whole group of learners by factoring out the association between prior knowledge and learning outcomes and then looking to see how interaction with the tutor relates to outcomes. When you do this, you find that better learning outcomes are generally associated with a slower rate of working with the tutor. So, if anything, we might want to advise any child who is struggling to solve problems to slow down, not to speed up.

I don't think we've ever been able to address the question about how prior knowledge, learning under tutoring, and learning outcomes all relate to each other in such depth and detail before. It's one of the advantages of working with computer-based systems. Whether or not the same findings would hold true in face-to-face tutoring, I can't say. I will have to leave you to make up

your own minds without any useful evidence. However, I suspect that, in the present company, we are going to agree that the relationship between the kind of learning outcomes we are looking at and the moment-to-moment microstructure and the timing of learning under tutoring is crucial.

Learners who learned more (after individual differences due to test scores are partialled out) not only worked more slowly with the tutor, they also avoided staying locked in an impasse and avoided error by seeking help. The impact of seeking help rather than risking error on error was in fact significant for the lower achievers. So here's real evidence that seeking help was actually more important for you if you're a struggler. It is also good news for us because it suggests that our helping helped. It had a beneficial effect on learning. It also follows from this that children can't have been abusing online help. If they were, we would expect either to find no relationship with learning outcomes or a negative one (because they'd be using the tutor to avoid putting effort into learning, i.e. cheating). It is what I had suspected: as soon as the child has enough of a clue to work on, he or she will work on it to generate answers.

Finally, some thoughts about children's use of time on task. See if this is part of your intuitions: it will come as no surprise to find that higher-achieving children get their right answers to a question more quickly than lower scorers (does this imply that low scorers are, quite properly, giving themselves more time to respond?). What about wrong answers: do these tend to be faster or slower? They are faster. In fact, for over 90% of the children, error times were faster than correct responses, and higher achievers made faster errors than lower achievers. I find this really quite puzzling. Does it mean that errors come from impulsive behavior? Are the higher achievers, then, more impulsive? Or is it the case that children look at a problem situation, recognize very quickly that they don't know what it demands, and resort to search or trial and error? I can't say, though we think it is more likely to be the latter explanation. More importantly, we find that the average time to seek help with QUADRATIC took longer than either errors or successful moves. This, for me, implies that thinking before deciding to seek help is going on. Learners are generally taking time to think about the problem before asking for help. If you agree, it implies that as a tutor, you need to ask yourself if you leave enough time for children to decide to seek help before you provide some help! Also, ask if you leave enough time for attempts at self-correction after an error.

DATA: ARE LOWER ACHIEVERS GENERALLY POORER AT REGULATING THEIR OWN LEARNING ENVIRONMENT?

We have just completed a series of studies using another contingent tutor, DATA (for dynamic assessment and tutoring in arithmetic). We developed this tutor to see if we could get a grip on domain contingency. DATA is designed to

assess each learner's stage of development in the mastery of basic concepts of number, addition, and subtraction. After a detailed online assessment, it offers a tutorial program that is contingent upon each child's pattern of task mastery, offering problem-solving experiences that are tailored to each child's performance. DATA always starts out from the simplest problems that each child seems to be having difficulty with and then moves on to more advanced problems on the basis of each child's performance with the tutor.

If DATA succeeds in achieving a degree of domain contingency, then we should expect to find that the relative levels of difficulty of the tutorial problems it sets for each learner map on to individual differences in math achievement. We have tested the tutor extensively, and we are confident that this is the case. So we have some evidence that it is possible to design and implement principles of domain contingent tutoring—a tutor that wraps around learners by finding tutorial tasks that fit their current levels of knowledge and achievement.

But the main reason for presenting DATA to you today are some surprising findings about help seeking, self-correction, and the use of time on task by high and low achievers. Basically, when we look at these performance characteristics of children in DATA—where we achieve a degree of domain contingency—we find no differences associated with prior achievement. *Low scorers were just as likely as their higher-scoring peers to seek help, to self-correct, and to work at a slow or fast pace with the tutor.* Children who work more slowly with the tutor, as in QUADRATIC, tended to achieve more successful problem solving, but the rate of working on the tutorial problems was quite independent of prior achievement.

These findings, which we had not expected, are really quite important if they can be generalized. They indicate that low achievers are not, as we had thought, generally poorer at regulating the tutor or their own learning.

When you look at the literature on face-to-face interaction and help seeking, it all points in one direction: low achievers tend to be poorer than higher achievers at self-regulation—particularly in help seeking and self-explanation. This leads to the general and widely held idea that self-regulation plays an important part in *causing* such individual differences, which is what I believed. Our findings from DATA bring this idea into question. Why are our findings so different from other work in this area?

In every single study in the literature that I have looked at, nobody has honored the demands of domain contingency. And that's because we do experiments. In experiments, we want to compare like with like, so we give every participant in the experiment the same task. In fact, if you don't make sure that all your participants receive the same treatment, editors may refuse to publish you on the grounds that the data are tainted by bad experimental hygiene—as I have found to my cost. We need ways out of this situation if we are to avoid

potential pitfalls like our (invalid) generalizations about learning and metacognition. In experiments, test situations, and in much teaching, all learners receive the same problems, similar kinds of help and, perhaps, similar amounts of time to work on tasks. This means that learners who know less to start with are essentially faced with harder problems to solve than ones who know more. If, as seems likely, they turn out to show evidence of poorer help seeking or less acute knowledge of their own knowledge and skill, how should we interpret their problems? Are such differences due to something that's *inside* learners, in terms of weaker strategies of self-regulation or whatever? Or is it just down to the fact that they are more confused and uncertain by problems that are more difficult, experience less frequent success, and get help that is not contingent on their real problems? Each of these states of affairs could arise if any learning situation lacks domain contingency. Perhaps, then, DATA portrays a different picture of children learning because it provides the only situation in which each learner is compared with others when each is working on problems that are contingent upon individual levels of task mastery.

In Reading Recovery, of course, you strive to achieve domain contingency all the time. You adapt and adjust the demands of reading and writing tasks to fit your records and knowledge of each child. You don't work children through a single program or expose all to the same learning sequence. You adjust and adapt the demands you place on the learner in an effort to achieve a level of challenge that is appropriate to each: you are contingent, or you strive to be. So you should be well placed to think about, digest, and, if you feel the need, to evaluate and challenge what I have been saying. Do you find, for instance, that when you adjust the levels of demand in reading and writing situations to the needs of each learner, you find no obvious connection between effective regulation of learning (help seeking, use of time on task, self-correction) and children's levels of prior achievement? Of course, individual children will vary one from another in these aspects of self-regulation, but is it the case, as with DATA, that such individual differences are quite independent of prior levels of achievement? If so, the educational and psychological implications would be really quite important. Are you willing to help us find out if, like me, you were attributing to poor learning skills in children what might really be due to our lack of knowledge about the what, when, and the how of tutoring?

ACKNOWLEDGEMENTS

I am deeply indebted to Emily Rodgers who volunteered for the taxing and tedious task of transcribing the talk on which this account is based. Without her efforts, this paper would never have appeared. However, I don't expect Emily to bear any of the responsibility should a reader find the final product less than satisfactory; any weaknesses are entirely my responsibility. I also thank

the Atlanta team who helped to design the program, to make the Institute work, and to educate me in Reading Recovery tutoring. I also acknowledge with gratitude the collaboration of my wife and colleague, Heather Wood, who undertook most of the difficult work on which the theoretical and empirical parts of this paper relating to contingent tutoring are based. Finally, I acknowledge the extensive financial support of the Economic and Social Research Council who funded CREDIT during the period that much of the research reported was undertaken.

REFERENCES

Dawkins, R. (1989). *The selfish gene.* New York: Oxford University Press.

Elbers, E. (Ed.). (1996). Cooperation and social context in adult-child interaction. *Learning and Instruction, 6*(4), 281–286.

Koslowski, B., & Bruner, J. S. (1972). Learning to use a lever. *Child Development, 43*(3), 790–799.

Mead, G. H. (1950). *Mind, self and society.* Chicago, IL: Chicago University Press.

Nunes, T., & Bryant, P. (1996). *Children doing mathematics.* Oxford, UK: Blackwell.

Wood, D. (1998). *How children think and learn* (2nd ed.). Oxford, UK: Blackwell.

Wood, D., Bruner, J., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry, 17,* 89–100.

Wood, H., & Wood, D. (1999). Help seeking, learning and contingent tutoring. In Wood, D. (Ed.), *Computers and Education, 33*(2/3), 153–170.

Wood, D., Wood, H., Ainsworth, S., & O'Malley, C. (1995). On becoming a tutor: Towards an ontogenetic model. *Cognition and Instruction, 13*(4), 565–581.