

---

THE PSYCHOLOGY OF EDUCATION AND INSTRUCTION

---

A series of volumes edited by:  
Robert Glaser and Lauren Resnick

---

- CARROLL • *Perspectives on School Learning: Selected Writings of John B. Carroll*  
BUSSIS, CHITTENDON, AMAREL, KLAUSNER • *Inquiry into Meaning: An Investigation of Learning to Read*  
CHIPMAN, BRUSH, WILSON • *Women and Mathematics: Balancing the Equation*  
DUSEK • *Teacher Expectancies*  
FOORMAN and SIEGEL • *Acquisition of Reading Skills*  
PARIS, OLSON, STEVENSON • *Learning and Motivation in the Classroom*  
Thinking and Learning Skills  
SEGAL, CHIPMAN, GLASER • *Relating Instruction to Research, Vol. 1*  
CHIPMAN, SEGAL, GLASER • *Research and Open Questions, Vol. 2*  
BEREITER and SCARDAMALIA • *The Psychology of Written Composition*  
WEINERT and KLUWE • *Metacognition, Motivation, and Understanding*

# The Psychology of Written Composition

CARL BEREITER

MARLENE SCARDAMALIA


*Ontario Institute for Studies in Education*



Lawrence Erlbaum Associates, Publishers  
1987 Hillsdale, New Jersey London

# Chapter I

## Two Models of Composing Processes



Chomsky (1980) has argued that if we want to advance our understanding of the mind as a biological entity, we should study what it does easily and well. In particular, we should study those abilities that people acquire most naturally, with the least dependence on the environment. If that is to be the role of cognitive psychology, what becomes of instructional psychology? For if instructional psychology has a province, it is the province of things that have proved difficult to learn, that are believed to require substantial and purposely arranged contributions from the environment.

We would suggest that there are, in fact, complementary roles for a psychology of the easy and a psychology of the difficult, and that a complete cognitive science needs to encompass both. Research on cognition in infants and young children makes it increasingly evident that human beings come into the world already primed for certain kinds of knowledge. That is, they either already possess the knowledge in embryonic form or possess some kind of innate readiness to acquire it. Such appears to be the case not only with grammar, which is Chomsky's prime example, but also number (Gelman & Gallistel, 1978) and perceptual organization (Spelke, 1982). On the other hand, human beings are distinguished by their ability to acquire expertise—that is, to develop high levels of ability and knowledge of kinds that do not arise naturally out of everyday living but that require sustained effort over long periods of time. Hayes (1981) has concluded that in all those areas where top levels of expertise are equated with genius, it takes 10 years of sustained effort to achieve the necessary knowledge and skill.

In the preceding paragraph we conflated easy with natural and difficult with requiring special provisions for learning. Obviously there is no strict correlation between the naturalness and the difficulty dimensions. Walking comes naturally, yet infants work hard at it. One must be taught how to drive a car, but it is not difficult for most people to learn. The key distinction, for which we have found no convenient labels, is between those kinds of abilities that are almost inevitably acquired through ordinary living (including ordinary living in school classrooms) and those that require some special effort to transcend naturally occurring limitations. With due recognition that the terms do not fully condense the intended meanings, we shall refer to the two kinds of abilities as *natural* and *problematic*.

A complete cognitive science needs to account for both ends of the natural-to-problematic continuum. But more than that, it needs to consider possible interactions, such as the following:

To what extent are the more problematic kinds of human capabilities built up from more naturally acquired capabilities?

To what extent may naturally acquired abilities stand in the way of development of more expert ways of performing the same functions?

In this book we want to consider written composition from the standpoint of naturally acquired and more problematic human capabilities, with a view toward issues like the two just raised. By looking at both the easy and the hard aspects and their interaction, we hope to contribute to understanding both the mind's natural capabilities and what is involved in going beyond those natural capabilities.

## WRITING AS BOTH NATURAL AND PROBLEMATIC

Writing—by which we mean the composing of texts intended to be read by people not present—is a promising domain within which to study the relationship between easy and difficult cognitive functions. On the one hand, writing is a skill traditionally viewed as difficult to acquire, and one that is developed to immensely higher levels in some people than in others. Thus it is a suitable domain for the study of expertise. On the other hand, it is based on linguistic capabilities that are shared by all normal members of the species. People with only the rudiments of literacy can, if sufficiently motivated, redirect their oral language abilities into producing a written text. Indeed, children lacking even the most rudimentary alphabetism can

nevertheless produce written characters that have some linguistic efficacy (Vygotky, 1978).

There is, indeed, an interesting bifurcation in the literature between treatments of writing as a difficult task, mastered only with great effort, and treatments of it as a natural consequence of language development, needing only a healthy environment in which to flourish. Convincing facts are provided to support both views. On the one hand we have evidences of poor writing abilities, even among relatively favored university students (Lyons, 1976) and professional people (Odell, 1980). On the other hand we have numerous reports of children taking readily to literary creation when they have yet scarcely learned to handle a pencil (Graves, 1983). While children's writing is unquestionably recognizable as coming from children, it often shows the kind of expressiveness and flair that we associate with literary talent.

One could perhaps dismiss such contradictory findings as due to the application of different standards of quality. It may, in short, be easy to write poorly and difficult to write well. But that is a half truth which obscures virtually everything that is interesting about writing competence.

The view of writing that emerges from our research is more complex than either the "it's hard" or the "it's easy" view or any compromise that might be struck between them. We propose that there are two basically different models of composing that people may follow. It is possible to write well or poorly following either model. One model makes writing a fairly natural task. The task has its difficulties, but the model handles these in ways that make maximum use of already existing cognitive structures and that minimize the extent of novel problems that must be solved. The other model makes writing a task that keeps growing in complexity to match the expanding competence of the writer. Thus, as skill increases, old difficulties tend to be replaced by new ones of a higher order. Why would anyone choose the more complex model? Well, in the first place it seems that not very many people do, and it is probably never used to the exclusion of the simpler model. But for those who do use it, the more difficult model provides both the promise of higher levels of literary quality and, which is perhaps more important for most people, the opportunity to gain vastly greater cognitive benefits from the process of writing itself.

One way of writing appears to be explainable within a "psychology of the natural." It makes maximum use of natural human endowments of language competence and of skills learned through ordinary social experience, but it is also limited by them. This way of writing we shall call *knowledge telling*. The other way of writing seems to require a "psychology of the problematic" for its explanation. It involves going beyond normal linguistic endowments in order to enable the individual to accomplish alone what is normally accomplished only through social interaction—namely, the reprocessing of

knowledge. Accordingly, we shall call this model of writing *knowledge transforming*.

A two-model description may fit many other domains in addition to writing. Everyday thinking, which is easy and natural, seems to follow a different model from formal reasoning, which is more problematic (Bartlett, 1958). Similar contrasts may be drawn between casual reading and critical reading, between talking and oratory, between the singing people do when they light-heartedly burst into song and the intensely concentrated effort of the vocal artist.

In each case the contrast is between a naturally acquired ability, common to almost everyone, and a more studied ability involving skills that not everyone acquires. The more studied ability is not a matter of doing the same thing but doing it better. There are good talkers and bad orators, and most of us would prefer listening to the former. And there are surely people whose formal reasoning is a less reliable guide to wise action than some other people's everyday thought. What distinguishes the more studied abilities is that they involve deliberate, strategic control over parts of the process that are unattended to in the more naturally developed ability. That is why different models are required to describe these processes.

Such deliberate control of normally unmonitored activity exacts a price in mental effort and it opens up possibilities of error, but it also opens up possibilities of expertise that go far beyond what people are able to do with their naturally acquired abilities. In the case of writing, this means going beyond the ordinary ability to put one's thoughts and knowledge into writing. It means, among other things, being able to shape a piece of writing to achieve intended effects and to reorganize one's knowledge in the process. The main focus of this book is on the development of these and other higher-level controls over the process of composition.

### FROM CONVERSATION TO KNOWLEDGE TELLING TO KNOWLEDGE TRANSFORMING

Although children are often already proficient users of oral language at the time they begin schooling, it is usually some years before they can produce language in writing with anything like the proficiency they have in speech. Longitudinal studies by Loban (1976) suggest that the catch-up point typically comes around the age of twelve. The most immediate obstacle, of course, is the written code itself. But that is far from being the only obstacle.

Others of a less obvious nature are discussed in Chapter 3. These less obvious problems have to do with generating the content of discourse rather than with generating written language. Generating content is seldom a problem in oral discourse because of the numerous kinds of support provided by conversational partners. Without this conversational support, children encounter problems in thinking of what to say, in staying on topic, in producing an intelligible whole, and in making choices appropriate to an audience not immediately present.

In order to solve the problems of generating content without inputs from conversational partners, beginning writers must discover alternative sources of cues for retrieving content from memory. Once discourse has started, text already produced can provide cues for retrieval of related content. But they are not enough to ensure coherent discourse, except perhaps of the stream-of-consciousness variety. Two other sources of cues are the topic, often conveyed by an assignment, and the discourse schema. The latter consists of knowledge of a selected literary form (such as narrative or argument), which specifies the kinds of elements to be included in the discourse and something about their arrangement. Cues from these two additional sources should tend to elicit content that sticks to a topic and that meets the requirements of a discourse type. In essence, the knowledge-telling model is a model of how discourse production can go on, using only these sources of cues for content retrieval—topic, discourse schema, and text already produced.

The main features of the knowledge-telling model are diagrammed in Figure 1.1. The diagram indicates a composing process that begins with a writing assignment. It could also begin with a self-chosen writing project, however, so long as there is some mental representation of the task that can be analyzed into identifiers of topic and genre or discourse type. The task might, for instance, be to write an essay on whether boys and girls should play on the same sports teams. Depending on the sophistication of the writer, the topic identifiers extracted from this assignment might be *boys, girls, and sports* or *amateur sports* and *sexual equality*. According to the model, these topic identifiers serve as cues that automatically prime associated concepts through a process of spreading activation (Anderson, 1983). This process does not ensure that the information retrieved will be relevant, but there is a built-in tendency toward topical relevance. As Anderson explains, "spreading activation identifies and favors the processing of information most related to the immediate context (or sources of activation)" (Anderson, 1983, p. 86). Naturally, the appropriateness of the information retrieved will depend on the cues extracted and on the availability of information in memory. For instance, one would expect that the cues, *amateur sports* and *sexual equality*, would have a greater likelihood of eliciting information fitting the intent of the assignment than would the cues, *boys, girls, and sports*, provided the writer had knowledge stored in memory related to those

cues. In either case, however, the retrieval is assumed to take place automatically through the spread of activation, without the writer's having to monitor or plan for coherence.

Cues related to discourse type are assumed to function in much the same way. The assignment to write an essay on whether boys and girls should play on the same sports teams is likely to suggest that what is called for is an argument or *opinion essay*. Again, the cues actually extracted will depend on the sophistication of the writer. Some immature writers may have an opinion-essay schema that contains only two elements—*statement of belief* and *reason*

(see Chapter 4). Others may have more complex schemas that provide for multiple reasons, anticipation of counterarguments, and so on. In any case, it is assumed that these discourse elements function as cues for retrieval of content from memory, operating in combination with topical cues to increase the likelihood that what is retrieved will not only be relevant to the topic but also appropriate to the structure of the composition. Thus, the cues, *boys, girls, sports*, and *statement of belief* would be very likely to produce retrieval of the idea that boys and girls should or should not play on the same sports teams, an appropriate idea on which to base the opening sentence of the essay.

According to the model shown in Figure 1.1, an item of content, once retrieved, is subjected to tests of appropriateness. These could be minimal tests of whether the item "sounds right" in relation to the assignment and to text already produced or they could be more involved tests of interest, persuasive power, appropriateness to the literary genre, and so on. If the item passes the tests it is entered into notes or text and a next cycle of content generation begins. Suppose, for instance, that the first sentence produced in our example is "I think boys and girls should be allowed to play on the same sports teams, but not for hockey or football." The next cycle of content generation might make use of the same topical cues as before, plus the new cues, *hockey* and *football*, and the discourse schema cue might be changed to *reason*. A likely result, therefore, would be retrieval of a reason why boys and girls should not play hockey or football together. Content generation and writing would proceed in this way until the composition was completed.

This way of generating text content was described for us by a 12-year-old student as follows:

I have a whole bunch of ideas and write down until my supply of ideas is exhausted. Then I might try to think of more ideas up to the point when you can't get any more ideas that are worth putting down on paper and then I would end it.

Knowledge telling provides a natural and efficient solution to the problems immature writers face in generating text content without external support. The solution is efficient enough that, given any reasonable specification of topic and genre, the writer can get started in a matter of seconds and speedily produce an essay that will be on topic and that will conform to the type of text called for. The solution is natural because it makes use of readily available knowledge—thus it is favorable to report of personal experience—and it relies on already existing discourse-production skills in making use of external cues and cues generated from language production itself. It pre-serves the straight-ahead form of oral language production and requires no significantly greater amount of planning or goal-setting than does ordinary

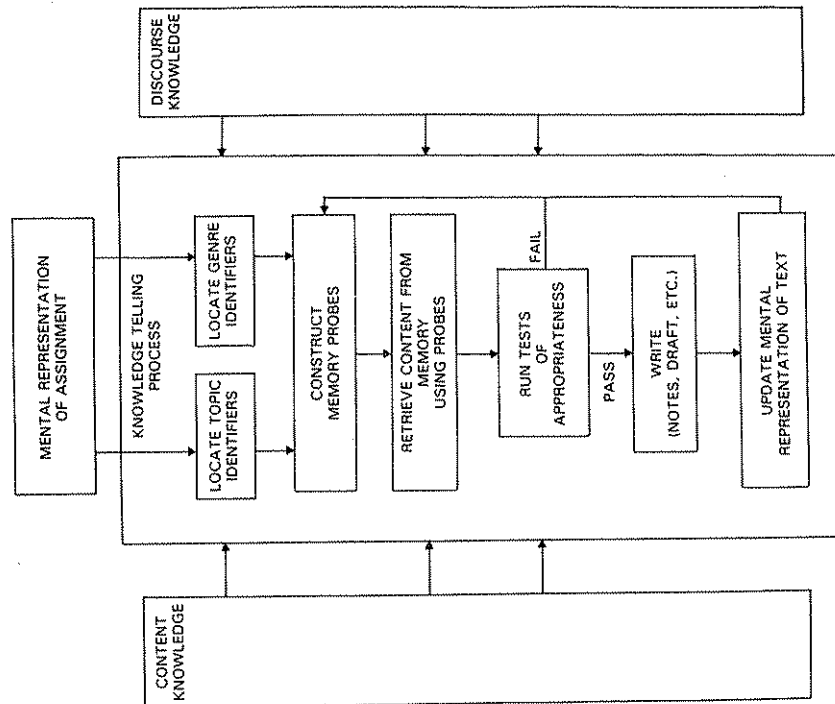


Figure 1.1. Structure of the knowledge-telling model.

conversation. Hence it should be little wonder if such an approach to writing were to be common among elementary school students and to be retained on into university and career.

### KNOWLEDGE TELLING VERSUS KNOWLEDGE TRANSFORMING

In the preceding discussion of the knowledge-telling model, it was allowed that there could be large differences in outcome depending on the writer's knowledge of the topic of discourse and on the writer's sophistication in the literary genre. In addition, of course, quality of the written product will vary depending on language abilities, such as diction and syntactic fluency, that are not dealt with in the knowledge-telling model. With all this allowance for individual differences and for improvement through learning, it is not obvious that a second model is required to account for the different ways writers go about generating text content.

Consider, however, the following description by Aldous Huxley of his composing process:

Generally, I write everything many times over. All my thoughts are second thoughts. And I correct each page a great deal, or rewrite it several times as I go along. . . . Things come to me in dribblets, and when the dribblets come I have to work hard to make them into something coherent. (Cited in *Writers at Work*, 2nd series, 1963, p. 197.)

The process described here does not sound like merely a more sophisticated or elaborate version of the process sixth-graders describe of writing down thoughts that they already have in their minds. The process Huxley describes is one in which the thoughts come into existence through the composing process itself, beginning as inchoate entities ("dribblets") and gradually, by dint of much rethinking and restating, taking the form of fully developed thoughts. This is the process that we shall call "knowledge-transforming." It is a process that cannot be accounted for by the knowledge-telling model and that seems to require a differently structured model.

This reworking or transforming of knowledge has been described in a variety of ways by professional writers (Lowenthal, 1980; Murray, 1978; Odell, 1980). But is it, then, a process found only in exceptionally talented people who have made writing their life's work? No. As studies to be reported in later chapters will show (see especially Chapters 8 and 14),

evidence of a knowledge-transforming approach to writing can be found even among people who have no particular talent for or commitment to writing, some of whom would even be judged to be bad writers by literary standards.

Where are writers who use knowledge-transforming strategies to be found? We find them among talented young students, undergraduate and graduate students in psychology, education, and English, but they could probably be found among people at advanced levels in any intellectual discipline. These are people who, like Huxley, actively rework their thoughts. While they may not have Huxley's skill in expressing those thoughts, they are used to considering whether the text they have written says what they want it to say and whether they themselves believe what the text says. In the process, they are likely to consider not only changes in the text but also changes in what they want to say. Thus it is that writing can play a role in the development of their knowledge.

To account for this interaction between text processing and knowledge processing, it is necessary to have a model of considerably greater complexity than the model of knowledge telling. Such a model is sketched in Figure 1.2. It will be noted that the knowledge-telling process, as depicted in Figure 1.1, is still there, but it is now embedded in a problem-solving process involving two different kinds of problem spaces. In the content space, problems of belief and knowledge are worked out. In the rhetorical space, problems of achieving goals of the composition are dealt with. Connections between the two problem spaces indicate output from one space serving as input to the other. For instance, a writer might be working in the rhetorical space on a problem of clarity and might arrive at the decision that she needs to define the concept of *responsibility* that she is building her argument around. This is a content problem, however, and so one might imagine a message going from the rhetorical problem space to the content problem space, saying "What do I really mean by *responsibility*?" Work on this problem within the content space might lead to determining that responsibility is not really the central issue after all but that the issue is, let us say, *competence to judge*. This decision, transferred to the rhetorical space, might initiate work on problems of modifying the text already written so as to accommodate the change in central issue. This work might give rise to further content problems, which might lead to further changes in the writer's beliefs, and so on until a text is finally created that successfully embodies the writer's latest thinking on the subject.

It is this kind of interaction between problem spaces that we argue, in Chapter 12, is the basis for reflective thought in writing. Writing is not always problematic, of course, and often we write things that have been so thoroughly thought out and rehearsed on other occasions that there is no need to reflect on them. Some writers, furthermore, may intentionally

## OBSERVABLE DIFFERENCES BETWEEN KNOWLEDGE TELLING AND KNOWLEDGE TRANSFORMING

The studies to be reported in this book are for the most part experimental studies that try to bring hidden aspects of the composing process to light. In this introductory chapter, however, we will confine ourselves to more directly observable manifestations of different approaches to composition.

It might seem that the way to begin is by showing pieces of writing that exemplify knowledge telling and knowledge transforming. That would be misleading, however. Knowledge telling and knowledge transforming refer to mental processes by which texts are composed, not to texts themselves. You cannot tell by reading this chapter whether we have engaged in problem-solving and knowledge-transforming operations while writing it or whether we have simply written down content that was already stored in memory in more or less the form presented here. You would have had to overhear, for instance, our deliberations about referring to knowledge telling as "easy" versus referring to it as "natural" to judge the extent to which a rhetorical problem led us to revise our thinking about a matter of substance.

When we see a typical example of what Macrorie (1976) calls "English"—a string of vacuous assertions dressed up in the student's impression of academic diction—we feel fairly confident that we are looking at a product of knowledge telling. But it is impossible to be sure. The student's assertion that change is the norm in this modern world of today might, in fact, express an insight sharpened by the struggle for a sententious tone. On the other hand, if it could be established that that assertion appeared on the page within 30 seconds after the essay was assigned, we might with greater confidence judge what kind of composing process generated it. It is such overt indicators of composing processes that we now survey briefly—start-up times, planning notes, thinking-aloud protocols, and revisions. The processes to which these overt indicators point are discussed in subsequent chapters.

### Start-up Times

According to the knowledge-telling model, the time it should take to get started writing upon receipt of a writing assignment is the time it takes to retrieve a first item of content fitting requirements of the topic and genre. This would vary, of course, depending on the writer's familiarity with the topic and genre. Where these are very familiar, start-up time could be expected to be very fast indeed. A corollary is that start-up time should not vary with other requirements. According to the knowledge-transforming

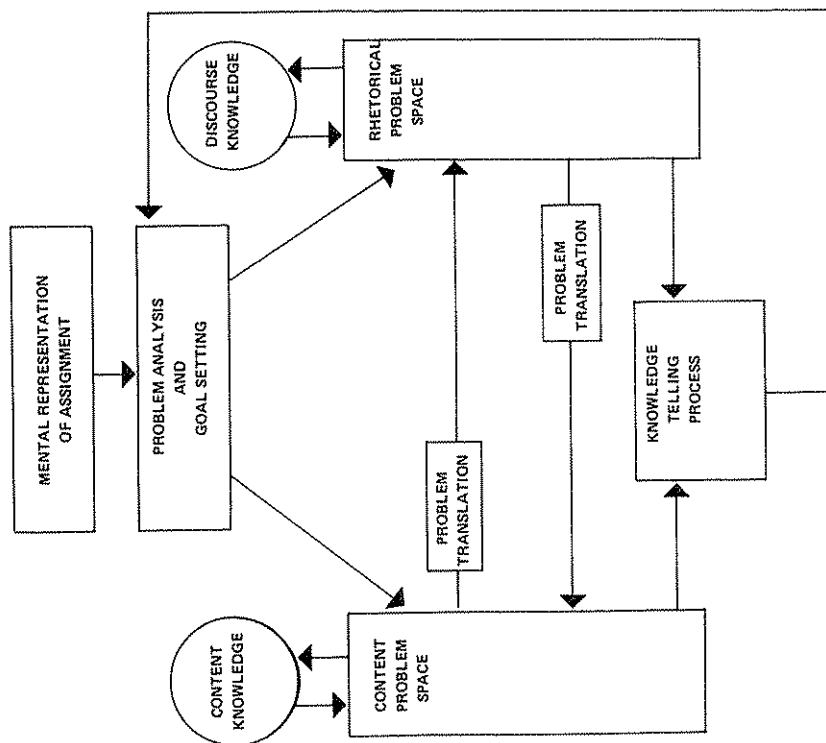


Figure 1.2. Structure of the knowledge-transforming model.

suppress problem-solving operations until a first draft is completed. In all of these cases, knowledge telling might function much as we described it in the preceding section. In this way, knowledge telling remains one of the capabilities of the knowledge-transforming model. But the distinctive capabilities of the knowledge-transforming model lie in formulating and solving problems and doing so in ways that allow a two-way interaction between continuously developing knowledge and continuously developing text.

model, on the other hand, time to start writing should, in general, depend on goals set by the writer, the kinds of problems that have to be solved in advance, and the complexity of the plan constructed. Start-up times would thus be highly variable, but they should tend to increase as more time is available and, unless there is some requirement of condensation, should be greater for an anticipated long composition than for a short one.

Figure 1.3 (from Zbrodoff, 1984) shows how long people spent between the time they were given a simple story-writing assignment and the time they started to write. The top half of Figure 1.3 shows how people responded to time constraints, which varied from being allowed only 2.5 minutes for completing the story to being allowed 20 minutes. The bottom half of the figure shows comparable data for conditions in which there were no time constraints but where length was controlled instead, with required lengths ranging from 6 to 48 lines. It may be seen that grade 5 children behaved exactly as would be expected from the knowledge-telling model. Their start-up times were very brief—just a few seconds—and they did not vary with either the amount of time allowed or with the size of text they were to produce. Adults, on the other hand, showed the kinds of adaptations to be expected from the knowledge-transforming model. The more time they were allowed, the more time they spent before beginning to write; and when no time constraints were imposed they spent much more time than under even the most liberal time limits. Their start-up times also increased depending on the length of the story they were required to produce. It may be noted that the grade-10 students were more like fifth-graders than like adults in these respects, although there is some suggestion that they took longer to get started on longer stories.

#### Notemaking

The data just discussed deal with spans of time in which text planning of some kind is presumably going on, but the data give no evidence as to what sort of mental activity is taking place. In a study that will be reported at greater length in Chapter 8, we tried to bring mental activity more into the open by instructing people to plan a composition in advance of writing it, encouraging them to take notes but urging them not to write actual text. Figure 1.4 shows notes from a graduate student planning a story on an assigned topic. What we have here is best described as a worksheet. On it the writer enters ideas at several different levels of abstraction, evaluates, and builds a structure out of them. Figure 1.5 shows, by contrast, notes typical of a 10-year-old. Except for their telegraphic form, these notes are best described as constituting the first draft of a text. As further analysis in Chapter 8 will




Figure 1.3. Time taken to start writing simple script-based stories on assigned topics (e.g., "Rick goes to a restaurant"). In Experiment 1 stories were written under different time constraints. In Experiment 2 they were written under constraints specifying the number of lines to be written (with unlimited time). From *Writing stories under time and length constraints* by N. J. Zbrodoff, 1984, unpublished doctoral dissertation, University of Toronto, Toronto. Reprinted by permission.

show, in fact, going from notes to final text is primarily a matter of editing in the case of younger students. For adults, on the other hand, going from notes to text involves going from a multi-level data structure, often set out in nonlinear form, to the creation of a linear text—a major transformation (cf. Beaugrande, 1984b).

These differences between younger and older writers are again congruent



The kid<sup>d</sup> how lost thing

There was a kind ~~man~~ named bob

He lost a bick  
on street.

He cant see it

He is sad

He got home

His mother was mad  
and what to his room bob did'4 have  
supper ~~(the)~~ in the morning he got it.  
from a big kind

the big kind ~~got~~ stole it.

His mother was ~~happye~~ happie

the big kind was punished from His  
friends.

Figure 1.5. Notes from a grade-4 student planning a story on the assigned topic, "The Kid Who Lost Things."

(superficially) but cannot alter the process. This is what would be predicted if these young writers were following the knowledge-telling model. For more expert-like writers, however, composing is a complex goal-directed activity, significant parts of which do not involve the actual generation of text content or language. Instructions to plan rather than to write bring out evidences of these other activities in the form of nonverbal symbols (arrows, etc.), comments on ideas, and other types of notes not intended to form part of an eventual text.

The kid who Lost Things: remember  
little boy he!

- kid new to the neighborhood - new to the country  
+ school - Spanish/Viet Nam

- artist - loses his brushes, erasers, papers,  
toys, hat, gloves  
- his address, balls, hockey cards,  
given things by other children → loses them

- loses his friends - others don't like  
him because they can't rely on him.

- draws a fantastic picture  
& makes a model of an airplane ship.  
model of himself - loses it - back pocket

magic  
- little model talks to him  
- helps him to find things  
one kid likes his artwork - helps him  
giving things away - kid  
helps him out in locating his losses.

→ { recognition of one's qualities }  
to learning from others & help.

Figure 1.4. Notes from a graduate student planning a story on the assigned topic, "The Kid Who Lost Things."

with the two models proposed. For novices, composing a text is a matter of generating a series of appropriate content items and writing them down. Instructions to plan rather than to write can affect the form of the output

Darkened bars show the mean number of words in thinking-aloud protocols, plain bars the mean number of words in corresponding essays. At all ages writers say more than they write, but the difference for adults is proportionately much larger (Scardamalia, 1984).




Figure 1.6. Mean number of words in thinking-aloud protocols of subjects' planning compositions, compared to mean number of words in the resulting texts.

A look at the content of the protocols indicates that the adults are not simply being more verbose, but seem to be thinking about things that the younger writers do not consider. The following is a portion of an adult protocol. Material printed in boldface duplicates or paraphrases material that actually ends up in the story being written. The remainder, and in this case by far the bulk of the protocol, consists of provisional ideas, goal

Such operations on knowledge, only indirectly related to text generation, are suggested by T. S. Eliot in a response to questioning about his composing processes:

That's one way in which my mind does seem to have worked throughout the years . . . doing things separately and then seeing the possibility of fusing them together, altering them, and making a kind of whole of them. (Cited in *Writers at Work*, 2nd series, 1963, p. 100.)

How such a process looks from the standpoint of knowledge telling is suggested by this sixth-grader's comment:

Well, if he got all those ideas . . . well, first of all he wouldn't know what order to put them in. He could put them in a really bad order. And sometimes he might forget half of the information. . . . Then he starts writing the next paragraph on something else, then suddenly remembers something more about the [first thing]. . . . [That's happened to me] once or twice—very confusing for the reader who reads it.

From his own experience, this student concluded that it is best not to get involved with separate clusters of information that must then be fused together. Stick to one source:

I think it's better if you're reading on somebody to get a book that is generally about him. So in one chapter or in one thing it would tell you everything.

#### Thinking-Aloud Protocols

Under normal circumstances, much of the planning that goes on in composition takes place during writing, not in advance of it (Hayes & Flower, 1980). Protocols obtained from people instructed to think aloud while they write provide a means of access to this kind of planning. The knowledge-telling model would suggest that what goes on mentally in the novice writer would bear a close resemblance to what appears on the page, and this is indeed what our protocol analyses of school-age writers have found (see Chapter 8). The knowledge-transforming model, on the other hand, would suggest that among more expert writers there should be a great deal of activity revealed in the thinking-aloud protocols that is not directly represented in the text. This is what our protocol analyses of graduate students show.

Figure 1.6 gives a gross quantitative indication of adult-child differences.

statements, comments, and problem-solving attempts, in relation to which the boldface statements are like tips of icebergs:

... Right now, he's isolated—and how I would... If I have a connection made there—how I want to do that: Do I want an adult to intervene? Or do I want this to be that realistic? Or fairy tale-ish? Or... because I can make it any way I want. Okay, maybe I... weird!! Ah, let me see... I know. **He makes this model of a ship, and on this ship he makes a little model of himself, and he loses it!** And this little model of himself happens to end up in his back pocket. Oh, why not? I can do anything I want with this story! Okay, so he just doesn't have any friends, and he's still losing things, and he doesn't know where he's put his ship and this little model he made of himself. But—magic!! The little model starts to talk to him, and helps him to find things! Let me see now... I want to get some other kids involved, here. There's always one kid that shines through. **Okay, one kid likes his work, his art work, and, and, helps him. so, this kid who loses things and ends up giving—instead of losing things, he gives things away!** What a mixed-up story! Imagine, giving things away!! Okay, it's not too clear right now because there's kind of... I have different directions—whether to go with the notion of this other child helping him out, or staying with the idea of some magic happening for him. Maybe I'll have both. You wanted a story? Once upon a time... Okay. Not quite, not quite! I am just writing the title again, just to get a feel for it. "The kid who lost things." I think what I'm going to do here is have development of sharing that he gives, not everything away, but gives things away to this other kid. **This other kid, in return, without knowing it, helps him out in locating his losses.**...

Here, by contrast, is the complete thinking-aloud protocol of a 10-year-old in the same study, coded in the same way. Except for a couple of procedural comments, the protocol consists entirely of content subsequently put into the story.

I could put him going to school and he probably loses a shoe. And then he's trying to find it and someone else finds it. And he goes home and tells his mother and his mother... and then the person that finds it gives it back and the next day, and then the next day, the person that found it, so the boy says thank you to the person that found it. Then the next day he goes to school, he loses something else. And the

teacher asks him what he lost and he says his short pants. He said his short pants. And they were in the washroom. And he goes home and brings them back. And then it's Saturday and school is over. And that's all. He goes back to school on Monday, he goes and plays, comes back in, he does his work, and he loses his gold ring and his mother says, "you lose everything." And he says it's true. And then he goes back to school and he never loses anything else. And his mother was happy because she doesn't want to buy him any more stuff. And they were happy from that day on. That's finished.

With a protocol like this one, it is fairly easy to reconstruct the composing process, using only the mechanisms posited by the knowledge-telling model. The first idea, "I could put him going to school," appears to be directly triggered by the narrative genre requirement of a setting. "And he probably loses a shoe" reflects the genre requirement of an "initiating event" (Stein & Glenn, 1979), combined with the assigned topic of the story (a kid who lost things). The next protocol statement suggests cuing from two more standard elements of story structure—the need for the protagonist to have a goal (finding the shoe) and the need for an obstacle to attainment of the goal (someone else's finding the shoe first). From that point on, it appears that much of the cuing of memory comes from items already generated, with continuing input from the genre schema, leading to a rambling string of ideas which nevertheless culminate in an appropriately formed story that sticks to the assigned topic. Although these processes are all detectable in the adult protocol as well, it would be impossible to account for the protocol on the basis of knowledge-telling mechanisms. A much more elaborate planning mechanism, involving goal setting and problem solving, would have to be posited—which is the sort of mechanism that we have tried to represent in the knowledge-transforming model.

The mental activities displayed by the adult and the child can both be regarded as planning, of course, and both are presumably directed toward attaining goals. But what the adult displays fits a definition of planning as working through a task at an abstract level in advance of working through it at a more concrete level (Anderson, 1983; Newell, 1980). The child's activity, on the other hand, might more precisely be labeled "rehearsal" (cf. Graves, 1975). It is working through the task at approximately the same level of concreteness as in the actual carrying out of the writing task.

## Revising

The quotation from Aldous Huxley in an earlier section, about writing things many times over and all his thoughts being second thoughts, refers to a knowledge-transforming process that often reveals itself in substantive modifications of previously written text. Revision, indeed, seems to hold an almost hallowed place among professional writers (Murray, 1978). Student writers, on the other hand, are famous for their avoidance of revision and for confining it to a cosmetic level of little more than proofreading (National Assessment of Educational Progress, 1977; Nold, 1981).

A typical elementary school manuscript might well show no revisions at all. With certain kinds of support or instruction, however, elementary school children do begin to make revisions of some consequence (see Chapter 11, this volume, and Graves, 1979). Let us therefore look at the kinds of revisions made by a student who has had the benefit of such treatment. Figure 1.7 represents the original text by a sixth-grade student, plus the revisions she made to the text after a six-week program of instruction in ways of diagnosing and remedying text problems. (The instructional study from which this protocol is taken is described in the postscript to Chapter 11.) The student has made a number of stylistic improvements, but beyond that has made a major structural improvement through the addition of a topic sentence that ties together the rest of the text. With respect to the distinction between knowledge telling and knowledge transforming, however, what we must note is that even revisions of this comparatively high level represent alternative ways of saying the same thing or additions to, rather than transformations of, information. Through the process of revision, the student in question has produced a better composition and, as a consequence, a more compelling statement on the issue of children's television-watching. But there is little to suggest that these rhetorical operations led to reconsideration of the thoughts that informed the original composition.

C. Day Lewis summed up the knowledge-transforming function of writing in the aphorism, "We do not write in order to be understood; we write in order to understand" (quoted in Murray, 1978). The response of one of our articulate sixth-grade informants to this aphorism was:

My main idea . . . is to make my ideas as clear to someone else—but only to someone else and not to me. . . . It's automatically going to be clear to me, especially if I put myself in someone else's shoes—because I wrote it.

Thus, from the perspective of the knowledge-telling model, knowledge is something one already has and that remains intact; writing is a matter of conveying a selection of this knowledge to someone else.

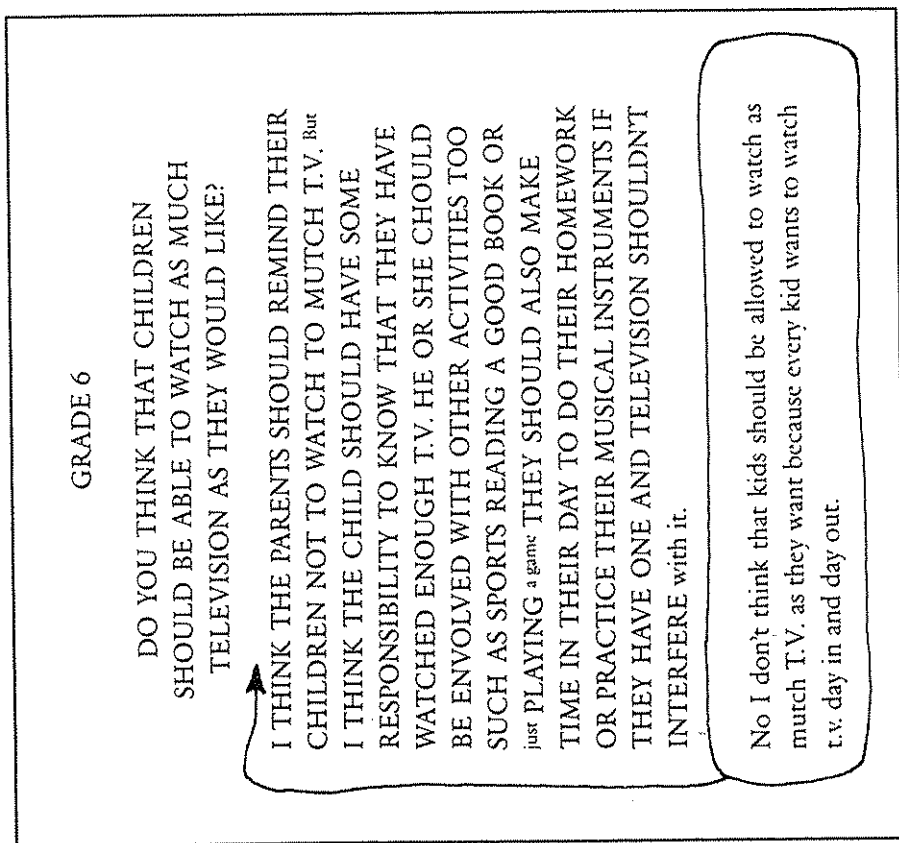


Figure 1.7. Original text and revisions by a grade-6 student who had been through a six-week program of instruction in diagnosis and revision of texts. Original text is in upper case; revisions are in lower case, placed as they actually were by the student.

## TESTING THE MODELS

The knowledge-telling and knowledge-transforming models fall somewhere in the middle among several types of models being used in cognitive research, and this raises questions about how one judges whether the models are valid, reasonable, a step in the right direction, or whatever. One type of model, often arising from work on artificial intelligence, amounts to an algorithm for executing a cognitive task (cf. Newell, 1980). Claims of merit for such models are generally based on demonstrations that the algorithm can actually be executed on a computer, producing something that bears a resemblance to observable human performance. The resemblance may be very crude and still be counted as evidence that the model is on the right track. For instance, a model of subtraction procedures intended to account for the kinds of systematic errors children make was able to account for only about half the "bugs" actually observed in children (Brown & Van Lehn, 1980). Yet the model is rightly regarded as a significant accomplishment, because it shows that a coherent set of rules can generate a substantial amount of what had previously appeared to be idiosyncratic error. The knowledge-telling model reflects a similar ambition of showing how writing behavior that has usually been attributed to deficiencies of skill or knowledge, faulty motivation, and the like, can be accounted for by a coherent procedure that runs according to reasonable rules. The model does, in fact, describe an algorithm; but the algorithm is nowhere near fully enough specified to be executed by machine, and so the possibility of validating it by showing that it can churn out something resembling student prose is not available. (Readers to whom this kind of talk is strange should be assured that there is no thought of reducing writing to something like long division, which would then be taught as a routine. Instead, the talk is about algorithms that represent partial theories of how a process works.)

Another type of model, and one of which there are already several examples in the literature of composing processes, is essentially descriptive. Such models categorize the kinds of mental processes that go on in writing and indicate something about how these processes are interrelated, but they stop short of suggesting algorithms by which composition is carried out. Examples are the models of writing proposed by Cooper and Matsubashi (1983) and Beaugrande (1984b). Such models are often useful as frameworks for research and discussion. The knowledge-telling and knowledge-transforming models sketched in this chapter resemble these other models in that they characterize processes at a rather global level and do not specify procedures in detail, but their purpose is fundamentally different. The models we have presented are not intended to describe or to summarize current knowledge about the composing process. They attempt to explain phenomena by setting forth hypothetical systems that could generate these

phenomena. Thus they come closer to what are generally recognized in science as theories. They are, for instance, falsifiable—that is, vulnerable to contrary evidence (Lakatos, 1970). We would expect any serious readers at this point in the book to have doubts and reservations about the validity of the two models, on the basis of facts which they feel to be inconsistent with what the models seem to imply. While we may hope that such doubts will be quieted as reading proceeds, it is worth noting at this time that such doubts would not occur if the models were merely descriptive.

The frequently-cited model of Hayes and Flower (1980) also occupies a middle ground between algorithmic and descriptive models, but in a different way from the models we have presented. Their model exists in two versions, representing different levels of detail. The more global version, shown in figure 1.8, is descriptive. This is the version most commonly cited, and true to its descriptive nature, it is mainly used as a framework for discussing the sorts of things that go on in composition. Figure 1.9 shows the portion of Hayes and Flower's detailed model that deals with the generating process. It may be noted that it has considerable similarity to the model of the knowledge-telling process as shown inside the central box of Figure 1.1. Even though models like this are not specific enough to be executed, they are specific enough to raise genuine questions of whether they are true or not.

The research reported in this book is not intended as verification of the knowledge-telling and knowledge-transforming models. The research ranges

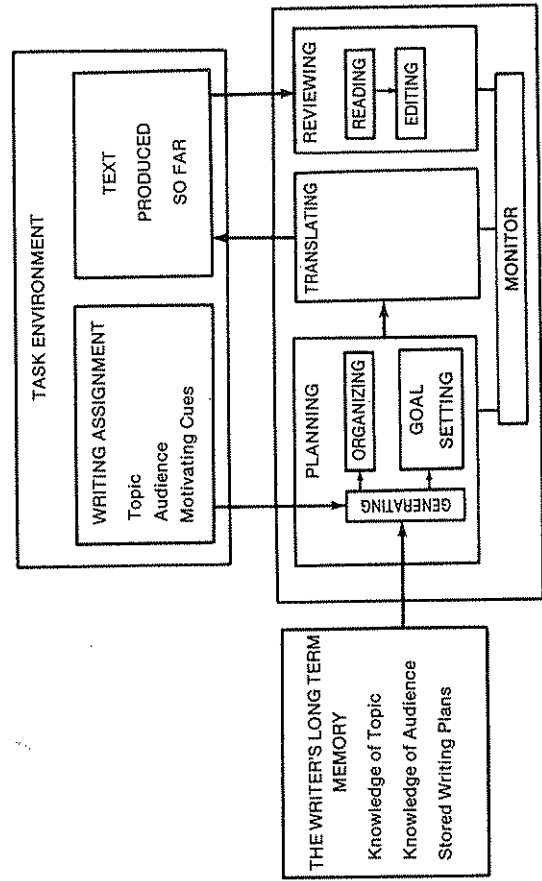


Figure 1.8. Structure of the writing model. From "Identifying the Organization of Writing Processes" by J. R. Hayes and L. S. Flower. In *Cognitive Processes in Writing* ed. by L. W. Gregg and E. R. Steinberg, 1980: Lawrence Erlbaum Associates.

## CAUTIONARY NOTES

In this chapter we have set forth two different models showing how we believe people carry out written composition. We have then tried to apply these models and give them some provisional plausibility by discussing several observable differences between the composing behaviors of children we have studied (most of whom appear to be following the knowledge-telling model of composing) and university students (at least some of whom appear to be following the knowledge-transforming model). Concrete examples always introduce the risk of unintended generalization, however, and so we devote this concluding section to trying to head off several possible misunderstandings that could becloud not only what has been said so far but also what is to follow.

I. Child versus adult competence. If one conclusion has emerged dramatically from the past decade of cognitive developmental research, it is that age differences in performance can have many sources and almost never lend themselves to straightforward interpretation in terms of intellectual competence. It is therefore important to be clear about what kinds of inferences we try to draw from the comparisons of younger and older writers that figure prominently in this book. As for what has been said so far, the matter is simple: We have not tried to say anything at all about child-versus-adult competence. The use of age comparisons has been purely incidental to the purpose of showing that there are distinct models of composing that have some correspondence to real behavior. It so happens that miscellaneous collections of elementary school students and of graduate students, when working under similar conditions, give fairly clear evidences of these contrasting models.

If we were to remain aloof from issues about the development of competence, there would be little likelihood of confusion arising from age comparisons: such comparisons would always be incidental. But the fact is that we are intensely interested in how the kind of composing process that graduate students exhibit gets established, whether what children show is an early version of the same model or whether it is qualitatively different, and what stands in the way of attaining expert-like competence. Such questions inevitably involve one in the perils of trying to identify age differences in competence among a tangle of confounding factors. One is going to make mistakes, and therefore needs working principles in order to minimize the damage done by mistakes. The following are the principles we have chosen to work by:

- a. Not to treat age (or more precisely, maturation) as an explana-




Figure 1.9. The structure of the GENERATING process. From "Identifying the Organization of Writing Processes" by J. R. Hayes and L. S. Flower. In *Cognitive Processes in Writing* ed. by L. W. Gregg and E. R. Steinberg, 1980: Lawrence Erlbaum Associates.

over many aspects of writing process and writing competence. The picture that emerges has much more to it than is covered by the models, but we believe it will become increasingly apparent as the findings unfold that something like the two models presented here is necessary in order to make sense of the picture. This will not constitute proof, of course, but it will, we hope, have the effect of raising criticism to a theoretically interesting level. That is, we hope that by the end readers who are critical of the models presented here will feel constrained to consider what other kind of model, or what revision of the existing models, would do a better job of making sense of the aggregate of findings.

tory variable. Biological maturation may, in fact, have something to do with determining the complexity of intellectual processes children can sustain at different ages (see Chapters 5 and 6). But since this question is very much open, and since the kind of research currently possible on written composition is unlikely to contribute much to resolving the question, it has seemed advisable to adopt merely as a working assumption the position that all adult-child differences are due to learning experience.

b. To try to explain differences rather than explaining them away. When older people display a higher type of performance on a task than younger people do, the difference can often be attributed to factors such as the older person's test-wisdom, interpretation of task instructions, motivation, or knowledge of relevant subject matter. It is also usually possible to show that by manipulating these factors one can reduce or even eliminate age differences (e.g., Chi, 1977). But if you show, for instance, that children produce more adult-like writing when their interest is high and when they are in a supportive environment, you have not really explained adult-child differences, you have only explained them away. For it remains that adults can write like adults even when their interest is low and when they receive no external support. A real explanation of adult-child differences must treat this differential dependence on situational factors as *part of what needs to be explained*, not as an easy way out of the explanatory muddle.

2. Knowledge telling versus dumping memory. The computer memory dump (in which the contents of the computer's memory locations are written out seriatim) provides a convenient but *completely wrong* metaphor for the knowledge-telling process. We bring the matter up only because the metaphor is apparently very attractive to people who approach writing from an information processing standpoint. Human beings cannot list the contents of their memories—least of all young children, with their more limited memory-search procedures (see Chapter 3). There is a sort of “memory dump” task in the Stanford-Binet intelligence test, calling for children simply to say all the words they can think of. Children below the age of ten typically fail to reach the criterion of retrieving 28 words in one minute. Quite to the contrary of the memory dump metaphor, the knowledge-telling model is adapted to children's difficulties with memory search; it makes maximum use of external prompts, self-generated prompts, and the apparently powerful support of discourse schemas (story schemas, etc.) for memory retrieval.

3. Good writing through knowledge telling. Much of the bad

writing that one sees is explainable by the knowledge-telling model—writing that seems to lack purpose, plan, or consideration of the reader. It should also be recognized, however, that very good writing can arise from knowledge telling. If the writer has some distinctive content already organized and available in memory—and especially if this content has affect attached to it—then the knowledge-telling process may be sufficient to generate an organized, well-elaborated, and even eloquent composition. Virtually all of the procedures that are claimed to boost the quality of student writing may be understood as ways to achieve better results through the knowledge-telling process (Scardamalia & Bereiter, 1985b).

Even though knowledge telling relies on readily available material from memory, it is probably nevertheless the case that putting such material into words has a knowledge-transforming effect. Church (1961) has speculated on the value of such activity, which he called “thematization”—the processing of relatively raw experience into verbalized form. Thus the knowledge-transforming model cannot be claimed to have an absolute advantage over the knowledge-telling model either in terms of quality of writing or in terms of cognitive benefits to the writer. Why, then, make such an issue of the distinction between models? Answers to that question should emerge gradually from the chapters that follow. For the present, we may simply recall the initial discussion about the need to study both abilities that come naturally and abilities that come with difficulty and effort. The two models of composing appear to cover that range. It is from this standpoint irrelevant whether one model is judged “better” than the other. What matters is that both are models of significant human abilities, and that investigating their acquisition promises to show something about the different tracks that cognitive development can take.

4. Models as idealizations. Do we really claim there are two models of composing? Why not three? Why not a continuum, running from relatively little strategic control over the composing process to a great deal? At a certain level of description it is no doubt true that every person who writes follows a different model (and, indeed, follows different models on different occasions). But models of the kind we have proposed are not intended to describe. They are more like representations of design concepts. They are intended to capture core ideas, which can be elaborated in different ways to correspond to real-world variability. Continuums are fine for describing the variability, but discrete models have an advantage for showing distinctions between design concepts.

As for there being two rather than three or more models, the

possibilities for other models are open. At present we are studying talented writers of elementary-school age. We are interested in seeing whether their composing resembles that of experts, whether it resembles that of less talented students their own age, or whether it suggests some different model. The several quotations sprinkled through this chapter are from students in this study. They suggest a strong and even conscious adherence to the knowledge-telling model; but on the other hand these students show a degree of consideration of audience uncommon in writers their age. At this point we do not have enough information to say whether such talent is best accounted for by an elaboration of the basic knowledge-telling design or whether it calls for a distinct model. It has seemed to us that the idea that there are two distinct models of composing rather than one was radical enough that we should not be in a hurry to add still further models.

## Chapter 2 An Integrative Schema for Studying the Composing Process

### PREFACE

There seem to be three levels at which inquiry relevant to the composing process has gone on. One level consists of personal reflections, often by writers themselves (see, for instance, *Writers at Work*, 1963; a critique of this tradition is given in Bereiter, 1984). Another, carried out mainly by educational researchers, generally follows the empiricist conventions popular in all the behavioral sciences up until the 1970s, of defining quantifiable variables and testing statistical hypotheses about them. (An influential monograph, setting forth principles and topics for this kind of research on writing was Braddock, Lloyd-Jones, & Schoer, 1963.) And then, of course, there is the long tradition of literary scholarship, usually devoted to the close analysis of individual texts. In this chapter we define three additional levels of inquiry that are important to understanding the composing process—process description, theory-testing experimentation, and simulation. These are commonplace activities in cognitive research, but just beginning to be found in writing research. Our argument is not that these three kinds of inquiry should replace the earlier three. On the contrary, the argument is that all six