CHAPTER ONE

....

The nature of reading

Introduction

I am not the first person to say that an overview of the study of the
nature of reading is impossible. The sheer volume of research on the
topic belies any individual's ability to process, much less to synthesise,
everything that is written. Similarly, the number of different theories
of reading is simply overwhelming: what it is, how it is acquired and
taught, how reading in a second language differs from reading in a first
language, how reading relates to other cognitive and perceptual abili-
ties, how it interfaces with memory. All these aspects of reading are
important, but will probably never be brought together into a coherent
and comprehensive account of what it is we do when we read. Added
to this are the inevitable complications when we consider the com-
plexities of analysing texts: since the nature of *what* we read must have
some relation to *how* we read, then text analysis must be relevant to
theories of reading and to research into reading. Yet the simple phrase
'text analysis' covers an enormous range of study within linguistics,
which again no individual can hope to overview.

 Any review, therefore, of 'the nature of reading' is bound to be
somewhat pretentious, and this introductory chapter will inevitably
be selective, rather than exhaustive. Yet consider the dilemma for
anybody wishing to assess reading. In order to assess the **construct** –
the ability we wish to test – we need to know what the construct is. In
order to devise a test or assessment procedure for reading, we must
surely appeal, if only intuitively, to some concept of what it means to

1

read texts and to understand them. How can we possibly test whether somebody has understood a text if we do not know what we mean by 'understand'? How can we possibly diagnose somebody's 'reading problems' if we have no idea what might constitute a problem, and what the possible 'causes' might be? How can we possibly decide on what 'level' a reader is 'at' if we have no idea what 'levels of reading' might exist, and what it means to be 'reading at a particular level'? In short, those who need to test reading clearly need to develop some idea of what reading is, and yet that is an enormous task.

The fact is, however, that if we wait until we have a perfect understanding of our constructs before we begin to devise assessment instruments, then we will never begin test construction. Some might say: 'Good. Better not to start than to design something invalid that may do harm.' And we might have sympathy with such a position, yet the plain fact is that assessment of reading is necessary – we will look at the multitude of real-world needs for this throughout this book. To refuse to get involved in designing instruments would thus be irresponsible, and risk the danger that others, with a lesser understanding of what is involved in reading, might design the instruments instead, with more calamitous results. Thus, testers have to get involved in test construction even though they know in advance that their understanding of the phenomenon – the construct – is faulty, partial and possibly never perfectible.

The consolation, however, is that by designing admittedly imperfect tests, we are then enabled to study the nature of the tests and the abilities that appear to be being measured by those tests. This will in turn hopefully lead to a better understanding of what one has assessed, which should feed back into theory, and further research. Thus by doing testing, provided that we research what we design, we can contribute to a growing understanding of the construct.

This is a fundamental tenet of this volume and other books in the series: it is only by trying to operationalise our theories and our understandings of the constructs through our assessment instruments that we can explore and develop our understanding. The corollary is that we need to look to theory in order to have some idea of what it is we are trying to test. This is what I shall do shortly. Before I begin, however, I should acknowledge that another approach to test design seems possible, and indeed, potentially more practical, and that is, rather than starting with theory, to begin with **target situation language use**. In other words, to begin by determining the situations in

which the persons to be assessed will need to 'read'; to analyse such situations; and then to devise assessment instruments which reflect reading in those target situations; and 'see' 'how well' our assessees can 'read'. Indeed, such approaches will be illustrated later in this book. Note, however, that even such an approach needs some crude notion of what we mean by the words in quotation marks: 'read', 'see' and 'how well'. 'How well' implies some sort of standard, at the very least some notion of comparison with how others read; 'see' implies that there are acceptable ways of externalising either how people are reading, or what they have understood of what they have read; 'read' implies that we know what it means to read, to process text meaning through some process of interaction with print.

Rather than continue in this vein indefinitely, we need to start somewhere, and I shall do so by considering the nature of reading.

Process and product

It is commonplace to make a distinction between the **process** of reading, and the result of that process, the **product**. The process is what we mean by 'reading' proper: the interaction between a reader and the text. During that process, presumably, many things are happening. Not only is the reader looking at print, deciphering in some sense the marks on the page, 'deciding' what they 'mean' and how they relate to each other. The reader is presumably also 'thinking' about what he is reading: what it means to him, how it relates to other things he has read, to things he knows, to what he expects to come next in texts like this. He is presumably thinking about how useful, entertaining, boring, crazy, the text is. He may be consciously reflecting on the difficulties or ease he is experiencing when reading, and on ways of overcoming the difficulties or of continuing the pleasure. He may be completely unconscious of how he is reading, and of what is happening around him: he may be fully absorbed in 'reading'.

Evidently, many different things can be going on when a reader reads: the process is likely to be dynamic, variable, and different for the same reader on the same text at a different time or with a different purpose in reading. It is even more likely, then, that the process will be different for different readers on different texts at different times and with different purposes. Understanding the process of reading is presumably important to an understanding of the nature of reading,

but at the same time it is evidently a difficult thing to do. The process is normally silent, internal, private.

Research has focused on examining the eye movements of readers, and interesting insights have been gained from eye movement photography. Watching what the eyes are doing, however, may not tell us what the brain is doing if, in Smith's (1971) terms, 'What the Brain Tells the Eye is More Important than What the Eye Tells the Brain.'

Asking the reader to read aloud is an alternative to eye movement photography as a means of externalising the reading process, and **miscue analysis** (which analyses the mistakes readers make when reading aloud – for details see Goodman, 1969) is one method of investigating the reading-aloud process. Yet reading aloud is not the 'normal' way in which people read, and the process of reading aloud may be very different from reading silently. Externalising the private process of reading may be the only way to inspect it, yet such externalising risks distorting and changing the nature of the process.

Introspection, through think-aloud protocols or verbal retrospection in interviews, is an increasingly frequently used method of investigating the reading process, and researchers have identified different strategies that good and poor readers appear to use when reading; they have investigated the parts of text that cause problems when reading; and they have also looked at the affective issues that arise when readers are processing particular texts. Introspective methodologies have their critics and are obviously limited in how much light they can throw on the process, but, equally obviously, such methodologies have their uses.

Other research methodologies are also possible and indeed used; it is not the purpose of this chapter to review research methodologies (see Chapter 9), but simply to indicate both the importance and possibilities of examining the reading process in order to understand it, and to understand the limitations that such research must, perhaps inevitably, have.

An alternative approach to examining the process of reading is to inspect the product of reading and, often, to compare that product with the text originally read. It is sometimes said that, although different readers may engage in very different reading processes, the understandings they end up with will be similar. Thus, although there may be many different ways of reaching a given understanding, what matters is not *how* you reach that understanding, but *the fact that* you reach it, or, to put it another way, what understanding you do reach.

The problem of potentially infinite variation in processes of inter-preting text is then supposedly reduced by a focus on what one has understood. Product approaches to reading have been unfashionable in recent years as research efforts have concentrated on under-standing the reading process, and as teachers of reading have endeav-oured to improve the way in which their students approach text. However, a great deal of research into reading earlier this century used essentially product approaches to reading, and much research into the effect of linguistic variables still concentrates on the product of reading. Both a growing realisation that processes of reading are more complex than originally assumed, and the inevitable pendulum swing in research and teaching fashions, have led to revived interest in the product of reading.

As mentioned above, earlier research into reading used a product approach. This means that researchers would typically design tests of understanding of particular texts, administer the tests to suitable informants, using particular research designs, and then inspect the relationship between the results of the tests and variables of interest.

For example, readability researchers would relate scores on reading tests to measures of the linguistic complexity of particular texts, in order to arrive at estimates of text difficulty. Researchers interested in understanding reading ability would devise text comprehension ques-tions at various 'levels of understanding' (see below) and would then see how readers fared on these different questions. Other researchers, wishing to understand what distinguished one type of reader from another (boys versus girls, first-language readers versus second-language readers, children taught by 'whole-word approaches' versus children taught by 'phonics' methods, and so on), might compare and contrast the summaries made by their subjects after reading parti-cular texts. What these studies have in common is that they take some measure of text understanding – test questions, summaries, even interviews – and relate that measure to other relevant variables.

There are at least two limitations to, or problems with, product approaches to reading: one is the variation in the product, the other is the method used to measure the product.

To take the matter of variation first. As we shall see in more detail in Chapter 2, it is clear that what readers understand from text varies. Obviously what people *remember* of what they have read will be affected by their ability to remember. Leaving aside variations in memory, however, and assuming that our measures of understanding

do not depend upon readers' memories, it is still the case that different readers will develop somewhat different understandings of what a text 'means'. This is at least in part because a text does not 'contain' meaning which is waiting to be discovered by an able reader. Rather, meaning is created in the interaction between a reader and a text: the text has what Halliday (1979) and Widdowson (1979) call **meaning potential**, and the potential is realised – in the product of understanding – only by readers reading. Since, as we shall see in Chapter 2, readers' knowledge and experiences influence the realisation of this meaning potential, and since readers may differ in their knowledge and experiences, then the products of reading will also necessarily differ.

Given such differences in understanding – the products – the issue is: how are we to determine (if at all) which product, which understanding is 'correct', and which is 'incorrect'? One approach popular among post-modernists is to say that all products are possible and equally 'correct', or that none are correct, and that the notion of correctness is inappropriate, or theoretically misguided. Without wishing to take sides in this somewhat philosophical argument, which clearly has some force – how else can we account for the fact that people do have legitimately different interpretations of text? How else can we account for the existence of lawyers as a profession? – there must also be some acceptance at a common-sense level that some interpretations of text are simply 'wrong': they do not represent any plausible interpretation of an author's possible intentions. The problem remains, for researchers, theorists and test constructors alike: how to decide which interpretations are acceptable and which are not? Test constructors in particular will need to be able to answer that question, since it is surely not adequate to say that somebody has only understood a text when he agrees with the test constructor's interpretation. Yet this is all too often what happens.

The second problem alluded to above is the method by which one has assessed the product of understanding. This issue will be addressed in more detail in Chapter 7, since it is central to concerns in the testing of reading. It is mentioned here to show the inevitable limitations in theories as well as tests.

If the method of assessing reading product – comprehension – involves a reader recalling what he has read without further recourse to the text (as happens, for example, in the use of recall protocols and interviews, or in some kinds of summary test), then it will be difficult

to distinguish understanding from remembering. If the method of testing is unfamiliar to readers (as happens in some cultures with multiple-choice tests, for example), then one risks a test-method effect. Similarly, if the method – as seems to happen in the case of cloze techniques and gap-filling – induces some readers to read in a particular way (paying close attention to individual words, for instance, or reading the text preceding the gap, but not the following text), then it will be difficult to generalise from a specific test performance to an ability to read, especially when assessed by other methods. It may be the case that some understandings can be assessed by some methods and not by others: can the cloze procedure, for instance, assess whether the reader has read a text critically, rather than passively? If not, obviously the view of understanding derived from the product assessed by such a method will be limited.

What is not always realised when building theories of reading upon the results of such research is that the theories do depend rather centrally on the validity of the measures of understanding used, and the 'accuracy' of the researcher's definition of 'adequate understanding'. This, incidentally, is a nice illustration both of the centrality of some means of assessing reading to the development of a theory (and the limitations therefore of such theories), and of the near circularity of using test results to build theories on which to base test construction. I shall return to this issue in later chapters.

To summarise thus far: it is possible to see reading as a process, or to examine the product of that process. Any theory of reading is likely to be affected by the emphasis that is placed on process or product. Product is easier to investigate than process, although this is not without its problems.

Levels of understanding

It is commonplace in theories of reading as well as in everyday talk about reading to distinguish different **levels of understanding** of a text. Thus, some may distinguish between a literal understanding of text, an understanding of meanings that are not directly stated in text, or an understanding of the main implications of text. Similarly the distinction between understanding details and understanding the main idea of a text is familiar enough to teachers of reading, as is Gray's (1960) distinction between reading 'the lines', reading

'between the lines', and reading 'beyond the lines'. The first refers to the literal meaning of text, the second to inferred meanings, and the third to readers' critical evaluations of text.

Such distinctions clearly relate to the product of reading, and enable us to describe some of the observed differences in understanding among readers. They also enable the evaluation of such differences, since it is believed that inferred meanings are somehow 'deeper' than literal meanings, and that a critical understanding of a text is more highly valued by society than a 'mere' literal understanding. Such value judgements lead to an implicit (at times explicit) hierarchy of levels of understanding: the literal level being considered somehow 'lower' than critical understanding. This in turn leads to an assumption that it is more 'difficult' to reach a critical understanding of text than it is to infer meanings, and that both of these are more difficult than 'merely' understanding the literal meaning. Thus the notion of levels of understanding becomes overladen with an ordered hierarchy of increasingly valued and increasingly difficult 'meanings'. The next logical leap is from this ordered hierarchy of difficulty and value to a hierarchy of acquisition: it is very frequently assumed that readers first learn how to understand texts literally, then to infer meanings from text, and only later do they learn how to approach text critically, to evaluate text, and so on. Thus it is often asserted that the levels are ordered: i.e. one must understand the lines in order to read between them, and one had better understand both before adventuring beyond them. In fact, the empirical justification for such assumptions is very slim indeed, as we shall see in Chapter 2, but the theoretical notions are persuasive, especially to teachers of reading, and they are thus pervasive.

However, although intuitively appealing, such distinctions among 'levels of understanding' are not always easy to define. Since language is rarely completely explicit, normal language processing requires the reader to make inferences. As Bransford *et al.* (1984) show, readers of the sentence '*The floor was dirty because Sally used the mop*' will readily – some would say automatically – infer that '*the mop was dirty*', yet this statement was not made 'literally'. Similarly, writers must make assumptions about their readers' knowledge, since total explicitness would lead to enormously unwieldy use of language, and would probably make communication impossible. If readers do not possess the knowledge that writers assume, then difficulties in literal understanding will occur, even if inferences can be made.

In summary, a consideration of the nature of reading must include recognition of frequently made distinctions among levels of meaning and understanding in and from text. Test constructors, thus, must also consider the level of meaning that they believe readers ought to 'get out of' a particular text when assessing 'how well' they have understood the text in question.

What does it mean to be able to read?

Discussions of 'levels of understanding' frequently merge into a discussion of a reader's ability to understand at certain levels. Kintsch and Yarbrough (1982), for instance, distinguish levels of comprehension: it is possible to comprehend the words but not the meaning of a sentence, and sentences but not the organisation of the text. Kintsch and van Dijk (1978) relate the former to 'microprocesses' and the latter to 'macroprocesses': microprocesses have to do with local, phrase-by-phrase understanding, macroprocesses with global understanding. In fact, as mentioned above, reading researchers have frequently attempted to identify reading skills or abilities by giving subjects a series of passages, and asking them questions intended to test different levels of understanding of the passages. Thus 'the ability to make inferences' becomes defined as 'the ability to answer a question relating to meanings not directly stated in text'. There is, of course, a degree of circularity in such definitions, but that has not stopped researchers and theorists from positing the existence of reading skills and subskills from the answers to such questions. It is common to factor-analyse the results of such answers, and then to state that questions that load on the same factor measure the same skill or subskill. In such a fashion, many different lists, taxonomies and even hierarchies of skills have been developed, as Alderson and Lukmani (1989) point out. The New York City Board of Education is cited by Lunzer and Gardner (1979) as identifying thirty-six different skills. Davis (1968) defines eight skills, as follows:

1 recalling word meanings
2 drawing inferences about the meaning of a word in context
3 finding answers to questions answered explicitly or in paraphrase
4 weaving together ideas in the content

5 drawing inferences from the content

6 recognising a writer's purpose, attitude, tone and mood

7 identifying a writer's technique

8 following the structure of a passage

As we shall see in Chapter 2, however, there is a considerable degree of controversy in the theory of reading over whether it is possible to identify and label separate skills of reading. Thus, it is unclear (a) whether separable skills exist, and (b) what such skills might consist of and how they might be classified (as well as acquired, taught and tested). Nevertheless, the notion of skills and subskills in reading is enormously pervasive and influential, despite the lack of clear empirical justification.

Bloom's 'Taxonomy of Educational Objectives in the Cognitive Domain' (Bloom *et al.* 1956) appeals to similar theorising about the components of educational achievement, and his taxonomy has been enormously influential in the devising of curricula, instructional material and tests. In second-language education, Munby's taxonomy of microskills has been influential in syllabus and materials design as well as the design of language tests. Munby (1978) distinguishes the following reading 'microskills':

- recognising the script of a language
- deducing the meaning and use of unfamiliar lexical items
- understanding explicitly stated information
- understanding information when not explicitly stated
- understanding conceptual meaning
- understanding the communicative value of sentences
- understanding relations within the sentence
- understanding relations between parts of text through lexical cohesion devices
- understanding cohesion between parts of a text through grammatical cohesion devices
- interpreting text by going outside it
- recognising indicators in discourse
- identifying the main point or important information in discourse
- distinguishing the main idea from supporting details