

CHAPTER THREE

Approaches to Learning

FERENCE MARTON AND ROGER SÄLJÖ

Explaining Differences in Outcome

How can we account for the qualitative differences in the outcome of learning described in the previous chapter? Let us take the first of our own studies as an example. Students were asked to read an article on university reform intended to bring the pass rates of universities more in line with those of polytechnic institutes (see Chapter 2). As there were substantial differences in pass rates between different groups within universities, the author argued that improvement in pass rates at universities, if necessary at all, would depend on taking selective measures, i.e. measures directed towards the groups with low pass rates and not towards those whose results were satisfying. As we have seen, four qualitative different outcomes were identified (Marton, 1974; Marton & Säljö, 1976a):

- A. Selective measures should be taken
- B. Differential measures should be taken
- C. Measures should be taken
- D. There are differences between different groups of students

Now, how did these differences in understanding come about? Those whose answer was of the C-variety, for instance, obviously thought that the author was arguing for something which, in reality, he was arguing against (i.e. general measures). This observation could be seen as a reminder of the kinds of problems one finds when analysing in detail how people read texts and how they learn. Those with a D-kind of understanding, furthermore, seem to have totally missed the point that the author was arguing for anything at all. Probably, they assumed that he simply wanted to describe something, to convey information.

The most obvious explanation of why such variations in understanding arise would be to argue that learning depends on prior knowledge. Thus the differing outcomes could be explained in terms of differing levels of knowledge or linguistic skills. Although such an argument may be true in a general way, it cannot explain the results of this experiment. The article here was chosen specifically because the language used was simple (it was an article taken from a daily newspaper), and because the prerequisite knowledge could reasonably be assumed to be available to all the students (it was about a widely debated university reform). After having, at least tentatively, ruled out that explanation, the next one again seems fairly obvious. The students understand the text in different ways because the students themselves differ; above all some are brighter than others. Apart from the fact that the text could hardly be called intellectually demanding, the proposed “explanation” could be no more than a statistical correlation at best.

This type of explanation does not illuminate the fundamental question of how the different ways of understanding the text have come about.

If the outcome of learning differs between individuals, then the very process of learning which leads to different outcomes must also have differed between individuals. This is a fundamental assumption underlying the line of reasoning pursued in these studies. The most obvious explanation of the differences in outcome should derive from a description of the differences in the process that led to the different outcomes.

After having reached such a position, we still face a highly important question concerning the strategy of research: What does it take to describe differences in the learning process? Of course, we could have had a theory or a model of processes involved in learning by reading a text. We could have attempted to test that theory or model by *ceteris paribus* varying one factor at a time between one or several experimental and control groups. In fact we did not have any such theory or model to test in an experimental setting. On the contrary, we started from the assumption that the functional background of differences in the outcome of learning in natural study situations was still largely unknown. In consequence we had to try to find out in what way students function differently in such situations. But what sources of information could we find in order to be able to answer this question? Observing students engaged in studying, is really not a very rewarding research method. There is simply not much to observe. We can measure the time spent on reading the text, we can examine the underlinings and notes made, but such data do not provide useful information. External data of this type do not form a pattern systematically related to the outcome of learning (see Svensson, 1976). And, indeed, there are no good reasons why they should.

The Original Experiment

There are, then, basically two alternatives left for collecting data about what the students actually “do” when trying to learn from a text. One is to scrutinise the various qualitative aspects of the students’ performance, of the outcome of learning, in order to be able to make inferences of the type “to recall this, or understand that, the student must have proceeded in this or that way”. This was one of the methods used by Svensson (1976, 1977) which led him to the conclusions discussed in the next chapter. Here we shall consider results derived from the alternative strategy (Marton, 1974). Students were asked to recount how they had been handling the learning task and how it appeared to them. To ask the students to describe how they had been handling the learning task is to some extent tantamount to asking them how the learning task and the learning situation appeared to them, because it is the only language in which questions about what we do when we try to learn by reading a text can be answered. Answers to such questions are not of an introspective nature; there is just no way in which we can look into ourselves. What we can do instead is to say how the world appears to us and this was exactly what the students did in our experiments.

The basic methodology was introduced in the previous chapter. Students were asked to read the article, knowing they would be asked questions on it afterwards.

Besides the questions about what they remembered of its content, students were also asked questions designed to discover how they had tackled this task. They were asked, for example:

Could you describe how you went about reading the text?

Was there anything that you found difficult?

Did you find it interesting or not?

While reading, was there anything that struck you as particularly important?

Each student participated in an individually run session and all the conversation between him or her and the experimenter was recorded and transcribed verbatim subsequently. The transcripts of the students’ answers to these and other similar questions made up the data base for our attempt to answer the main question dealt with in this chapter: how did the students arrive at those qualitatively different ways of understanding the text read?

Methods of Analysis

In Chapter 1 the idea of *phenomenography* was introduced as the method which grew out of the approach developed in this study. Here, more detail is provided to show how the forerunner of this method was used in the analysis of the interview data. Of course, the actual term *phenomenography* post-dates the original work reported here.

The first phase of the analysis, which was supposed to lead to the results searched for, was a kind of selection procedure based on criteria of relevance. Comments which seemed in any way relevant to our enquiry were identified and marked. The meaning of a comment could occasionally lie in the words themselves but, in general, the interpretation had to be made in relation to the context within which that comment had been made. Svensson and Theman (1983) offer an illuminating example of the way in which the same utterance may take on different meanings in different contexts. The phenomenon in question (“differences in the learning process accounting for the differences in outcome”) was thus *delimited* and interpreted in terms of interview extracts which were selected quotes from the interviews with the students, while the quotes themselves were delimited in terms of the context from which they were taken.

The selection procedure resulted in a collection of comments relevant from the point of view of the phenomenon to be delimited. Extracts were then brought together into groups on the basis of similarity and the groups were delimited from each other in terms of differences. In very concrete terms it meant sorting the quotes into piles, trying to extract a core meaning common to all the quotes in a certain pile, examining the borderline cases and eventually making explicit the criteria attributes defining each group, not the least in contrast to the other groups. In such a way the group of quotes were turned into categories defined in terms of core meaning, on the one hand, and borderline cases, on the other. Each category was exemplified by a selection of appropriate quotes.

An important difference between the way in which phenomenographic analyses are conducted and how traditional content analysis develops is that the categories

into which the comments are sorted are predetermined. The analysis is dialectical in the sense that bringing the quotes together develops the meaning of the category, while at the same time the evolving meaning determined which of the categories are included or omitted. This means of course a lengthy and painstaking iterative procedure with continual modifications in which quotes are assembled, and consequently further changes in the precise meaning of each group of quotes take place. There is, however, a decreasing rate of change and eventually the whole system becomes stabilised. Each category is then as homogeneous as possible. The outcome is a hierarchical structure of categories, chiefly related to each other in terms of similarities and differences.

A second important difference in relation to content analysis is the level at which we examine what the students say about their experience of learning.

The next crucial step in this particular analysis was to consider the whole set of quotes selected to delimit the various categories of description. The researcher's attention now shifted from the individual students (the interviews which had lent meanings to the quotes by being their contexts), to the meanings embedded in the quotes regardless of whether these different meanings originated from the same individuals or not. The boundaries between individuals were thus abandoned and interest was focused on the "pool of meanings". In this way, each quote had two contexts in relation to which it had to be interpreted. First it depended on the interview from which it was taken and then on the "pool of meanings" to which it belonged. The interpretation was thus an iterative procedure which went back and forth between the two contexts for each unit of analysis. The first phase of the analysis was thus a selection procedure carried out within each interview (though taking the other interviews into consideration as a background). The second phase was the shift of attention from the individual interview to the "pool of meanings" consisting of the relevant quotes selected. Then came a third phase which involved a decision about the specific level at which the quotes should be seen in relation to each other. It should be noted at this stage, however, that these different phases were not, of course, strictly sequential. There was a good deal of overlap as the iterative procedure progressed. The differences in the outcome of learning referred to in the previous chapter were described at a certain level and if we were to find the differences in the process of learning accounting for those, we would have to aim at a description at the same level, and outcome and process would have to be described in a comparable way.

There are obviously differences at different levels. There are differences in the way people express themselves and there may be differences in their general orientation, but our interest did not focus on either of these two levels. Differences in outcome had been described in terms of the different ways in which the message of the text read had been understood. We were now searching for differences in the process of learning leading to these differences in outcome. We thus had to look for the different ways in which the process leading to these outcomes had been experienced.

When scrutinising "the pool of meanings" at this particular level, a pattern of a hierarchy of similarities and differences in meaning may ultimately emerge. We

do not believe there is any uniform technique which would allow other researchers to go from "the pool of meanings" to the emerging pattern of a hierarchy of similarities and differences. It is essentially a *discovery procedure* which can be justified in terms of results, but not in terms of any specific method. In each study the discovery process will inevitably be different, depending on the specific purpose and the context of the research. Yet, whatever specific method is adopted, the crucial point is that there is what was described in Chapter 1 as 'rigorous qualitative analysis' in identifying and describing the categories of description, and in examining the relationships between them. Only then can the method be described as *phenomenography*.

Levels of Processing

In the specific case we are dealing with here, all our efforts, all our readings and re-readings, our iterations and reiterations, our comparisons and groupings finally turned into an astonishingly simple picture. We had been looking for an answer to the question of why the students had arrived at those qualitatively different ways of understanding the text as a whole. What we found was that *the students who did not get "the point" failed to do so simply because they were not looking for it.*

The main difference we found in the process of learning concerned whether the students *focused on the text in itself or on what the text was about; the author's intention, the main point, the conclusion to be drawn.* Their focal point of attention was on the pages in the first case and beyond them in the second. The first way of setting about the learning task was characterised by a blind, spasmodic effort to memorise the text; these learners seemed, metaphorically speaking, to see themselves as empty vessels, more or less, to be filled with the words on the pages. In the second case, the students tried to understand the message by looking for relations within the text or by looking for relations between the text and phenomena of the real world, or by looking for relations between the text and its underlying structure. These learners seemed to have seen themselves as creators of knowledge who have to use their capabilities to make critical judgements, logical conclusions and come up with their own ideas.

Some quotes will serve to illustrate the first way of experiencing the learning situation:

... the only thing I was thinking about was that I'd got to hurry. What happened was that I read a couple of sentences and then I didn't remember what I'd read because I was thinking all the time, "I've got to hurry to get this done"... I kept on thinking that I'd got to remember what I'd just read, but (then I would wonder) "How am I going to remember this now". "I won't remember anything" is what I thought more or less in several places.

Well I only concentrate on trying to remember as much as possible . . .

You get distracted. You think "I've got to remember this now". And then you think so hard about having to remember it: that's why you don't remember it.

In a later study carried out at Lancaster University, similar extracts were used to describe equivalent experiences (see Entwistle, 1988b).

In reading the article, I was looking out mainly for facts and examples. I read the article more carefully than I usually would, taking notes, knowing that I was to answer questions about it. I thought the questions would be about the facts in the article This did influence the way I read; I tried to memorise names and figures quoted, etc.

I tried to concentrate — too hard — therefore my attention seemed to be on “concentration” rather than on reading, thinking, interpreting and remembering — something I find happening all the time I’m reading textbooks.

These students did not try to understand the text, they tried to memorise it. Their awareness skated along the surface of the text. Their only aim was to be in a position to remember it later when they would be asked questions about it. It should be remembered that the instructions put no time constraints on the students and yet the quotations show experiences of heavy time pressure. The intention to memorise the text, however, contains a paradox. The students often have the feeling that they will not remember, just because they are trying so hard to remember. And, indeed, this is exactly what happens. (This most extreme form of concentrating on the surface of the presentation, characterised by a failure to learn due to over-anxiety to perform well, has been called *hyperintention*). So we found that many students were not even trying to understand the message and, so, in consequence, they did not understand it. On the other hand, they tried hard to remember the text, yet failed to do so. (This is because the less meaning something has for us the harder it is to remember it.)

The analysis also allowed us to identify an entirely different way of tackling the article. Some students were trying to understand the message. They were not trying to memorise the text and yet they remembered it very well. The quotes below, when compared to those above, illustrate the fact that the two groups of students seem to have been engaged in fundamentally different activities in a situation which, from the point of view of an external observer, appeared to be the same for both groups.

... and what you’re thinking about then, it’s, sort of, what was the point of the article.

... or perhaps I stopped and thought about what they were actually saying ... if there was something I thought wasn’t right, and so on. You also stop and then (wonder) if that really follows that, sort of, is it really logical, what they’ve written. That sort of thing is what you stop for.

Well, it was sort of the whole aim of (the article) — if that is what is meant. The whole aim of the article was what I was thinking of, sort of.

Or, from a Lancaster student:

I read more slowly than usual, knowing I’d have to answer questions, but I

didn’t speculate on what sort of questions they’d be. I was looking for the argument and whatever points were used to illustrate it. I could not avoid relating the article to other things I’d read, past experience, and associations, etc. My feelings about the issues raised made me hope he would present a more convincing argument than he did, so that I could formulate and adapt my ideas more closely, according to the reaction I felt to his argument.

In these cases, the text is not considered as an aim in itself (as in the earlier quotes) but rather as a means of grasping something which is beyond or underlying it—the author’s intention, what it is all about. (For a more detailed and systematic description of these two modes of learning see Marton, 1982).

The ‘depth’ dimension which was implicit in the hierarchies of learning outcomes was thus also distinguishable in the case of the student’s account of the way in which they went about the learning task. The qualitative differences in the outcome of learning were referred to as levels of outcome and the qualitative differences in the process of learning were, in these initial stages of analysis, called levels of processing. This term was chosen on the basis of a metaphorical, but only metaphorical, resemblance to Craik and Lockhart’s (1972) “levels of processing” concept. (They discussed the correlation between the likelihood of the retention of a stimulus material (usually words) and the level at which the learner has attended it, for instance, in terms of shapes, sounds or meanings.) As we shall see, the use of Craik and Lockhart’s term was subsequently abandoned, due to the false equivalence otherwise created with a quite different concept.

Relationships between Process and Outcome

Our search for differences in the process of learning was motivated by the intention to find the functional correlates of the qualitative differences in the outcome of learning. To what extent had we succeeded? First of all, it should be said that even though we were able to discern some differences in the student’s experience of learning which we believed to be fundamental, it certainly did not imply that we could use these differences to meaningfully classify all the cases. Quite obviously, there are cases, on the one hand, where the information available in the interview transcript is just not sufficient to tell whether the learner had adopted deep or surface level processing, and there are cases, on the other hand, in which signs of both strategies can be found on the same occasion. In spite of these limitations, we were still able to conclude that there was a very close relationship indeed between process and outcome. Svensson (1976, 1977, and also in Chapter 4) has convincingly argued that the main dividing line, as far as levels of outcome are concerned, lies between categories A and B on the one hand and categories C and D, on the other. This is so, Svensson says, because the fact-conclusion structure, on which the whole article is based, is understood in the first case, but not in the second. And, indeed, the figures in Table 3.1 seem to support this point.

TABLE 3.1
Relationship between level of processing and outcome
 (from Marton and Säljö, 1976a)

<i>Level of outcome</i>	<i>Level of processing</i>			<i>Sub-totals</i>
	<i>Surface</i>	<i>Not clear</i>	<i>Deep</i>	
A	0	0	5	5
B	1	6	4	11
C	8	0	0	8
D	5	1	0	6
Sub-totals	14	7	9	30

One could argue, of course, that the high correlation between process and outcome reflects the way in which the two main categories of levels of processing were found. We had indeed been looking for differences in the students' accounts of their ways of learning which would correspond to the differences found in their understanding of the article read. The close correlation between depth of processing and quality of outcome in learning has, however, during the decade that has passed since the first publication of the first results, been confirmed several times, even in investigations in which process and outcome of learning have been independently assessed (see Watkins, 1983; van Rossum and Schenk, 1984).

Levels of Processing and Approaches to Learning

As was pointed out above, there are two main alternatives for obtaining information about what the students "do" when they are trying to learn from a text. One relies more on the analysis of outcomes; the other identifies processes. So far we have identified two levels of processing on the basis of what the students said about their experience of the learning process. Then, subsequently, process was related to outcome. Svensson (1976) combined these two sources of information in making an independent and simultaneous analysis of the same set of data with the same intention—to explain the differences in outcome. The implications of his findings are discussed fully in the next chapter; here we introduce on this alternative form of analysis. Svensson concentrated first on students' accounts of what they remembered, and from the characteristics of these outcomes, he drew conclusions about the nature of the processes that accounted for what was remembered. The students' own accounts of how they perceived and experienced that process were used only to complement the analysis of performance data. Svensson's analysis again relied on the iterative procedure of rigorous qualitative analysis, but with different "pools of meaning".

In spite of this very different procedure, a similar distinction was reported but with different terminology. Svensson described the main variation in *cognitive*

approach to be between *holistic* and *atomistic*. In the holistic approach during reading, students showed indications of a general direction towards understanding the text as a whole – a search for the author's intention, relating the content to a larger context and delimiting the main parts of the text. The indications of an atomistic approach were: focusing on specific comparisons in the text, focusing on the sequence of the text, but not the main parts, memorising details and, in contrast, clear evidence of a lack of an orientation towards the message as a whole (see Svensson, 1976, p. 93).

Svensson and Marton both used the term 'approach' to describe two distinct forms of understanding, but the categories of process were different (deep/surface, holistic/atomistic). The defining features of the two distinctions were very similar, but there was an important difference in emphasis which led to the differing terminology. The first difference was in terms of epistemological assumptions. Svensson, as he explains in the next chapter, was concerned to retain evidence of both outcome and process within his initial analysis. Marton concentrated first on process, before examining relationships with outcome. The differing sets of data yielded by the two distinct research strategies produced one dichotomy (deep/surface) which emphasised *referential* aspects of students' experiences - their search for meaning or not, while the other (holistic/atomistic) concerned *organisational* aspects – the ways in which they organised the informational content of the article in their reading.

The two aspects are normally inextricably mixed. In order to understand a text, we have to integrate, to reorganize, to see the passage as a whole. It was therefore no surprise to find a close empirical relationship between the two sets of categories. In terms of outcome 29 out of 30 cases were categorised in the same way, while there was complete agreement for 25 out of the 30 categorisations of cognitive approach or levels of processing (see Marton, 1976b, p. 17).

Subsequently, the distinction between the surface and deep levels of processing was included in an SSRC research programme at Lancaster directed by Noel Entwistle (Entwistle and Ramsden, 1983, and Chapter 13). Early work there led him to the conclusion that the term "processing" was too narrow in relation to the differences in learning described (Entwistle *et al.*, 1979a). He was concerned that the crucial intentional component was not a part of its connotation, for instance. He preferred to use the term *approach*, based on Svensson's description, but retaining Marton's categories of *deep* and *surface*. This change also fitted the altered theoretical framework of the work of the Gothenburg group, which had become less and less oriented towards the human information processing school of thought. This new terminology was thus accepted, and has since been widely accepted as the most appropriate label for these qualitative differences. To avoid confusion, subsequent sections of this chapter will thus use the term *approaches to learning*.

At about the same time, Laurillard (1978) and Ramsden (1981) started to investigate approaches to learning in normal study situations in various subjects. As we shall see in later chapters, their analyses produced rather different definitions of the deep/surface dichotomy in everyday learning. The instances occasionally

came closer to the difference between an atomistic and a holistic approach than to the original distinction from which they arose. Considering this very close conceptual and empirical relatedness and considering the fact that Svensson had been using the term “approach” from the very beginning, it would appear more correct to talk about surface/atomistic and deep/holistic approaches to learning. Still, we believe that the analytic separation of the referential (“what”) aspect, which is the heart of the surface/deep dichotomy, and the organisational (“how”) aspect, which is the heart of the atomistic/holistic distinction, remains highly meaningful. Only when they are identified separately can the relationship between them be demonstrated. Indeed Roger Säljö, in Chapter 6, has done just this in showing how a certain meaning-orientation leads to a certain way of organising (segmenting, delimiting) the text and parts of it and how that way of organising the text leads to a certain referential meaning being abstracted from it.

Approaches to Learning in Normal Studies

The main investigation discussed in this chapter had been aimed at illuminating important functional differences in how students carry out their studies. This was done by setting up an experimental situation intended to resemble everyday conditions in important respects. Consequently it is a most reasonable question to ask, whether the differences found in learning experiments represent central differences in studying at university.

This question has been answered in the affirmative many times. In interviews about everyday studying Marton (1974) found that the same variation in approaches to learning could be discerned, even if the “flavour” was slightly different.

In everyday learning situations “text” takes on a metaphoric sense. The studies as a whole can be seen as the “text”, on which attention is focused, and which is entirely separate from the “real” world. The relationship between surface approach (in this case certainly less intense and immediate, and more spasmodic) and the examination seems to resemble that between the surface approach adopted in the experiment and the retention test in that context:

Interviewer How did you read these books?

S: (Well) I studied in a way that’s typical for what you do when you study for an exam. Well, basic knowledge, sort of, to get in as much as possible.

Those adopting a deep approach in their studies seem to believe, on the other hand, that the idea of these studies is to learn something about reality, to change one’s way of thinking about it. For instance:

When you read something, then just afterwards, you’re not really not quite sure about ... things. But, after perhaps a day or an hour or so, as a result of experiences, or events which jog your mind, it sort of works its way into a more solid perception in some way, which you stick to. You don’t stick to a conception just after you’ve read it, then you haven’t got any clear conception of it and can’t defend it afterwards... And other people’s comments provide an impetus to get you thinking along different lines. And they help, even if they don’t give you the idea directly.

It was found that there was a close association between a deep approach to studying adopted by freshmen and their success in social sciences courses, in spite of the fact that examination results and qualitative differences in learning can hardly be said to be described in an equivalent way. Svensson (1976, 1977) found a similar relationship between holistic approaches and examination results. He concluded, however, that the relation between approach and examination results is indirect. The direct functional relationship was between approach and study habits, (how much one studies, when one studies, etc.). A holistic orientation is not a necessary prerequisite for being successful in all the subjects at university. If students have an atomistic orientation, i.e. if they read the course literature without understanding large parts of it, then studying will be extremely boring. Consequently, students with an atomistic approach tend to acquire bad study habits; in particular they do less work, besides using an inefficient approach, and so do not succeed in their studies. The students who combine an atomistic approach with hard work are just as successful in their studies (at least in certain subjects) as are students with a holistic approach. We are not arguing that the deep/holistic approach is always ‘best’: only that it is the best, indeed the only, way to *understand* learning materials.

The ways in which students vary in the processes and outcomes of learning has emerged clearly not just in the Gothenburg studies, but also those reported in later chapters (particularly Chapters 9 and 13). The differences between approaches to learning in different subject areas and in contrasting academic contexts will be considered there, and also in the concluding, integrative chapter.

Now that we have found a difference in the way students learn which we believe to be of fundamental importance, and since this difference goes between two approaches to learning, of which one is clearly preferable to the other, should we not try to make the students who tend to adopt the less appealing approach, change to the more highly valued one?

Changing Approaches to Learning

A significant prerequisite for attempting to influence how people act in learning situations is to have a clear grasp of precisely how different people act. In our case, the problem could be phrased as: what is it that a person using a deep approach does differently from a person using a surface approach. Or to use a slightly different terminology; in what sense do these two approaches differ as “cognitive projects” (i.e. what the learner is trying to accomplish)? The picture outlined so far indicates that a significant component of a deep approach is that the reader/learner engages in a more active dialogue with the text. It is as if the learner is constantly asking himself questions of the kind “How do the various parts of the text relate to each other?”; “Is the argument consistent or are there any logical gaps?”; “How does this relate to what I already know?”; and so on. Since one of the problems with a surface approach is the lack of such an active and reflective attitude toward the text, a fairly obvious idea would be to attempt to induce a deep approach through giving people some hints on how to go about learning.

Questions in the text

In a study aimed at testing this idea (Marton, 1976a), 30 students taking a one-year integrated course in political science, economy and sociology served as participants. The learning material used was the first chapter of an introductory text-book in political science that these students were about to read in their normal studies. The participants were randomly assigned to an experimental and a control group.

The procedure adopted for influencing the approach to learning in the experimental group was to have the students answer questions of a particular kind while reading. These questions were of the kind that students who use a deep approach had been found to ask themselves spontaneously during their reading. The questions, which were interspersed between each of the five sections of the chapter, were of the following kind:

What sub-sections do you think there are in this section? (Say where they start and finish).

Can you summarise the content of each of these sub-sections in one or two sentences?

What is the relationship between the various sub-sections?

Can you summarise the content of the whole section in one or two sentences? (ibid. p. 43).

It should also be added that the design of this study included an immediate, as well as a delayed, retention test. The latter session took place on the average more than two months after the first one.

This attempt to induce a deep approach through forcing people to answer questions found to be characteristic of such an approach, yielded interesting but contra-intuitive results. At one level it was obvious that the approach taken was influenced by the treatment to which the experimental group was exposed. However, this influence was not towards a deep approach: instead it seemed to result in a rather extreme form of surface learning.

TABLE 3.2
Number of items correct in retention test
(from Marton, 1976a)

<i>Retention Test</i>	<i>Group</i>						<i>t'-test</i> <i>p <</i>
	<i>Experimental</i>			<i>Control</i>			
	<i>Mean</i>	<i>SD</i>	<i>N</i>	<i>Mean</i>	<i>SD</i>	<i>N</i>	
Immediate	34.5	8.8	15	45.7	8.5	15	0.01
Delayed	21.0	5.5	15	30.2	10.5	14	0.01

Note: Sum values for three independent raters.
SD indicates *standard deviation*

The results on both the immediate and delayed retention measurements (which were quantitative measures of knowledge in this case) showed that the control group, which had not been exposed to any attempts at influencing approach, performed significantly better (see Table 3.2).

How can this rather clear difference in performance be accounted for? The explanation, in our view, reveals a fundamental aspect of how students adapt themselves to the demands they are exposed to. What happened was that the participants invented a way of answering the interspersed questions without engaging in the kind of learning that is characteristic of a deep approach. The technique they used was simply to read the text in such a way that they were able to mention (c.f. Säljö, 1975) the contents of various parts of each section in a rather superficial way. Thus, the task is transformed into a rather trivial and mechanical kind of learning, lacking the reflective elements found to signify a deep approach. What allowed the participants to transform the learning in this way, was obviously the predictability of the task. They knew that they would have to answer questions of this particular kind, and this allowed them to go through the text in a way which would make it possible to comply with the demands (i.e. summarising the various parts of the sections, stating the main point in them, dividing them into sub-sections, etc.) without actually going into detail about what was said. As Marton (1976a) concludes, "this process can be seen as a special case of the common human experience of transformation of means into ends" (p. 47). The questions which were intended as means of helping the students to adopt a deep approach, instead became the objective towards which the learning was geared. In this transformation, the attempts to deal with them in an expedient way became detrimental to learning.

The outcome of this study raises interesting questions about the conditions for changing people's approach to learning. The "demand structure" of the learning situation again proved to be an effective means of controlling the way in which people set about the learning task. Actually, it turned out to be too effective. The result was in reality the reverse of the original intention when setting up the experiment. The very predictability of the "demand structure" in our view played the central role in generating the paradoxical outcome. Therefore, in the next study to be described, while again using the learners' expectations of the questions subsequent to the reading as the independent variable, we tried to restrict their expectations, more indirectly, to the kind of questions, instead of to exact questions.

Questions after reading

In the study reported by Säljö (1975; see also Marton and Säljö, 1976b) 40 university students were divided into two groups. The factor varying between the two groups was the nature of the questions that the groups were asked after reading each of several chapters from an education text-book. One set of questions was designed to require a rather precise recollection of what was said in the text. The questions focused on were, for example, enumerations and listings of causes and consequences of certain events, factual information such as names, percentages, terminology, and so on. The idea behind this was, of course, to see if the participants

would adopt a surface approach to meet the demands they could predict they would have to face after reading each successive chapter.

In the second group, the questions were directed towards the major lines of reasoning. The demand for remembering exact information of the kind mentioned above was much lower, and instead the participants had to give evidence that they had understood how the conclusions followed from the developing argument, and they also had to provide judgements as to whether the reasoning seemed consistent and correct. In addition, they also had to recall the text and give a short summary of its main points.

After reading a final chapter, both groups were exposed to both kinds of questions and they were also required to recall the text and summarise it in a few sentences. The outcome here thus served as the major dependent variable of the study. The results show that a clear majority of the participants reported that they attempted to adapt their learning to the demands implicit in the questions given after each successive chapter. This could be seen both in students' subjective reports about how they set about learning as they went from one chapter to the next, and in the way the final chapter was recalled and the questions answered.

Was it, then, possible to influence people to use a deep approach in this more indirect way? Unfortunately, the answer to this question cannot be a simple 'yes'. The crucial idea of this study, that people would respond to the demands that they were exposed to, was verified. In the group which was given "factual" questions this could be clearly seen and, as expected, both their recalls and the way that the questions were answered, as well as the reports about how they set about learning, showed that they reacted to the questioning through adopting a surface approach. However, in the other group, the reaction did not simply involve moving towards a deep approach. Some students did, others did not. A fundamental reason underlying this was differing *interpretations of what was demanded of them*. Only about half the group interpreted the demands in the way intended. The other students handled the task in very much the same way as was found in the study by Marton described above. By focusing their attention on the most conspicuous tasks, they were able to foresee what they would have to face after reading (recalling the text and summarising it in a few sentences). These participants then *technified* their learning, again concentrating solely on perceived requirements. They could summarise, but not demonstrate understanding. This concept of technification requires some additional comments, since it reflects a rather fundamental observation that has been made in the studies where attempts have been made to influence the approach people use.

A common idea in these studies has been to observe and describe, in as great a detail as possible, what characterises a deep and a surface approach respectively in terms of the kind of learning that people engage in. On the basis of this knowledge, it ought to be possible to influence people who do not spontaneously adopt a deep approach to behave in a way similar to those who do this in a given situation. For example, since it could be observed that it was characteristic for students using a deep approach to clarify for themselves in an explicit way the main points and lines of reasoning in the text that they were reading, it is very

common to assume that if one makes other students perform the same kinds of activities, this would mean that they too would be using a deep approach. However, it is obvious from the two studies reported here, and from other similar investigations (see, for instance, Dahlgren, 1975), that this kind of logical reasoning does not always lead to the expected results when applied to human behaviour. It is important to realise that the indicators of a deep approach, isolated in the research, are symptoms of a rather fundamental attitude towards what it takes to learn from texts. Thus, one cannot treat these observations on what characterises a deep approach as pointing to casual factors that can be isolated and manipulated through rather simple means to achieve the desired end. Instead, if we take the study by Säljö as an example, what happened was that some students *made it an end in itself* to be able to give a summary of the text after each chapter. In this way, their learning was geared towards the objective of fulfilling this particular demand, and again the task can be solved through *mentioning* the various parts of the text at a very superficial level. This is thus an example of the process of *technification* of learning resulting in poor performance. The functional mechanism underlying this process is that the perceived demands become so predictable that students believe they can handle them through a very shallow interaction with the text. They simply use their knowledge about what is going to happen later to economise on their efforts. This, as was pointed out earlier, is a very common human reaction, and we should not be surprised to find students behaving in such a way.

Taken together, these studies illustrate that although in one sense it is fairly easy to influence the approach people adopt when learning, in another sense it appears very difficult. It is obviously quite easy to induce a surface approach and enhance the tendency to take a reproductive attitude when learning from texts. However, when attempting to induce a deep approach the difficulties seem quite profound. How are we to understand this?

If we return for a moment to the nature of this distinction as it emerged in our studies, the fundamental difference between approaches has been described as one of whether students interpreted the text itself as what was to be learned, or conceived the text as the *means* through which they sought to grasp the meaning underlying the words and so to change their conceptions about historical developments, economic processes, or whatever. The fact that even when students have been encouraged through various means to adopt a deep approach redefine the situation in a way which will make it expedient to use a surface approach, should tell us something about the strong mechanisms operating within educational contexts in support of this reproductive mode of learning.

Learning and Motivation

One of the factors contributing to the partial failure of these attempts to induce a deep approach by manipulating the "demand structure" of the learning situation is the relation between the learners' motives and the ways they go about learning. Learning or reading out of interest, a wish to find something out (i.e. due to intrinsic motivation), can reasonably be expected to be linked with a deep approach. On the other hand comments from students who had adopted a surface approach

showed that they had tried to memorise the text because they felt that this was required of them. Surface approach and the motive of fulfilling the demands raised by others (i.e. extrinsic motivation) seem to go together. This relation between approach to learning and motivation to learn was the topic of Fransson's (1977) study.

His premise was that intrinsic motivation is not so much something one creates but rather something one finds. If we want to utilise people's intrinsic motivation, we must focus on what they are interested in and link the study material to it. Once again, the material used was in the form of a text, but the text was chosen in such a way that it could be considered to be of immediate interest to one of the groups that participated in the experiment, but not the other. The text was about the examination system in the Education Department. The group which was assumed to be interested in the text was made up of first-year students in the department. The other group consisted of sociology students who were not taking Education. In addition, each group was randomly divided into two subgroups. One of these subgroups was subjected to treatment that was assumed would create extrinsic motivation while the other subgroup was left alone. This subgroup were told that after they had read the text they would be asked to give an oral report and that their report would be video-recorded (the equipment was prominently displayed). The source of the extrinsic motivation was thus the utilization of people's fear of "making a fool of themselves". When the text had been read, each group was, in fact, treated in the same way; they were all asked to write down what they remembered of the text.

After the sessions the students were asked to fill in questionnaires indicating how interested they were in the text, and how anxious they felt while reading it. Of course, not all the students from the Department of Education were interested in the text, nor were all the sociology students uninterested in it. Not all the subjects found the idea of being video-taped particularly threatening, but on the other hand some of those who were simply told they would have to write down what they could remember after having read the text, became very nervous. Some of the subjects who did not usually become nervous in other situations of a similar nature, did not become nervous this time either while others did. All this came to light during the interviews that were held after the experiment. The analysis showed that the main effect on approach to learning came not from the experimental situation *per se*, but from the reported experiences of the students – whether they *felt* interested, threatened, or anxious. Then the results produced a clear picture. Intrinsic motivation, absence of threat (extrinsic motivation) and absence of anxiety, both independently and together, were associated with a deep approach. Threat (extrinsic motivation), anxiety and absence of intrinsic motivation similarly correlate with a surface approach. In fact *all* students in the 'relaxed' condition who were interested and non-anxious used a deep approach, while *all* those who felt themselves threatened, anxious, and uninterested adopted a surface approach when they read. The conclusion that can be drawn from this experiment is that if we want to promote a deep approach, we should above all keep in mind the students' own interests at the same time as we should try to eliminate the factors that lead to a surface approach (irrelevance, threat and anxiety).

Approaches to Learning and Conceptions of Learning

There is, however, a paradoxical circular relation between approach to learning and motivation to learn. As the results quoted in the previous section suggest, intrinsic motivation (interest) seems to lead to a deep approach and extrinsic motivation (concern with demands) to a surface approach. On the other hand, adopting a surface approach means that the learner focuses on the "text" or tasks in themselves and not on what they are about. But it is hardly possible to be interested in a "text" unless one is paying attention to what it is about. Not being motivated by an interest in the "text" tends thus to lead to the adoption of a surface approach, and the adoption of a surface approach tends to block any interest in the "text". In order to unravel this circularity we may need to move to a superordinate level of description.

Säljö's (1975) above-mentioned study showed that the two groups participating in the experiment on the whole behaved differently because of the difference in the kind of questions they expected to follow their reading. On the other hand, there was a great variation both in the process and outcome of learning within the group which was given questions intended to induce a deep approach. In our view, this was due to the participants' different perceptions of what was required of them. The differences between contrasting experimental groups reflect the effect of context on learning. The differences within the same treatment, however, logically must originate from a variation in something which the participants "brought with them" to the experiments. Their perceptions of the task reflect their past experiences of similar situations, and so mirror differences in their preconceived ideas of what it takes to learn.

Säljö (1979) carried out an interview study in which he asked a group of adults what learning meant to them. Analyses of the transcripts produced five qualitatively different conceptions, to which a sixth has subsequently been added (Marton *et al.*, 1993). Learning was seen as:

1. A quantitative increase in knowledge.
2. Memorising.
3. The acquisition, for subsequent utilisation, of facts, methods, etc.
4. The abstraction of meaning.
5. An interpretative process aimed at understanding reality.
6. Developing as a person (added subsequently).

According to the previous line of reasoning, in a learning situation, which is the "same" from an external point of view and to which the participants adopt different approaches, we would expect the observed variation in approaches to be closely linked with a variation in conceptions of learning held by the participants on a more general level.

This was indeed one of the questions which van Rossum and Schenk (1984) set out to illuminate. They used an open-ended questionnaire to identify the conceptions of learning held by the students. The answers showed a clear correspondence to Säljö's categories. The students also had to read a piece of argumentative prose, give an account of its content, and report about their experience of learning. The design of this part of their study was very similar to

that of the Gothenburg experiment already extensively discussed in this chapter. And van Rossum and Schenk also found that students' experiential accounts of how the learning task was carried out could be classified in terms of deep and surface approaches respectively. Furthermore a close correlation between conceptions of learning and approaches to learning was found (see Table 3.3).

TABLE 3.3
Relation between conceptions of learning and approaches
(from van Rossum and Schenk, 1984)

<i>Conceptions</i>	<i>Approach to learning</i>		
	<i>Surface</i>	<i>Deep</i>	<i>Sub-totals</i>
1. Increase in knowledge	6	0	6
2. Memorisation	19	4	23
3. Fact acquisition for utilisation	8	7	15
4. Abstraction of meaning	1	11	12
5. Understanding reality	1	12	13
<i>Sub-totals</i>	35	34	69

In addition to matching the expectation of a relationship between conceptions and approaches, Table 3.3 seems to support Säljö's (1979) remark that the main dividing line goes between the third and fourth conception. Or modifying Säljö's earlier view somewhat, we can say that the first and second conception on the one hand, and the fourth and fifth, on the other, seem to have similar relations to approaches to learning. The third conception appears to be intermediate between the others. Though the conceptions refer to what learning is seen as, on the next subordinate level within the two pairs, there are "what-how" relations between the conceptions. The quantitative increase in knowledge (the first conception of learning) is achieved reasonably by memorisation (second conception). On the other hand, we improve our understanding of reality (fifth conception) by abstracting meaning from what we read, see, hear (fourth conception). The second and the fourth conceptions thus represent the "how" aspect of the two different answers to the "what" questions reflected by the first and fifth conceptions respectively. The first one of these two pairs of conceptions is closely linked to the surface approach, not only empirically, but conceptually as well. Indeed, it appears as a generalised version of the special experience of learning, which has been termed "surface approach" (orientation towards memorisation, seeing learning as a more or less passive transmission of what is on the paper into the head of the learner).

A similar relationship seems to hold between the second pair of conceptions (the fourth and fifth) and the deep approach. Especially in the context of normal studies, the distinction between conception (aiming at a better understanding of reality by abstracting meaning from what is presented) and approach (focusing attention on what the presentation refers to) seems to become blurred. As we pointed out earlier, a deep approach, in the context of everyday studying, primarily refers to the realisation of the fact that the studies one is engaged in deal with some aspect of the "real world" and thus by studying, one is trying to improve one's understanding of it. This is a view more or less identical with the one expressed in the fourth and fifth conceptions of learning; its investigation became the focus for much of the subsequent work in phenomenography (Marton, 1994).

The most recent work on conceptions of learning has been examining possible cultural differences in the conceptions themselves and in the relationships between them. For example, there has been considerable interest in what has been called the 'Chinese paradox' – that Chinese students seem to put much effort into memorising the material they are asked to learn, and yet they also seem to acquire a good understanding of it. Interviews with students in Hong Kong (Marton *et al.*, 1995) have indicated that high school students there seem first to seek understanding, and then try to commit that understanding to memory – a phenomenon which is not dissimilar to the ways British students carried out their revision for Finals (Chapter 9 and Entwistle, in press).

In this latest work, Marton has explored not only the previous depth dimension in the hierarchy of conceptions, but also a temporal dimension which separates *acquiring* from the subsequent *knowing* and *using*. The depth dimension was also elaborated to distinguish between committing *words to memory*, committing *meaning to memory*, understanding *meaning (text based)*, and understanding the *phenomenon (reality based)*.

Committing meaning to memory is more complex than committing words to memory, because meaning implies a distinction between words and meaning. Understanding (meaning) is more complex than remembering (meaning) because both imply permanence, but the former implies an act of constituting (or finding) meaning as well. (Marton *et al.*, 1995, pp.28-29)

In the responses of students, there also seemed to be a shift away from surface approaches and towards deep approaches, as pupils met increasing volumes of material in the later years of secondary education.

Conclusion

In Chapter 2 it was argued that the outcome of this kind of learning should be described in terms of the conceptions of the phenomena learned about which have been reached through learning. As one of the most immediate implications for teaching, on the other hand, attention is drawn in Chapter 15 to the importance of knowing what conceptions students already hold about the phenomena to be introduced to them. And it is exactly in transitions between preconceived ideas of the phenomena and an improved understanding of those phenomena, where the most important form of learning in higher education is to be found. Chapters 1

and 2 both showed, however, that this conceptualisation of learning differs radically from the one underlying the way in which learning has been dealt with in psychological research.

The research described in this book is an attempt to draw attention to outcomes of learning which seem to coincide with the aims of university lecturers (as seen in Chapter 1), and yet are currently under-emphasised in the teaching methods most commonly adopted, and not reached by the many students who still conceptualise learning in terms of reproductive requirements. To the extent to which students and teachers share an understanding of what it means to learn, we should expect improvements in the quality of learning in higher education – and indeed in education in general. The third part of this book explores the implications of this research in more detail, after looking at the experiences of teaching and assessment which students report.

Acknowledgements

The research reported here was financially supported by two separate grants from the Swedish Council for Research in the Humanities and Social Sciences, one to each of the authors.