The M.A.K. Halliday Library Functional Linguistics Series

M.A.K. Halliday
Edited by Jonathan J. Webster

Aspects of Language and Learning

Springer

M.A.K. Halliday
Author

Aspects of Language and Learning

Edited by Jonathan J. Webster

Springer

Author
M.A.K. Halliday
The University of Sydney
Sydney, NSW
Australia

Editor
Jonathan J. Webster
City University of Hong Kong
Kowloon, Hong Kong SAR
China

ISSN 2198-9869 ISSN 2198-9877 (electronic)
The M.A.K. Halliday Library Functional Linguistics Series
ISBN 978-3-662-47820-2 ISBN 978-3-662-47821-9 (eBook)
DOI 10.1007/978-3-662-47821-9

Library of Congress Control Number: 2015943443

Springer Heidelberg New York Dordrecht London

Printed on acid-free paper

Springer-Verlag GmbH Berlin Heidelberg is part of Springer Science+Business Media (www.springer.com)

Contents

Chapter 1
Language, Learning and 'Educational Knowledge'

In these lectures, I plan to roam around a fairly broad terrain. First, I will spend some time venturing into history, in various senses of the word 'history', beginning with the life history of the human child and his experiences as a baby before we might think he has any language at all. Next, I will look at the history of the language that we use in order to encode our knowledge, particularly our scientific knowledge—the history of scientific English, in other words. Then, putting these two together, I shall ask how a growing child comes to master this kind of language and put it to his own use as a means of learning. After that, I shall explore one or two questions of the relationship among language, education and culture, again taking the language of science as the focal point for the discussion, and finally, I shall try to draw these various themes together to construct a sort of language-based picture of experience—a linguistic interpretation of how we learn and how we learn how to learn. In all of this discussion, I am deliberately putting language at the centre of the stage. It is hard for any of us to keep language in the focus of attention for very long: we tend to fly off from it in all directions, to study thought processes, behaviour patterns, aesthetic values and so on. But I shall try to resist this tendency and shall ask you to **think linguistically**, that is, to use your conscious and unconscious understanding of language as a means of thinking about the world, and in particular—since this is my unifying concern in these lectures—to use language as a tool for exploring how people learn. The more deeply we understand the processes of learning, the more likely it becomes that we shall be able to help people to learn more effectively.

So in bringing together the concept of 'language' and the concept of 'learning', I am not focusing exclusively on questions to do with learning language, whether mother tongue, first language, second language or any others. Learning language is obviously an essential part of the picture; but I want to see it as part of a broader conception in which learning—all learning—is itself linguistic activity. Whatever you learn, you are engaged in language; learning involves 'languaging', if you will allow me this license with English (we have no word for the general process that lies behind the traditional four skills). But learning is always learning **something**, and since I cannot range over the whole terrain of human knowledge, I shall use scientific learning as a kind of focus for the general discussion. Insofar as I shall be considering language in the context of education, and of educational knowledge,

M.A.K. Halliday, *Aspects of Language and Learning*,
The M.A.K. Halliday Library Functional Linguistics Series,
DOI 10.1007/978-3-662-47821-9_1

I shall use scientific discourse as the principal domain in which to explore and from which to draw my illustrations.

But first we should start at the beginning—or one of the beginnings. In fact, there are two beginnings to language and learning: one is in the origins of the human species, the phylogenetic beginnings, but these unfortunately we cannot observe. We cannot go back in time to observe how language evolved in the early history of the human race. The other is the birth of a human child, the ontogenetic beginnings, and since these can be and have been observed, I would like to take this as the starting topic of this lecture. The question I want to address is: if we take a **linguistic** view of learning, how does this process start? And how then does it evolve into a special kind of learning process that we refer to as 'education'?

Until fairly recently, it was customary to think that a child has no language until he starts to say things in his mother tongue, some time on in the second year of his life. It always struck me as surprising that people would continue to hold this view when the evidence seemed so very clear that children are communicating from birth. They are born as communicating beings, and even if they do not begin to speak in English, or in Chinese, or in any other one of our adult languages till they are 12–15 months of age or more, they are typically responding to language almost from the day they are born and *using* language to communicate with from somewhere around 6–9 months of age. What they use first, however, is not our adult language, but rather a little language they create for themselves: we could refer it as a child tongue, not a mother tongue. But they cannot create it **by** themselves; no one could, because all 'languaging' is interactive, and this child tongue, or 'protolanguage', is created by a child together with those around him—mother, father, older brothers, sisters, etc.—when they listen to the child and understand what he is saying. The adults conduct their part of the dialogue in their own adult language; perhaps a little bit simplified although not necessarily so. Now, the point I want to establish first of all is that this protolanguage, as exchanged between the child and those who share in his acts of meaning, is already a highly effective medium of learning. Let me give you some examples from my own observations of Nigel, when he was aged 9–14 months, to show how he was already using his language to make sense of the world—to build up a picture out of his small but growing body of experience.

Text 1.1: Nigel at nine months—using language to learn

Nigel had just learnt to sit up on his own, and was now ready to start meaning in earnest.

He had a little floppy rabbit; I was holding it on my hand and stroking it, then making it jump in the air. When I stopped, Nigel put out his hand, and touched the rabbit, firmly but without pushing it. It was a gesture which meant 'go on, do that again'—the same meaning that he has later to express vocally as "ùh".

He had two other gestures. If he meant that he wanted something, he would grasp it firmly in his fist, without pulling it towards him, and then let go. If he meant he did not want it, he would touch it very lightly and momentarily with the tip of his finger.

These gestures were true acts of meaning. Nigel was not acting directly on the objects; he was addressing the other person, enjoining him to act.

In addition to the three meanings conveyed by gesture, Nigel had two other meanings which he expressed vocally. The two expressions were almost the same: one was "èu", the

other, slightly higher pitched at its starting point, was "èu". The first meant 'let's be together', and was used in conversation: "Nigel!"—"èu"—"*There's* a woozy woozy woozy"—"èu", and so on ad inf. The other meant 'look—a commotion', and was the successor to " ' ", the tiny high-pitched squeak. Nigel used it to express interest in his surroundings, especially that part of the surroundings that went into violent movement, like a flock of birds taking off.

This was the opening scene of Nigel's language.

These were two, out of a little set of ten or twelve 'signs'—sounds, or gestures—used regularly, systematically, and in a meaningful way.

Now, this is not yet a kind of English; Nigel is not yet trying to learn English, and you could not tell from **his** contributions to the dialogue what his mother tongue is going to be. (Some children do use more sound imitations at this stage, but that is just a difference of strategy; Nigel on the other hand invented his own sounds and in fact used some sounds that do not occur in English at all.) When my Chinese colleague Qiu Shijin studied the early language development of children growing up in Shanghai, she found the same pattern: first a protolanguage which did not necessarily sound like a form of Chinese at all and then, from around the middle of the second year, the move into the mother tongue, in this case Shanghainese (Qiu 1985). But with the Chinese children, as with the English children, she found that the protolanguage was being used in very much the same range of meanings: it was what they too were constructing as a tool with which to learn.

But you may well want to ask: learning what? You may doubt that what is taking place is in any real sense learning at all. I think it is, and I think we can be rather precise about what it is that is being learnt. This is brought out by my own observations, by Qiu Shijin, and again by my colleague Clare Painter in her detailed study of her first child, Hal (Painter 1984). Consider now some further examples from Nigel.

Text 1.2(a): Nigel at ten and a half months

Ten and a half months

Nigel was sitting on my knee. On the table in front of us was a fruit bowl with an orange in it. Nigel struggled to reach it.

"nà nà nà nà," he said. It meant 'I want it', 'give it to me'.

I gave him the orange. He made it roll on the table; it fell off.

"nà nà nà nà," he said again.

When the game was over, he got down, crawled away and disappeared along the passage, going boomp-boomp-boomp as he went. Then silence. His mother began to wonder where he was.

"Nigel!" She called.

"è—e—eh" It was his special response to a call: 'Here I am'.

"Where is he?" said his mother. "Nigel!"

She went to look for him. He was standing, precariously, by the divan, looking at his picture cards that were hanging on the wall.

"dòh," he said as she came in. It meant 'hullo—shall we look at these pictures together? "dòh … dòh"

"Are you looking at your pictures?" his mother asked him.

"dòh … dòh"

Text 1.2(b): Nigel at twelve months

Twelve months

Nigel and I were looking at his book together.

Nigel took hold of my finger and pressed it lightly against one of the pictures. "èya," he said.

The meaning was clear: 'you say its name'. "It's a ball," I said.

"è—e—eh" Nigel gave his long-drawn-out sigh, meaning 'yes, that's what I wanted you to do'. He was pleased that his meaning had been successful, and he repeated the procedure throughout the book.

Later he was looking at it all by himself.

"dò … èya … vèu"

This was Nigel's first complex utterance and the only one for many months to come. But it made excellent sense. He had picked up the picture book, opened it at the ball page and pointed at the picture. It was just as if he had said, in so many words, 'Look, a picture! What is it? A ball!'

In the first of these examples, he is exploring the environment and his own relation to it: expressing curiosity, interest, pleasure and so on. This is language to think with, and we can in fact trace the path Nigel took, step by step, from those early protolinguistic utterances to the naming of objects around him and from there to the part of English **grammar**—"transitivity"—that enables him to combine the names into complex representations of experience, such as the clause *strange man gone* said spontaneously at 20 months when he saw someone pulling funny faces. There is a direct link from protolanguage to the use of the grammar and vocabulary of adult language as a **key** with which to interpret experience.

In the second example above, Nigel is not only exploring, but also exploiting: using language to get things done. This is language not to think with but to act with. He may be getting his mother or his father to play with him, or he may be asking for a drink or a favourite toy—but like every human baby, he knows perfectly well that you cannot get possession of an object by talking to that object: you have to talk to a person, who will then (if well disposed) pass the object to you. So this kind of language is essentially interpersonal language, the language of demands (and also of offers; children like to give as much as to get), and once again, it is possible to follow it through as it evolves into the adult language, this time not through the use of words as names of things but through intonation patterns and eventually into the grammar of mood.

So these two motifs—(1) language as a way of thinking about the environment and (2) language as a way of acting in the environment, via the people in it—are present from the start, from the very beginnings of the language of a human child. In systemic functional theory, we refer to these motifs as 'ideational' and 'interpersonal'. These are two of the three most general functions of language, and because these three functions are the underlying principle on which all human languages are built, we refer to them as 'metafunctions', to distinguish them from functions in the sense of just different uses of language. The ideational and interpersonal motifs—language to think with and language to act with—are more than

simply uses of language: they are the fundamental organizing concepts around which the whole of language has evolved over the past two to five million years. And the first significant thing that the human child learns through language is that this is what language is: it is a way of interpreting and of controlling the world he finds around him.

If he learnt nothing else in his first encounters with language, that would be a major achievement. But of course he does learn something else, because within these broad motifs, he is already making meaningful distinctions: between 'I want that' and 'I don't want it', between 'give me something' and 'play with me', between 'where are you?' and 'there you are!', and so on. We can observe all these distinctions evolving in the language Nigel uses himself; if we went further and observed what he understands of what is said to him in return, we should get an even richer impression of what it is that he has learnt through the medium of language—all this before he has ever said a word of his mother tongue at all.

Before following through on the experience of the child, I need to make one thing clear. I have spoken of the child as 'learning through language', and this implies a distinction between (1) learning a language and (2) using that language in order to learn. This is an important distinction and a valid one for our analytic purposes. But we should also acknowledge that from the point of view of the child, these are one process not two. The child is simply learning, and learning as we have stressed is a linguistic activity. He does not distinguish between learning the word *bus*, when he gets to the naming stage, on the one hand, and on the other hand using that word to interpret a particular experience—some noisy monster lumbering past him in the street—and to relate this to **other** experiences which are alike in some respects and different in others (e.g. the same monster when he happens to be sitting comfortably inside it). He may **practise** the word *bus*, out of context; so to that extent he could be said to be 'learning language'—but even there it is doubtful whether there is a very clear distinction between rehearing the sound and rehearsing the experience with which that sound is associated.

The reason we need to remind ourselves of the child's point of view in this regard is simply this: we tend to think as adults—if we think of language at all—that language is simply a passive element in the learning process. We think of our experience as something given and the language as a convenient mirror in which that experience gets reflected. Language 'reflecting' experience is indeed a common metaphor that people use. But it is a misleading metaphor, nevertheless. Language does not passively reflect experience; rather, it creates or 'construes' experience: language is an active participant in the semiotic cycle. It is language that enables us to order and interpret the flux of events in which we find ourselves, the 'mush of general goings-on' as Firth used to call it, so that instead of defining language as that which encodes experience, we can almost turn things round and define experience as the order that is created out of chaos by means of language. To say this is to take the first step towards a language-based theory of learning.

What kind of language are we talking about, in this regard? Clearly not the conceptually complex, tightly constructed metalanguages of science and philosophy (we shall come back to those later). Here, we are concerned with the language of a very small child; even before it has any words, or any structures, language—his protolanguage—is already at the foundations of his learning. Once the child does move into the mother tongue, his language becomes recognizable as language in the adult sense: it gains a 'lexicogrammar', an organization in the form of words-arranged-in-structures like that of *man clean car*. We can easily adapt this to an adult model, something like *there's a man cleaning his car*; indeed, the adult who is being addressed, who was also sharing the experience, typically rewords the child's observation (and did, in fact, on this occasion) in a related adult form: 'Yes, there was a man cleaning his car'. The child gets confirmation, in this way, that his construction of the experience matches that of the adult world: in other words, he has 'got it right'.

The adult's response as we have said is in normal English, with all the grammatical words and morphemes added in place; gradually, the child will incorporate these into his discourse too. But it will still be the ordinary everyday discourse of the home and the neighbourhood: the unselfconscious, unplanned and unwritten-down language of daily life. It is this that is at the foundation of our knowledge, in at least three different ways which we shall need to explore later on. But first I need to draw attention to one further aspect of the child's linguistic experiences. We have noted the twofold character of language in relation to the environment: that we think with language, and we act with language—it is the continuation of the two that enables us to learn. Here are some further examples, this time from the phase of Nigel's transition into the mother tongue.

Text 1.3: Nigel's early mother tongue, in mathetic and pragmatic functions

Pragmatic

chuffa stúck	'the train's stuck; help me to get it out'
high wáll	'let me jump off and you catch me'
háve it	'I want that'
play ráo	'let's play at lions'
squeeze órange	'squeeze the orange'
bounce táble	'I want to bounce the orange on the table, can I?'

Mathetic

Clever boy fix roof on lòrry	'this clever boy fixed the roof on the lorry'
Dada come bàck … Dada come on fast chùffa	'Daddy's come back; Daddy came on a fast train'
too dàrk … open cùrtain … lìght now	'it was too dark; you've opened the curtains, and it's light now'

Now, in all these instances, Nigel learns because the language refers: there is something going on, out there, that creates a context and so enables the child to construe an appropriate meaning. But in addition to this representational property of language, at the same time it is encoding, or encapsulating, for him a great deal of other potential information, about the material and behavioural environment, about the social structure with its interpersonal relationships and about his own place in this complex scheme of things. It does this not just by referring to what is going on but by participating; language enables people to **act out** the social and physical processes in which the child himself is involved. This point is one that is extremely difficult to illustrate, since it depends on the continuity of linguistic interaction over a long period, the ongoing dialogue in which the child is engaged throughout all his early years with those who are looking after him. But let me try.

Nigel has a game, high wall, in which he throws himself full tilt off whatever object he is standing on and has to be caught by his father before he hits the ground. This started harmlessly enough in the park near his home, where there was an old ruined abbey of which only some stunted walls remained; some just one or two courses of stones above the ground—these were the 'low walls'—and some at chest or shoulder height (the adult's, that is)—these were the 'high walls'. Soon any object, such as the arm of a settee or an upturned suitcase, could function as a 'wall' for 'jumping off'. As the months went by, Nigel's 'high walls' got higher and higher, until he was participating himself from a height way above his father's reach—and only barely waiting to ensure his father was there to catch him. This simple expression *high wall* had long since ceased to function as a name for a class of objects; but as a consequence of its extended use, it got a variety of different responses which provided Nigel with a great deal of miscellaneous information from which he was able to learn. Here are just a few:

High wall!

(response:) That's not a high wall; that's only a low wall

> Objects (e.g. wall) and their properties (high/low); properties continuous—no clear boundary high/low, one pole is negative ('*only* low'; one couldn't say 'that's not a low wall; that's only a high wall').

No that wall's too high—you'll hurt yourself

> Concept of 'too much' of some property—undesirable, and reason why—in what respect—undesirable (danger and consequences)—to himself

No Grandad doesn't play high wall; you're too heavy for him

> This time 'too much' is Nigel himself—too heavy; but the focus is on Grandad. He is old; old people are different—not so strong; these are things they cannot do.

Come on now, enough—it's time to go home and have your tea

Things have a time and a place—there is order in daily life, and an obligation to keep to that order, even if it means giving up on something you were enjoying.

Of course—not all is learnt from one event; but the child has hundreds of similar linguistic experiences.

Now, my examples are taken from English, so I should stress that there is no priority to this or that particular language as a vehicle for learning, for the child's construction of reality in this way. Every language is as good as every other language at transmitting the sociocultural environment in which a child is growing up. This after all is the context in which all languages have evolved, and it is important to make it explicit that languages did evolve, slowly and naturally over many hundreds of generations: nobody planned or designed them. The languages in which children live the first years of their lives, the world over, have evolved along with society itself, along with humanity itself, as the discourse of the family and the community—fluent, highly structured, rich in resources, and always relevant to its context. It is the language that people use to live.

So by the time a child is 3–4 years of age, he is well aware of the power of language in enabling him to control his environment—again, in both senses, of thinking about it and acting on it. Ruqaiya Hasan's pioneering research into the linguistic interaction between mothers and children at 3½ years of age shows how penetrating and persistent are the children's demands on language—and therefore their demands on whoever is talking and listening to them, as many a tired mother with attest!—their constant questioning and arguing and experimenting with language (Hasan 1983–1986). Here is an example from one of Hasan's transcripts:

Text 1.4: Example of child–mother interaction from Ruqaiya Hasan's research 'The role of everyday talk between mothers and children in establishing ways of learning'

M. D'you want to have a look at him?
K. Yeah.
M. Poor little moth!
K. Poor little moth! Poor little moth!
M. D'you want to put him on your hand?
K. Mmm.
M. There … what will we do with him?
K. Put him in there—em, leave him there.
M. Okay, we'll leave him up on the side.
K. [sadly] I want him.
M. There'll be other moths. I guess he was happy when he was flying around, and now he's not any more, is he?
K. [crying] no.
M. He can't hurt any more either. I'm afraid little moths do that all the time.
K. Was that a baby one or a big one?

M. No; moths are, em—[banging noise; to Ruth] Hey! Don't bang the oven!—Moths are quite old. When they're little, they're little worms, and—well you know the book about the hungry caterpillar that you've got?
K. Yes
M. He ate and ate until he became a big fat caterpillar.
K. Mm.
M. And then he built himself a cocoon, and then he came out of the cocoon; then he was a butterfly. Well moths also come out of cocoons; but they … they only come out of their cocoons to lay their eggs, and after they lay their eggs they die.
K. Why?
M. Well because they're very old by then. It's the end of their life when they come out of the cocoon, so if he hadn't died there he would have died in a few days anyway.
K. Why did he fall into the hot water?
M. I don't know. But I guess he probably flew low over it, and with hot steam coming up from the water it might have made it hard for him to fly.

Note in particular how much the child is asking questions and the complexity of the mother's answer, because this turns out to be a feature of their discourse which disappears when they get into the educational setting. As the researchers discovered in the 1970s, in Canada and Britain, and subsequently in the USA, children do not ask questions in school—not, at least, in the classroom, when they are occupying their status as pupils. But in the home they do, all the time, and they do listen to the answer, as a parent soon finds out if she gives a different answer next time the question comes round: 'But you said it wouldn't spoil; now you say it will!?' What Ruqaiya Hasan's work is revealing is how important this interaction is, the child's questions and the adult's responses, to the child's progress in learning. Learning is not an individual process; it is social and interactive—we shall always recognize this fact if we interpret learning in terms of language, as I am doing here, since languaging is also a social process. You may talk to yourself, in odd moments—many of us do. But it is only because you talk to others that you have any language to talk with, and even if you are talking to yourself, it is likely that you are simply playing both parts in a game that is still basically dialogic.

The time comes when children are sent to school. This usually happens around the age of five. There are of course very varied practices, in different countries, both in the age at which school is started and in the kind of preschool experience which the children have had—kindergartens, preschools, crèches, Filipino nannies and the like. But school proper usually begins around five or six, and if you ask what the children are sent there to do, apart from keeping out of mischief, the traditional answer was always 'to learn to read and write'. In other words, they go to school first of all for linguistic reasons: to learn to use language in new ways.

Now, there was a time, even as late as the early 1960s, when I was first involved in language education as director of the London project which produced *Breakthrough to Literacy*, when reading and writing were scarcely thought of as language. In fact, at that time, the word 'language' hardly figured at all in educational discussions, and educators like Bereiter and Engelmann in the USA explicitly denied that reading and writing were parts of a child's general experience of language development. We worked hard in the London project to put literacy in its

linguistic context, to bring out the fact that learning to read and write is an extension of the power that children already have, the rich resource of language in its spoken form. Writing extends their language into new contexts and new functions; but it can only do so because the language is already there.

Once this principle came to be accepted that reading and writing were still kinds of languaging, it became possible to ask the question: what is the relation of writing to speaking and of written language to spoken language? These are not simple questions—at least they may be simple questions, but they do not have simple answers. English is still English, whether you are writing it or speaking it, and Chinese is still Chinese. But is it the same English? The same Chinese? At certain times in history, writing and speech have grown very far apart; if, for example, you had lived in China in the Qing dynasty, you would have written only wényán 文言, a form of Chinese as different from the spoken varieties of the language as Latin is from Italian or French. There is much less distance between the two today and much less distance also between spoken and written English. At the same time, written English is definitely not spoken English written down. You **can** take a tape recorder, record conversation and then transcribe it; this is often done, for research purposes, and I have been using examples in my talk today—it is perfectly intelligible, but it is not what you would typically write. The same is true in Chinese, where we also have some transcriptions of spoken dialogues. Spoken and written languages tend to use different forms of expression and different ways of organizing discourse.

These differences are often referred to as 'conventions'. These are said to be different 'conventions' governing written language. This is true enough, in a sense; it is an understandable way of putting it. But at the same time, this formulation has a certain drawback: it suggests that these are some sort of arbitrary rules, a kind of etiquette or table manners—which you have to follow in order to write acceptably. Now, it is true that many written registers do, like those of speech also, acquire certain more or less ritual features—the forms for beginning and ending a letter, for example—which seem to have no function other than signalling that it **is** a letter: a purely indexical function, in other words. But instances of this kind are not the norm. Most features that distinguish some or other functional variety of a language are not arbitrary but motivated; they bear some meaningful relationship to what that variety has evolved to achieve. This will be important when we come to look more closely into the forms of written English, especially scientific English. Of course, it is always possible to have too much of a good thing: such features, even though they are there for good reason, may always be overused and so tend to **become** more or less ritualized. But even then it is important to understand why it was that they came into use in the first place.

To take a very simple and obvious example, we introduce punctuation into written text—capitals, full stops, inverted commas and the like—for a very good reason: we cannot write down the intonation, the rhythm or the voice quality. In speech, all these phonological devices serve to signal the grammatical structure; punctuation evolved to take over this function in writing.

A more complex example would be the order of expressions of cause and effect: in speaking, we usually put the effect first and the cause afterwards, whereas in writing, it is usually the other way round. Reasons for this are rather more complicated; but we will pick them up at a later stage.

Written language evolved, in the course of history, in certain human cultures where people had settled down, produced or at least husbanded their own food instead of hunting and gathering it and begun to engage in new forms of activity for which they needed to keep records: trade, with property owning and taxation; calendrical and astronomical studies; and perhaps certain ceremonial and religious activities. These new **forms** of social and economic life led to convergence between two different semiotic practices: language on the one hand, which up to that time had been purely spoken language, and various forms of pictorial representation on the other, such as paintings on the walls of caves. Over a period of time, the pictures came to stand not for things and events but for words—items in the language, and when that happened, you had reached the stage of writing. Writing then evolved along different lines with different languages; that is not something that will concern us here. What will concern us, however, is what lies behind these new activities for which writing, and written language, first evolved.

Keeping records is a simple enough idea; it makes us think of documentation and filing systems, invoices and bills of lading, tables of facts and figures—or in these days of the electronic office, computers and floppy disks. But think of these still as forms of language and compare them with the language you speak. They represent very different kinds of **demand** that we make on language. When language is used to store and accumulate information in this way, it is functioning very differently from the way it is functioning in spontaneous, natural conversation, or even in oral narrative or religious ritual. It represents a different form of knowledge: knowledge that is stored, but can be consulted; that is frozen, but can be accumulated. This is a new way of construing the social and material environment. The new forms of written language which evolved, over many centuries, in the context of these new functions constituted in important ways a different theory of reality: complementary to that one based on spoken language, neither more true nor less true, but bringing a different dimension to our picture of the world we live in.

So when children go to school and become literate, they are not simply extending their linguistic powers into a new channel; they are learning a new way to learn. They are moving from the world of 'commonsense knowledge', which is typically spoken, into the world of 'educational knowledge', which is typically written down. This distinction, between commonsense knowledge and educational knowledge, plays a significant part in shaping the social system and in defining the individual's access to and participation in social processes. Commonsense knowledge is fluid; it lacks clearly defined stages or clear-cut boundaries (which does not mean, on the other hand, that it lacks structure). Educational knowledge is compartmental, subject (as Bernstein described it) to classification and framing: that is, it is organized into subjects and presented in measured doses according to a master plan we call a curriculum. But the fact that it **can** be classified and framed in this way means that it is a different **kind** of knowledge; commonsense knowledge is

learnt ('picked up' is a favourite metaphor), and educational knowledge has to be taught. Another way of expressing this is to say that educational knowledge is 'technical' knowledge, which means that such knowledge has to have its own special language—so we talk about 'technical terms', or if we do not understand them, we call them 'jargon', which simply means technical language used by somebody else.

But behind the institutionalized learning, and the technical vocabulary, in which educational knowledge is enshrined is a further, deeper distinction which again we can see most clearly if we approach it through language. The specialized words that are listed in a technical dictionary are the clearest signs, the indices of a particular discipline; they are a necessary part of the register. But behind these lies the more general phenomenon of a technical **language**: a style of meaning that is expressed as much through the grammar as through the vocabulary. In fact, we cannot really distinguish vocabulary from grammar in language; the two are simply the same thing seen from different ends, so that it is more helpful to refer to a single, unified concept of 'lexicogrammar'. Technical language, then, embodies different grammar and vocabulary because it embodies distinct semantic patterns, different ways of organizing meaning. Its characteristic lexicogrammatical features are not just different ways of saying the same things; they are ways of saying different things, which cannot be said in the language of everyday discourse. So when children go to school to learn different forms of language, they are not simply being dressed up in neat new linguistic clothing. They are being initiated into new ways of learning and new forms of knowledge.

Up to now, we have assumed a simple dichotomy: commonsense knowledge is associated with speaking and listening, and educational knowledge is associated with reading and writing. This explains why the first task of a child coming into school is to learn to read and write. But expressed in this way, it is obviously oversimplified. For one thing, we need to make clear what distinction it is that we are talking about. On any particular occasion, we are operating in a given **channel**: the language is being transmitted either as sound, which we can hear, or as writing, which we can read (of course both may be involved, as when someone reads aloud from a written text). But the channel must be distinguished from the **mode**: that is, whether the language—the meanings expressed, and the grammar and vocabulary in which they are represented—is that which is typically written, or that which is typically spoken. In the case of the mode, unlike the channel, we have a continuum: in a clearly **written** mode, or in a clearly **spoken** mode, or in something that lies somewhere in between. And here of course the picture gets complicated, because what distinguishes the written from the spoken mode is not one simple feature but a whole syndrome, an assemblage of different features (mainly grammatical features) that typically go together but can in fact vary independently of one another. And different kinds of writing—for example, different subjects in the secondary school curriculum—use different mixtures.

Scientific English, of the kind that we find in learned journals, typically uses a rather extreme variant of the written mode. Even in textbooks and popular journals, we find a strongly marked package of features characteristic of writing, which