

has served its purpose. In the case of writing the text is unilaterally produced and remains as a permanent record. But it is still only a discourse trace, and what is meant by it has to be inferred by interpretation, and this inevitably raises the question of how far this interpretation corresponds with the intentions that informed the discourse which gave rise to the text in the first place.

In the normal circumstances of use, of course, we only pay attention to text in order to realize its discourse function and so we tend to think of the two as the same thing, as indeed do some linguists, and talk about the meaning of a text as shorthand for what it means to us or what it might mean to the writer or speaker. But although we normally experience text as part of the discourse process, it is perfectly possible to focus on the text alone. This is after all what proofreaders generally do when they scrutinize a piece of writing to identify typographical errors, or wordings that do not conform to established code conventions. More interestingly, texts can also be subjected to close analytic study to find out patterns of actual usage which those producing them may be quite unaware of, an area of linguistic description we shall be returning to in a later chapter. But before we get on to this, it will be necessary to say a little more about how the concepts introduced in this chapter—semantic and pragmatic meaning, sentence, utterance, text, context, and discourse—figure in a general model of communication.

2

Communication

Grammar and communication

As was pointed out at the beginning of the first chapter, linguists have traditionally focused their attention primarily on the internal properties of languages, on how meaning is **formally encoded** in lexis and grammar. The description of such properties can be said to be an account of what people know of their language, an account of their **linguistic competence**. It was this competence that enabled us to describe the grammatical features of the utterance in the train in the preceding chapter. Of course that description made use of terminology which might well be unfamiliar with people competent in English: knowing the grammar of a language is not the same as knowing how to describe it—that is the business of the grammarian. But the point is that anybody competent in English would recognize that the utterance conforms to the **encoding conventions** of the standard language, that it exemplifies a well formed sentence in English, which it would not do if it had taken the form, for example:

They has it in a safe place put and it will not to find.

The original utterance, we can say, exemplifies a possible sentence in English and this second one does not.

So on hearing this remark in the train, one judgement we can make on the basis of our linguistic competence is whether it is grammatically and lexically possible or not, that is to say, in accordance with the encoding conventions of the language. But we can also recognize degrees of possibility. So if the utterance had been:

They has put it in a safe place and it will not be found.

We would recognize that this is a closer approximation to the standard language. Indeed, an utterance like this might be perfectly possible in a non-standard variety of English. We have to note that we make judgements about what is possible in English only by reference to some agreed norm or other. So in the English that I myself spoke as a boy, *you was* and *he were* are grammatically possible, whereas the standard forms *you were* and *he was* are not. Usually, however, for reasons we shall touch on later, it is the standard that serves as the norm of **grammatical well-formedness**.

However, although linguists may take particular notice of such features of a linguistic form, ordinary language users generally focus their attention on what is meant by the use of such forms, and this, as we have seen, involves taking context into account. So another judgement that might be made about an instance of language use is whether and to what extent it is **appropriate** to its context. The issue here is not whether a piece of language is grammatically or lexically well formed as a sentence or not but how far it is pragmatically effective as an act of communication.

Three kinds of pragmatic meaning

Acts of communication can be pragmatically effective in three ways. First, the language can be used to talk about something, to express a **proposition** of some kind. This involves making a connection with context in such a way as to make an appropriate **reference**. Let us suppose, for example, that the following is part of the text of a conversation:

The taxi will be here in a quarter of an hour.

The definite article provides an appropriate contextual connection since it signals that what is being referred to is common knowledge between the people engaged in the conversation (*The taxi we talked about, the taxi you asked me to order ...*). The adverb *here*, and the prepositional phrase *in a quarter of an hour*, locate the utterance in a particular context of place and time that the participants share. Without these contextual co-ordinates, the referential possibilities of this expression would, of course, be endless (*here* could be anywhere, *in a quarter of an hour*, any

time). To the extent that the conversationalists recognize the co-ordinates, appropriate reference is achieved.

One kind of pragmatic meaning, then, is reference, but it is not the only kind. The person who utters this expression is not just referring to a future state of affairs, but is doing so in the process of performing a kind of **communicative or illocutionary act**. This, for example, might be a promise (*I have arranged everything. The taxi will be here in a quarter of an hour*), or advice (*You should pack your bags. The taxi will be here in a quarter of an hour*). Whether the utterance is intended as having the **illocutionary force** of promise, or advice, or anything else, will again depend on the context of knowledge and assumption that the speaker assumes to be shared.

To the extent that this expression makes an appropriate connection with context, then, it can be taken as having a certain reference and a certain force. But the speaker is doing something else as well. The speaker is not just acting, but acting upon the other person, to bring about a certain state of mind or course of action. In performing an illocutionary act, she is also bringing about a **perlocutionary effect**. In promising, she or he may be intending to reassure the other person (*Don't worry, I have arranged everything ...*), or the advice may be meant to stir him/her into action (*Hurry up!*)

We shall be returning to these kinds of pragmatic meaning in a later chapter, but the point to be made at present is that our recognition of the extent to which a piece of language keys into context appropriately so as to acquire a certain reference, force, and effect requires more than linguistic knowledge. We may be able to assign semantic meaning to a particular expression as a sentence but be quite unable to make pragmatic sense of it as an utterance, as an instance of language use. Knowing what a sentence means is one thing, but knowing what is meant by an utterance is another. The two are of course related: our knowledge of the encoded possibilities in the language delimits the range of pragmatic interpretation. So it is that in our example, our linguistic knowledge about what the word *taxi* denotes as a **lexical item** in English, and knowing that the definite article *the* signals shared knowledge, provide crucial pointers as to what the speaker is referring to. Our ability to infer what somebody is

possible
1 appropriate
doing with the language on any particular occasion will to some degree depend on our knowing what is encoded in the language itself, even though it is not determined by it. Our ability to communicate, our **communicative competence**, therefore, incorporates both a knowledge of what is encoded as possible in the language and a knowledge of how these encodings are used appropriately in context.

Four aspects of communicative competence

3 Dell Hymes
What is **possible** in the language, in the sense of what can be encoded in it, and what is appropriate in its use are two factors that are included in the well-known account of communicative competence proposed by the American scholar Dell Hymes. He includes two other factors in his account. One of these is **feasibility**. Our competence in a language, he argues, includes our recognition of the extent to which a particular expression can be readily processed. A piece of language might be a grammatically and lexically possible encoding in the language, but difficult to decipher. An example often given of this is of the multiple embedding of one structure in another. Thus the sentence *The dog chased the cat* can be grammatically rewritten, or transformed, into *This is the cat the dog chased*. Another sentence, *The cat killed the rat*, can then be embedded to form *This is the rat the cat the dog chased killed*. This already poses something of a problem, but if we go on to embed another sentence—*The rat ate the corn*—the resulting structure becomes almost impossible to process: *This is the corn the rat the cat the dog chased killed ate*. Though this can be taken as a possible structure in English, since it conforms to encoding conventions, it will also be recognized as relatively unfeasible. This is, of course, an extreme example, and one might object that although linguists might amuse themselves by inventing such structures, they are so perversely complex that it is unlikely that they would ever actually occur. But the point about feasibility can also be illustrated by kinds of structure that are by no means uncommon, namely those we recognize as **ambiguous**. **Ambiguity** occurs when two distinct structures converge into one single sequence of sentence constituents, as in the well-known example:

To visit aunts can be boring Aunts who visit can be boring

Visiting aunts can be boring.

The ambiguous sentence is grammatically entirely well formed but lacks feasibility in that it can be decoded in two completely different but equally valid ways. This, it is true, is an invented example but, unlike the embedded structures considered earlier, ambiguities of this kind are well attested in actual texts. Here are two examples:

Two sisters were reunited after eight years at a checkout counter.

The stolen painting was found by a tree.

Perfectly possible expressions in a language may, then, be relatively lacking in feasibility. Conversely, an expression may be entirely feasible in that it is easy to process, but not be grammatically well formed and so not possible with regard to the encoding conventions of the language. This is commonly the case with the utterances of spoken text which tend to dispense with grammatical features that are surplus to pragmatic requirement, although Mr Jingle, a character in Dickens' *The Pickwick Papers* perhaps takes this tendency to extremes:

'Mr Pickwick—deepest obligations—life preserver—made a man of me—you shall never repent it, sir.'

'I'm happy to hear you say so,' said Mr Pickwick. 'You look much better.'

'Thanks to you, sir—great change—Majesty's Fleet—unwholesome place—very,' said Jingle.

Here, Jingle's fragments of talk are easy to process: they are feasible enough.

The general point to be made, then, is that there is no direct correspondence between feasibility and grammatical well-formedness, any more than there is between grammatical well-formedness and appropriateness to context. The extent to which an expression is possible, processable, and appropriate are three separate things. One of the central questions that discourse analysis needs to consider is how they are related in actual use, and we shall be taking up this issue in later chapters.

But meanwhile we have a fourth factor to consider. Hymes suggests that to be communicatively competent in a language we must be capable not only of recognizing how far an instance of it is possible, feasible, and appropriate but also how far it is actually performed or produced, the extent of its frequency of occurrence. The language code provides the basic resources for communication, but we do not, of course, make use of these resources in equal measure. We may, for example, know the words *walk* and *perambulate*, and both will be listed in a dictionary of English, but we also know that *walk* is of relatively common usage, and *perambulate* relatively rare. Similarly, we may know that a structure like *he had been being careful* is possible in English but also that it does not actually occur all that often.

In these examples, we are consciously aware of relative frequency. But in innumerable other cases, we are not. Thus we may know perfectly well that the words *big* and *large* are semantically synonymous, as are the words *little* and *small*, but, if asked, we would not be able to say with any certainty which word in each pair is more frequently used. Nor what other words they tend to keep company with. In the usage of any language there is a tendency for words and structures to co-occur in patterns of relative probability. Some co-occurrences (or collocations) are fixed and familiar, as in idiomatic phrases like *to and fro*, *ebb and flow*, *by and large*, *by hook or by crook*, *by fair means or foul*, and so on. Some we know as proverbs: *Too many cooks spoil the broth*, *A rolling stone gathers no moss*, and so on. For the most part, however, we are unaware of the patterning of language in the texts we produce. With the development of corpus linguistics over recent years, such patterning can now be described with great precision.

By means of the computer, linguists can now provide details of the frequency and co-occurrence of words and structures which are simply not accessible to intuition. So they can tell us things about the language we produce that we did not know in the sense that we are not consciously aware of them. But since our usage is informed by these frequencies and patterns of co-occurrence, we must know them in some other subliminal sense and this must be a covert part of our competence in the language. We have

procedural knowledge of them, and what linguistic description can now do is to make this knowledge overt and declarative. We shall be considering corpus descriptions in more detail later (Chapter 8).

Hymes identifies the possible, the feasible, the appropriate, and the performed as four distinct aspects of communicative competence, and in the analysis of language in use we can focus attention on any one of them. Given an instance of language, we can, as we have already seen, concern ourselves with how it exemplifies what is formally possible in the language in terms of its encoded grammatical and lexical properties with little if any consideration of any other aspect. If, however, we are interested in the psycholinguistic issue of how these formal properties are actually processed in the mind, then we shall naturally focus attention on feasibility, bringing in the other aspects only to the extent that they are relevant to such processes. A focus on the appropriate will bring context into consideration as a crucial factor in the achievement of pragmatic meaning, and here we are centrally concerned with discourse, as this was earlier defined, with how expressions are assigned a particular reference, force, and effect. As pointed out earlier, this process of meaning making leaves a textual trace, and the patterns of frequency and co-occurrence that this reveals can be independently described without reference to the particular contextual conditions in which the text was produced, or the pragmatic meanings that were achieved in producing them. Here the focus of attention is on the fourth of the Hymes factors, the language that is actually attested as having been performed.

Conclusion

When we use language in the normal circumstances of everyday life, all of these four factors come into play in complex and interdependent ways. But the determining factor is that of contextual appropriateness. It is the discourse we want to achieve that regulates how we draw on the encoded resources of the language to make a text. The language that we produce and process is not designed to demonstrate what is possible or feasible or commonly occurring but to realize a discourse purpose. How

then is this done? What exactly does it mean to use language appropriate to context? These are questions that will be taken up in the next chapter.

3

Context

Conditions of language use

We experience language not as something separate but as an intrinsic part of our everyday reality. We do not, in normal circumstances, just display our linguistic knowledge: we put it to use to give shape to our internal thoughts and to give external expression to our communicative purposes. Indeed, we usually find it difficult to display our knowledge in dissociation from these natural conditions of use. So although you might be extremely competent in a particular language, if somebody were to ask you to show your competence by saying something in that language you would in all likelihood be at a loss to know what to say. We only produce language when we have the occasion to use it, and the occasions for use occur in the continuous and changing contexts of our daily life.

Context and shared knowledge

These contexts can be thought of as **situations** in which we find ourselves, the actual circumstances of time and place, the here and now of the home, the school, the work place, and so on. When people talk to each other, they will naturally make reference to what is present in such situations—present in the sense of both place (here) and time (now):

The chalk is over there.

Pass me the tape measure.

There's a page missing.

I like the look of that.

Is that the time?