

Another line of attack on a definition of this sort would start by questioning the notion of *language understanding*. How is this to be construed? A reasonable, and perhaps the only plausible, response would be to say that to understand an utterance is to decode or calculate all that might reasonably have been meant by the speaker of the utterance (cf. Strawson, 1964). Here the notion of speaker-meaning is best explicated, once again, by reference to Grice's concept of meaning-*nn*, for we are interested only in the inferences overtly and intentionally conveyed. So the definition really amounts to: pragmatics is the study of the role context plays in speaker- (or utterance-) meaning. But since we have failed to produce a clear notion of context, what we include in context is likely to be whatever we exclude from semantics in the way of meaning relations. And so we seem to be back to the idea that pragmatics concerns whatever aspects of meaning are not included in semantics. (In which case, it may be objected, the problematic concept of context has been gratuitously introduced.) Certainly the two definitions ((9) and (12)) are not far apart; but it might be claimed that at least the one that focuses on the nature of context makes clear that one of the goals of a pragmatic theory should be to explicate that nature.

Let us now turn to one of the definitions most favoured in the literature, albeit mostly in an implicit form. This definition would make central to pragmatics a notion of **appropriateness** or **felicity**:

- (14) Pragmatics is the study of the ability of language users to pair sentences with the contexts in which they would be appropriate

Such a definition should have a nice ring to it, from the point of view of those who wish to place pragmatics on a par with other aspects of linguistic inquiry. For if pragmatics is to be considered an aspect of linguistic competence in Chomsky's sense, then like other aspects it

constrained only by consistency ... The consistent sets of propositions that comprise contexts are to be interpreted as the unique speaker's own 'commitment slate' in the sense of Hamblin (1971: 136)" (Gazdar, 1979a: 130); while in a work concerned with literary interpretation, "a context is construed as a 'complex event', viz. as an ordered pair of events of which the first causes the second. The first event is - roughly - the production of an utterance by the speaker, the second the interpretation of the utterance by the hearer" (Van Dijk, 1976: 29). But clearly a general theory of aspects of context relevant to production and interpretation must be broader than either of these.

must consist of some abstract cognitive ability. Further, such a view provides a nice parallel with semantics: for just as a semantic theory is concerned, say, with the recursive assignment of truth conditions to well-formed formulae, so pragmatics is concerned with the recursive assignment of **appropriateness-conditions** to the same set of sentences with their semantic interpretations. In other words, a pragmatic theory should in principle predict for each and every well-formed sentence of a language, on a particular semantic reading, the set of contexts in which it would be appropriate.

Such a view enjoys much support, not only among linguists (see e.g. Van Dijk, 1976: 29; Allwood, Andersson & Dahl, 1977: 153; Lyons, 1977a: 574) but also among philosophers (originally Austin, 1962 and Searle, 1969). But unfortunately it is beset with many problems. First, as we shall see, most definitions of pragmatics will occasion overlap with the field of sociolinguistics, but this definition would have as a consequence exact identity with a sociolinguistics construed, in the manner of Hymes (1971), as the study of **communicative competence**. Secondly, it requires a fundamental idealization of a culturally homogeneous speech community or, alternatively, the construction of *n* pragmatic theories for each language, where *n* is the number of culturally distinct sub-communities. For example, in a village in South India, where there may be say twenty distinct castes, a single honorific particle may have just one meaning (e.g. speaker is inferior to addressee) but have twenty distinct rules for its *appropriate* usage: members of one caste may use it to their cross-cousins, others only to their affines, etc. (for the actual details see Levinson, 1977). Thirdly, speakers of a language do not always comport themselves in the manner recommended by the prevailing mores - they can be outrageous, and otherwise 'inappropriate'. So such a definition would make the data of pragmatics stand in quite an abstract relation to what is actually observable in language usage, whereas for many linguists one of the major contributions of pragmatics has been to direct attention once again to actual language usage. Fourthly, it seems to be a fact that pragmatic constraints are generally defeasible, or not invariable. So suppose we attempt, for example, to phrase accounts of the pragmatic notion of presupposition in terms of appropriateness conditions, we shall find that they wrongly predict conditions of usage. For instance, the verb *regret* seems to presuppose that its complement is true, and

so we could try the following characterization: the sentence *John doesn't regret cheating* can only be used appropriately in contexts where it is known (or believed) that John cheated. But unfortunately we can then easily imagine a context in which that sentence might be appropriately used, in which it is *not* assumed that John cheated: for example, you thought he had cheated, asked me whether he now repents, but I tell you he never did, and persuade you accordingly, and then I say *So John doesn't regret cheating* (Gazdar, 1979a: 105). The problem is quite general: when the pragmatic implications of an utterance do not match the context, then in general the utterance is not treated as in any way infelicitous or inappropriate or bizarre – rather the pragmatic implications are simply assumed not to hold. But the use of the notion of appropriateness-conditions would in that case simply make the wrong predictions.

Finally, and decisively, there is another problem with the use of the notion of appropriateness as a primitive or basic concept in pragmatics. For, there is a widespread phenomenon that Grice has called **exploitation**: in general, if there is some communicative convention C that one does A in context Y, then suppose instead one does B in Y, or does A but in context Z, one will not normally be taken to have simply violated the convention C and produced nonsense. Rather, one will generally be taken to have exploited the conventions in order to communicate some further pertinent message. For example, if I normally doff my cap only to my superiors, but on an occasion doff my cap to an equal, then I can effectively communicate an ironic regard, with either a joking or a hostile intent (the non-linguistic example is intended to draw attention to the great generality of the phenomenon; for a study of a particular linguistic practice and the jokes thus made available, see the study of the openings of telephone calls by Schegloff (1979a)). Irony is a good example of this exploitation and the difficulties such usages pose for a pragmatic theory based on appropriateness, for ironies take their effect and their communicative import, and thus their appropriateness, precisely from their inappropriateness. So the problem is in general that, in being grossly inappropriate, one can nevertheless be supremely appropriate! True, one may need some notion of 'normal practice' (in preference perhaps to appropriateness) even to describe such phenomena, but it would be a mistake to limit pragmatics to the study of that normal practice or appropriateness. Pragmatics should be

much concerned precisely with such mechanisms whereby a speaker can mean more than, or something quite different from, what he actually says, by inventively exploiting communicative conventions. We must conclude that, despite its initial attractions, the proposal that pragmatics be based on a notion of appropriateness should be discarded: language usage is too elastic to allow a pragmatic theory to be based on such a concept. If instead one accepts that the goal of a pragmatic theory is to predict the meaning, in the broad Gricean sense, of an utterance in a specified context, then none of these difficulties arises.

At this point, someone searching for a simple definition of pragmatics is likely to be exhausted. One possibility is to retreat to an *ostensive* or *extensional* definition, i.e. simply to provide a list of the phenomena for which a pragmatic theory must account (cf. Stalnaker, 1972). Such a definition might run as follows:

- (15) Pragmatics is the study of deixis (at least in part), implicature, presupposition, speech acts, and aspects of discourse structure

This list would certainly provide a reasonable indication of some central topics in pragmatics, but the definition scarcely helps those unfamiliar with these topics and has other more serious drawbacks. For in common with all extensional definitions, it provides no criteria for the inclusion or exclusion of further phenomena that may come to our attention; at best one can say that what warrants pragmatic treatment for some new topic is simply linguists' consensus based on intuitive 'family resemblance' to more familiar pragmatic topics. But surely such intuitive resemblance must be based on some underlying implicit common themes – our difficulty is that when we try to spell these out we arrive at the various problems experienced in our earlier attempts at definition.

At this point, we might step back and attempt some conceptual clarification from other angles. Katz & Fodor (1963) tried to delimit the scope of semantics by a boundary drawing exercise: the 'upper bound' of semantics was provided by the borders of syntax and phonology, and the 'lower bound' by a theory of pragmatics, understood as a theory of contextual disambiguation. Using the same strategy, we could say that the upper bound of pragmatics is provided by the borders of semantics, and the lower bound by sociolinguistics (and perhaps psycholinguistics too). Indirectly, we have already

explored this way of thinking in our consideration of the proposal that pragmatics is 'meaning minus semantics', and the idea that some distinction from sociolinguistics is necessary was responsible for some of the dissatisfaction with a number of the definitions above. We have already seen the difficulties of drawing a neat dividing line between semantics and pragmatics; given the cross-cutting criteria of conventionality and non-defeasibility (see again Table 1.1), the best strategy seems to be to restrict semantics to truth-conditional content. Assuming that this is accepted (and many linguists would resist it), we can turn our consideration to the lower bound, the border between pragmatics and sociolinguistics. Here things are even more problematic. Let us take two paradigmatic kinds of sociolinguistic phenomenon, and ask how they fall with respect to two of our definitions of pragmatics, namely, the most restrictive and the broadest definitions. Consider **honorifics**, most simply exemplified by the polite singular pronoun of address in European languages (like *vous* vs. *tu* in French – let us call this the V vs. the T pronoun). There are a number of sociolinguistic investigations of such honorifics and their usage (e.g. Brown & Gilman, 1960; Lambert & Tucker, 1976). If we take the view that pragmatics is concerned only with grammatically encoded aspects of context (see definition (8) above), then we might propose a tidy division of labour between pragmatic and sociolinguistic accounts of honorifics: pragmatics would be concerned with the *meaning* of honorifics (e.g. with the specification that V encodes that the addressee is socially distant or superior), while sociolinguistics would be concerned with the detailed recipes for *usage* of such items (e.g. the specification that amongst some segment of the speech community, V is used to aunts, uncles, teachers and so on, or whatever the local facts are). Such studies would be exclusive but complementary. Now, however, consider what happens if we take pragmatics to be the study of the contribution of context to language understanding: suppose normally an aunt gives her nephew T, but on an occasion switches to V, then in order to predict the intended ironic or angry meaning, a pragmatic theory must have available the detailed recipe for usage that tells us that V is not the normal usage, and thus not to be taken at face value. So on this broader scope for pragmatics, the neat division of labour collapses – pragmatic accounts of language understanding will at least need access to sociolinguistic information.

Taking another paradigmatic kind of sociolinguistic phenomenon, namely the variable phonological realizations associated with social dialects (see e.g. Labov, 1972a), let us ask how our definitions of pragmatics treat such facts. On the most restrictive view, that pragmatics is concerned with linguistically encoded aspects of context, such facts would seem to lie outside the purview of pragmatics. Such an exclusion would rely on the restricted sense of *encoding* that is required, *inter alia*, that the significances in question are (a) intentionally conveyed (and we can now say, meant-*nn*) and (b) conventionally associated with the relevant linguistic forms. For, as we noted, the association of particular accents (realized by proportions of phonological variables) with particular social or geographical communities is generally not part of an intentional message (Labov (1972a) argues that such variables are only very partially under conscious control), nor are such social significances associated with linguistic forms by arbitrary synchronic convention so much as by regular historical and social process. However, if we take the broader scope of pragmatics represented by the definition that relates context to language understanding, there may well be cases where sociolinguistic variables would be of relevance to language understanding. Gumperz (1977), for example, has argued that such variables can be used to invoke domains of interpretation, e.g. to mark transitions from chat to business. Or, consider the case of a comedian telling a joke about a Scotsman, an Irishman and an Englishman – he may well rely on mimicked features of accent to track which protagonist is talking. In short, drawing a boundary between sociolinguistic and pragmatic phenomena is likely to be an exceedingly difficult enterprise. In part this can be attributed to the diverse scopes that have been claimed for sociolinguistics (see Trudgill, 1978: Introduction), but in part it comes about because sociolinguists are interested in inter-relations between language and society however these are manifested in grammatical systems: sociolinguistics is not a component or level of a grammar in the way that syntax, semantics, phonology and, quite plausibly, pragmatics are.

Another angle from which we might attempt conceptual clarification of the issues is to ask: what are the **goals** of a pragmatic theory? The term *goal* is used here in the special way current in linguistic theorizing, and is to be distinguished from the ultimate goals or motivations that might prompt interest in a theory. Those ultimate

motivations will be the subject of the next section, but here we are interested in exactly what it is that we expect a pragmatic theory to do. One abstract way of thinking about this is to think of a pragmatic theory as a 'black box' (an as yet unexplicated mechanism), and to ask: what should be the input to such a theory, and what should be the output (or: what is the theory meant to predict, given what particular information)? We can then think of a theory as a *function* in the mathematical sense, which assigns one set of entities (the *domain*) to another set of entities (the *range*), and the question is, what are these sets of entities? Thinking the same way about syntax, we can say that a given set of rules (a syntactic analysis) is a function whose domain is the set of possible combinations of morphemes in the language L , and whose range has just two elements, denoting the grammatical and the ungrammatical in L ,¹⁸ or thinking about semantics, we might say that a semantic analysis of L has as its domain the set of well-formed sentences of L , and as its range the set of semantic representations or propositions representing the meaning of each of those sentences. It is by no means so obvious what the input and output of a pragmatic theory should be.

Two authors, at least, have been explicit on this subject. Katz (1977: 19) suggests that the input should be the full grammatical (including semi-antical) description of a sentence, together with information about the context in which it was uttered, while the output is a set of representations (or propositions) which capture the full meaning of the utterance in the context specified. Since a sentence plus its context of use can be called an utterance, Katz's suggestion amounts to the idea that a pragmatic theory is a function whose domain is the set of utterances and whose range is the set of propositions. Or symbolically, if we let S be the set of sentences in language L , C the set of possible contexts, P the set of propositions, and U the cartesian product of $S \times C$ - i.e. the set of possible combinations of members of S with members of C , and we let the corresponding lower case letters stand for elements or members of each of those sets (i.e. $s \in S$, $c \in C$, $p \in P$, $u \in U$):

$$(16) \quad f(u) = p \text{ (or: } f(s, c) = p)$$

¹⁸ Or, in alternative parlance, a function from the set of morpheme combinations to the well- vs. ill-formed sentences, or a function that maps the set of morpheme combinations into the well- vs. ill-formed sentences. See Allwood, Andersson & Dahl (1977: 9ff) for elementary exposition.

i.e. f is a function that assigns to utterances the propositions that express their full meaning in context

Gazdar (1979a: 4-5), on the other hand, wishes to capture the ways in which utterances *change* the context in which they are uttered; he shows that Katz's formulation is incompatible with that goal, and therefore suggests instead:

$$(17) \quad f(u) = c \text{ (or: } f(s, c) = c)$$

i.e. f is a function from utterances to contexts, namely the contexts brought about by each utterance (or: f assigns to each sentence plus the context prior to its utterance, a second context caused by its utterance)

The idea here is that the shift from the context prior to an utterance to the context post utterance itself constitutes the communicational content of the utterance. It suggests that pragmatic theory as a whole should be based on the notion of context change (see some applications in Chapters 4 and 5 below).

Both these formulations are consistent with the definitions of pragmatics as 'meaning minus semantics' or as the contribution of context to language understanding. Our other definitions might require slightly different formulations. For example, where pragmatics is construed as the study of grammatically encoded aspects of context, we might want to say:

$$(18) \quad f(s) = c$$

where C is the set of contexts potentially encoded by elements of S

i.e. f is a theory that 'computes out' of sentences the contexts which they encode

Or, alternatively, where pragmatics is defined as the study of constraints on the appropriateness of utterances, we could say:

$$(19) \quad f(u) = a$$

where A has just two elements, denoting the *appropriate* vs. the *inappropriate* utterances

i.e. f is a theory that selects just those felicitous or appropriate pairings of sentences and contexts - or identifies the set of appropriate utterances

Or, where pragmatics is defined ostensively as a list of topics, we could say:

$$(20) \quad f(u) = b$$

where each element of B is a combination of a speech act, a

set of presuppositions, a set of conversational implicatures, etc.
i.e. *f* is a theory that assigns to each utterance the speech act it performs, the propositions it presupposes, the propositions it conversationally implicates, etc.

Clearly, there are other possibilities, and it is far from obvious, at this stage of the development of the subject, just which of the many possible formulations is the best. But as the subject develops we can expect researchers to be more explicit about exactly how they expect a pragmatic theory to be formulated.

Let us sum up the discussion so far. We have considered a number of rather different delimitations of the field. Some of these seem deficient: for example, the restriction of pragmatics to grammatically encoded aspects of context, or the notion that pragmatics should be built on the concept of appropriateness. The most promising are the definitions that equate pragmatics with 'meaning minus semantics', or with a theory of language understanding that takes context into account, in order to complement the contribution that semantics makes to meaning. They are not, however, without their difficulties, as we have noted. To some extent, other conceptions of pragmatics may ultimately be consistent with these. For example, as we noted, the definition of pragmatics as concerned with encoded aspects of context may be less restrictive than it seems at first sight; for if in general (a) principles of language usage have as corollaries principles of interpretation, and (b) principles of language usage are likely in the long run to impinge on grammar (and some empirical support can be found for both propositions), then theories about pragmatic aspects of meaning will be closely related to theories about the grammaticalization of aspects of context. So the multiplicity of alternative definitions may well seem greater than it really is.

In any case, we embarked on this definitional enterprise with the warning that satisfactory definitions of academic fields are rarely available, and the purpose was simply to sketch the sorts of concerns, and the sorts of boundary issues, with which pragmaticists are implicitly concerned. As was suggested at the outset, if one really wants to know what a particular field is concerned with at any particular time, one must simply observe what practitioners do. The rest of this book will largely be concerned with an overview of some of the central tasks that pragmaticists wrestle with.

Before proceeding to a discussion of the motivations that lie behind the growth of the field in recent years, it would be as well to clarify the role that pragmatics might be seen to play within linguistic theory as a whole. There is no doubt that some workers see pragmatics as a running commentary on current linguistic methods and concerns, and its role as the juxtaposition of actual language usage with the highly idealized data on which much current theorizing is based. Viewed in this way, attempts to delimit pragmatics in the ways explored above would make little sense; pragmatics would not be a component or level of linguistic theory but a way of looking afresh at the data and methods of linguistics. In that case, pragmatics would be a field more akin to sociolinguistics than semantics. It is therefore worthwhile seeing that, whatever the merits of this view, there is a need for a kind of pragmatic theory that can take its place beside syntax, semantics and phonology within an overall theory of grammar.

The need for a pragmatic component in an integrated theory of linguistic ability can be argued for in various ways. One way is to consider the relation of the pragmatics-semantics-syntax trichotomy to the competence-performance dichotomy advanced by Chomsky (see Kempson, 1975: Chapter 9). In Chomsky's view, grammars are models of competence, where competence is knowledge of a language idealized away from (especially) irregularity or error and variation; to this, Katz influentially added idealization away from context (see Lyons, 1977a: 585-91, for discussion of kinds of idealization). On such a view, insofar as pragmatics is concerned with context, it can be claimed that by definition pragmatics is not part of competence, and thus not within the scope of grammatical descriptions. But suppose now we require that adequate grammatical descriptions include specifications of the meaning of every word in a language, and such a requirement has normally been assumed, then we find words whose meaning-specifications can only be given by reference to contexts of usage. For example, the meaning of words like *well*, *oh* and *anyway* in English cannot be explicated simply by statements of context-independent content: rather one has to refer to pragmatic concepts like relevance, implicature, or discourse structure (this claim will be substantiated in the Chapters below). So either grammars (models of competence) must make reference to pragmatic information, or they cannot include full lexical descriptions of a language. But if the lexicon is not complete, then neither is the syntax,

semantics or phonology likely to be. There are other arguments that have been made along the same general lines, to the effect that to capture regular processes (e.g. syntactic regularities) one must refer to pragmatic concepts (see e.g. Ross, 1975), arguments that will arise from time to time in the Chapters below.

Another more powerful kind of argument goes as follows. In order to construct an integrated theory of linguistic competence, it is essential to discover the logical ordering of components or levels. For example, Chomsky has elegantly argued that syntax is logically prior to phonology, in that phonological description requires reference to syntactic categories, but not vice versa; syntax is thus **autonomous** with respect to phonology, and phonology (non-autonomous with respect to syntax) can be envisaged as taking a syntactic input, on the basis of which phonological representations can be built up. Accepting for a moment this kind of argument, the question is, is it possible to argue that there is some accepted component of grammar that is non-autonomous with respect to pragmatics (i.e. some component requiring pragmatic input)? If so, pragmatics must be logically prior to that component, and so must be included in an overall theory of linguistic competence.

It seems fairly clear that it is possible to make this argument in a convincing way. For example, we have already noted the argument (and see Chapter 2 below) that if semantics is to be truth-conditional, then the truth conditions can only be assigned to utterances, not sentences – in other words, contextual specifications are a necessary input to a semantic component, and thus pragmatics is (at least in this respect)¹⁹ prior to semantics. Gazdar (1979a: 164–8) assembles a number of detailed arguments to this effect (and philosophers have long noted further such arguments – see e.g. Donnellan, 1966; Stalnaker, 1972; Kaplan, 1978; etc.). One of these, due to Wilson (1975: 151), will have to suffice here, and holds not just for truth-conditional semantics but for virtually any semantic theory independent of pragmatics. Consider the following sentence:

¹⁹ There are also simple arguments that pragmatics requires semantic input: for example, an ironic interpretation of an utterance can only be calculated if the semantic (or 'literal') reading is already available. So the two kinds of arguments together seem to show that neither semantics nor pragmatics is autonomous with respect to each other – information provided by the one component must be available to the other.

- (21) Getting married and having a child is better than having a child and getting married

Good arguments will be given in Chapter 3 to suggest that the word *and* of itself does not mean (have the semantic content) 'and then', but is neutral with respect to a temporal dimension. So, there is no difference in *semantic* content between *p and q* and *q and p*, or between 'getting married and having a child' and 'having a child and getting married'. How then are we to explain that (22) does not mean the same as (21)?

- (22) Having a child and getting married is better than getting married and having a child

We have to provide a pragmatic account, along the following lines. The 'and then' reading of both *ands* in the first sentence can be shown to be systematically 'read in' to conjoined reports of events by a pragmatic principle governing the reporting of events: tell them in the order in which they will or have occurred. If this is accepted, then the semantic content of (21) (and identically for (22)) would only allow the interpretation that A is better than B (where A is composed of *p and q* or *q and p*, neutral with respect to ordering). But such a reading is either necessarily false or meaningless, and in any case semantically anomalous. The sentence can only be assigned the right truth conditions, or alternatively be given the correct semantic representation, if the pragmatic significance of *and* in this sentential context (namely the 'and then' interpretation) is taken into account before doing the semantics. This amounts to a concise argument that semantics is not autonomous with respect to pragmatics, and that pragmatics provides part of the necessary input to a semantic theory. But if pragmatics is, on occasions, logically prior to semantics, a general linguistic theory simply must incorporate pragmatics as a component or level in the overall integrated theory.

1.3 Current interest in pragmatics

There are a number of convergent reasons for the growth of interest in pragmatics in recent years. Some of these are essentially historical: the interest developed in part as a reaction or antidote to Chomsky's treatment of language as an abstract device, or mental ability, dissociable from the uses, users and functions of language (an

abstraction that Chomsky in part drew from the post-Bloomfieldian structuralism that predominated immediately before transformational generative grammar). In looking for the means to undermine Chomsky's position, *generative semanticists* were then attracted to a considerable body of philosophical thought devoted to showing the importance of the uses of language to an understanding of its nature (work by Austin, Strawson, Grice and Searle in particular). To this day, most of the important concepts in pragmatics are drawn directly from philosophy of language. Once this broader scope for mainstream American linguistics was established,²⁰ pragmatics soon took on a life of its own, for the issues raised are interesting and important in their own right.

But there have also been powerful motivations of a different kind. In the first place, as knowledge of the syntax, phonology and semantics of various languages has increased, it has become clear that there are specific phenomena that can only naturally be described by recourse to contextual concepts. On the one hand, various syntactic rules seem to be properly constrained only if one refers to pragmatic conditions; and similarly for matters of stress and intonation.²¹ It is possible, in response to these apparent counter-examples to a context-independent notion of linguistic competence, simply to retreat: the rules can be left unconstrained and allowed to generate unacceptable sentences, and a performance theory of pragmatics assigned the job of filtering out the acceptable sentences. Such a move is less than entirely satisfactory because the relationship between the theory of competence and the data on which it is based (ultimately intuitions about acceptability) becomes abstract to a point where counter-examples to the theory may be explained away on an *ad hoc* basis, unless a systematic pragmatics has already been developed.

Alternatively, pragmatics and other linguistic components or levels can be allowed to interact. Arguments between these two positions have never been fully articulated, and because of their highly theory-dependent nature are dealt with, in this book, only in passing.

²⁰ It is worth noting that many other schools of linguistic thought had always taken for granted such a broader scope, e.g. the Prague school, the so-called London school, and even the glossesemantics. For a treatment of the historical developments in America, see Newmeyer, 1980.

²¹ A useful general list of pragmatic constraints on linguistic form can be found in Gazdar, 1980a (see also Green, 1978a).

(But see Gordon & Lakoff, 1975; Ross, 1975; Gazdar & Klein, 1977; Lightfoot, 1979: 43-4.)

On the other hand, concurrent developments in semantics have isolated intractable phenomena of a parallel kind: presuppositions, speech acts and other context-dependent implications, together with troublesome phenomena like honorifics and discourse particles that had long been given short shrift in the work of generative grammarians. Further, thought about the nature of the lexicon, and how one might construct a predictive concept of 'possible lexical item', has revealed the importance of pragmatic constraints (see Horn, 1972; McCawley, 1978; Gazdar, 1979a: 68ff). It is these issues, arising from the study of meaning, with which this book is centrally concerned.

In addition to these particular problems that seem to require pragmatic solutions, there are also a number of general motivations for the development of pragmatic theory. One of the most important of these is the possibility that pragmatics can effect a radical simplification of semantics.²² The hope is based on the fact that pragmatic principles of language usage can be shown systematically to 'read in' to utterances more than they conventionally or literally mean. Such regularly superimposed implications can then become quite hard to disentangle from sentence or literal meaning; in order to prise them apart, the theorist has to construct or observe contexts in which the usual pragmatic implications do not hold. For example, it seems perfectly natural to claim that the quantifier *some* in (23) means 'some and not all':

(23) Some ten cent pieces are rejected by this vending machine and that would be the basis of the natural interpretation of a notice with this message, attached to the machine. But suppose I am trying to use the machine, and I try coin after coin unsuccessfully, and I utter (23); I might then very well communicate:

(24) Some, and perhaps all, ten cent pieces are rejected by this vending machine

and indeed I could say this without contradiction. Faced with these facts the semanticist must either hold that *some* is ambiguous between the readings 'some and not all', and 'some and perhaps all', or allow a pragmatic account of the different interpretations. (Parallel

²² Hence the term *radical pragmatics*, as in the title of Cole, 1981, although the term *radical semantics* might be more appropriate.

arguments can be made for the word *all*, and indeed most of the lexical items in a language.) This pragmatic account would explain how principles of language usage allow addressees to 'read in' the 'not all' implication. Since such a pragmatic account is available, as will be seen in Chapter 3, we can let the semantics just provide a reading compatible with 'some and perhaps all'. Not only will such a division of labour approximately halve the size of the lexicon (by accounting for different interpretations of words by a general external principle), it will also immeasurably simplify the logical base of semantics – the word *some* can be equated directly with the existential quantifier in predicate logic (while the reading 'some and not all' taken as basic leads to serious internal contradictions: see Horn, 1973 and Chapter 3 below). In this way, by unburdening semantics of phenomena that are resistant to semantic treatment but tractable to pragmatic explanation, there is considerable hope that pragmatics can simplify semantic theories.

Another powerful and general motivation for the interest in pragmatics is the growing realization that there is a very substantial gap between current linguistic theories of language and accounts of linguistic communication. When linguists talk of the goal of linguistic theory as being the construction of an account of a sound-meaning correspondence for the infinite set of sentences in any language, one might perhaps infer that such a grand theory would *eo ipso* give an account of at least the essentials of how we communicate using language. But if the term *meaning* in this correspondence is restricted to the output of a semantic component, those interested in a theory of linguistic communication are likely to be greatly disappointed. For it is becoming increasingly clear that a semantic theory alone can give us only a proportion, and perhaps only a small if essential proportion, of a general account of language understanding. The substantial gap that remains to be bridged between a semantic theory (together with a syntactic and phonological theory) and a complete theory of linguistic communication will be demonstrated throughout this book. Where are we to account for the hints, implicit purposes, assumptions, social attitudes and so on that are effectively communicated by the use of language, not to mention the figures of speech (e.g. metaphor, irony, rhetorical questions, understatement) that have preoccupied theorists of rhetoric and literature? These communicated inferences can be quite diverse in kind. Consider, for example, the following

extracts from recorded conversations,²³ where the responses to an utterance indicate that for participants the utterance carried the implications (or something like them) indicated in brackets:

- (25) A: I could eat the whole of that cake [implication: 'I compliment you on the cake']
 B: Oh thanks
 (26) A: Do you have coffee to go?²⁴ [implication: 'Sell me coffee to go if you can']
 B: Cream and sugar? ((starts to pour))
 B: Hi John
 (27) A: How're you doing?
 B: Say, what're you doing [implication: 'I've got a suggestion about what we might do together']
 A: Well, we're going out. Why?
 B: Oh I was just going to say come out ...

There are also cases where the location of a verbal exchange in a particular kind of activity seems to warrant specific inferences:

- (28) (*In a classroom*)
 Teacher: Johnnie, how do you spell Ann?
 Johnnie: A, N, N
 ((intervening material))
 Teacher: Okay, Isobel, do you see a name on that page you know?
 Isobel: Ann
 Teacher: That's the one that Johnnie just named [implication: 'That doesn't count']
 (29) (*Beginning of a telephone conversation*)
 Caller: ((rings))
 Receiver: Hello
 Caller: Hello [implication: 'I know who you are, and you can tell from my voice who I am']
 Receiver: Oh hi [implication: 'Yes, I know who you are']

Each of these, or examples like them, will be treated in the pages below, together with more familiar examples of pragmatic implication. The point here, though, is that the existence of a great range of such implications, some of which have only the most tenuous relationship

²³ Drawn, with orthographic simplifications, from the following sources: (25), author's transcript; (26) from Merritt, 1976; (27) from Atkinson & Drew, 1979: 143; (28) from Gumperz & Herasimchuk, 1975: 109ff; (29) from Schegloff, 1979a. Double parentheses enclose descriptions that are not part of the verbal record; further conventions are developed in Chapter 6.

²⁴ American idiom for 'coffee to take away, rather than drink on premises'.

to the semantic content of what is said, emphasizes the need for a theory or theories that will complement semantics in order to give a relatively full account of how we use language to communicate.

Finally, another very important general motivation for the recent interest in pragmatics is the possibility that significant **functional** explanations can be offered for linguistic facts. Most recent linguistic explanations have tended to be internal to linguistic theory: that is to say, some linguistic feature is explained by reference to other linguistic features, or to aspects of the theory itself. But there is another possible kind of explanation, often more powerful, in which some linguistic feature is motivated by principles outside the scope of linguistic theory: for example, it seems possible that the syntactic processes known as *island constraints* (Ross, 1967) can be explained on the grounds of general psychological principles (see e.g. Grosu, 1972). This mode of explanation, by reference to external factors (especially causes and functions), is often called **functionalism** (see e.g. Grossman, San & Vance, 1975). Now the possibility exists that language structure is not independent (contrary to Chomsky's well-known views) of the uses to which it is put. That is to say, it may be possible to give powerful functionalist explanations of linguistic phenomena by reference to pragmatic principles. Indeed, to many thinkers such explanations seem to be obviously of the right kind (cf. Searle, 1974; Givon, 1979a, 1979b). For example, one might observe the fact that nearly all the world's languages have the three basic sentence-types: imperative, interrogative and declarative (Sadock & Zwicky, in press). On the grounds that these seem to be used paradigmatically for ordering, questioning and asserting, respectively, one might argue that it is pointless to search for internal linguistic motivations for these three sentence-types: they recur in the languages of the world because humans are, perhaps, specifically concerned with three functions of language in particular – the organizing of other persons' actions, the eliciting of information, and the conveying of information. (Such an explanation is of course suspiciously *post hoc*: we would need independent evidence that these three activities are indeed predominant in social life.) Or one might note that most languages have some, and many languages have elaborate, ways of encoding relative social status between participants: again, a functional explanation in terms of universal (or near universal) principles of social organization seems called for (see e.g. Brown & Levinson,

1978). Indeed, one might hope for still more in the way of functional explanation: much of the syntactic machinery of a language seems to be concerned with the linear re-organization of material in sentences (as in passive or topicalized constructions), a re-organization which does not seem substantially to affect the (truth-conditional) semantic content. What, then, is the purpose of such elaborate derivational machinery? It may be that it exists essentially for the purpose of meshing sentence-construction with pragmatic principles: for example, for 'foregrounding' and 'backgrounding' informational content²⁵ (or, in the terms preferred in this book, for eliciting certain pragmatic implications).

One of the motivations for research in pragmatics might then be to establish the effects of the uses of language on language structure. But such research requires a fundamental clarification of the *explicans*, i.e. the functional matrix that is to produce explanations of linguistic structure. Unfortunately, many recent examples of such work have utilized explanatory principles that have been left quite vague (see M. Atkinson, 1982). It is important, therefore, that there be sufficiently well-defined pragmatic principles and structures to make such functional explanations precise and testable.

How, therefore, should we think of the uses of language, in a way that could provide functional accounts of linguistic structure? We might turn to traditional approaches to the 'functions of speech' (see the summary in Lyons, 1977a: 50–6). Perhaps the most thoughtful of these is Jakobson's (1960) modification of earlier schemes (see especially Bühler, 1934). He suggests that the functions of speech can be to focus on any of the six basic components of the communicational event: thus the **referential** function focuses on the referential content of the message, the **emotive** function on the speaker's state, the **conative** function on the speaker's wishes that the addressee do or think such-and-such, the **metalinguistic** function on the code being used, the **phatic** function on the channel (establishment and maintenance of contact), and the **poetic** function on the way in which the message is encoded. Any such scheme, though, is of dubious utility to the pragmaticist in search of functional principles: the categories are of vague application, they do not have direct empirical motivation, and there are many other rival schemes built upon

²⁵ See Givon, 1979a; Foley & Van Valin, in press.

slightly different lines. Perhaps the only clear utility is to remind us that, contrary to the preoccupations of many philosophers and a great many semanticists, language is used to convey more than the propositional content of what is said. Certainly, very few linguists have produced analyses of linguistic facts that make use of gross functional categories of this sort (but cf. Halliday, 1973). A very similar sort of enterprise has been engaged in by philosophers interested in the notion of **speech act** (addressed in Chapter 5); either by examining a special set of verbs called **performative verbs**, or by more abstract conceptual analysis, they arrive at classifications of the basic purposes for which language can be used (see e.g. Searle, 1976). Again, such schemes seem to be far too broad to relate to detailed aspects of linguistic structure.

How else, then, might we proceed? One possibility, which has scarcely been explored, would be to take some large sample of the world's languages and ask what basic pragmatic distinctions are needed to describe their grammatical structures. (The procedure requires, of course, acceptance of the view that not all encoded features of meaning are semantic simply by definition.) We would note that many languages have, in addition to the three basic sentence-types mentioned above, others that appear to be similarly circumscribed in use: **exclamatives** that are used paradigmatically to express surprise, **imprecatives** to curse, **optatives** to express a wish, and so on (again, see Sadock & Zwicky, in press).²⁶ Some languages would motivate distinctions that, from the point of view of European languages, are quite exotic. For example, to describe the lexicon, morphology and syntax of Javanese one would need to distinguish three levels of respect to addressees and two levels of respect to referents (Geertz, 1960; Comrie, 1976b); to describe the particles of a number of South American Indian languages one would need to distinguish between sentences that are central versus those that are peripheral to the telling of a story (Longacre, 1976a); to describe the third person pronouns of Tunica one would need to distinguish not only the sex of the referent, but also the sex of the addressee (so there would be two words for 'she' depending on whether one is speaking to a man or a woman; Haas, 1964), while in

²⁶ But caution is necessary here – e.g. what are traditionally called optatives in Sanskrit and Greek do not necessarily, or perhaps even primarily, express wishes.

some Australian languages the pronouns encode the moiety or section (kinship division) of the referent, or the kinship relation between referents (e.g. there are sometimes two words one of which means 'you-dual of the same moiety' and another 'you-dual in different moieties from each other'; Dixon, 1980: 2–3; Heath et al. 1982); to describe the Quileute demonstratives one needs to make a distinction between objects visible and not visible to the speaker (Anderson & Keenan, in press); and so on.

From this profusion of language-specific material one then might be able to build up some idea of just which aspects of the context of utterance are likely in general to exert functional pressures on language. Further, taking features that are directly and simply encoded in one language, one may well be able to find the same features encoded in more subtle and less visible ways in either the structure or the use of other languages. For example, although we do not have in English the grammaticalization of the levels of respect that exist in Javanese, we do have means of expressing degrees of respect, largely by choices in the use of expressions: thus (31) would generally be a more polite request than (30):

(30) I want to see you for a moment

(31) I wondered if I could possibly see you for a moment

So by taking at first just the grammaticalized or encoded features of context in the world's languages, we would have both something like a 'discovery procedure' for relevant functions of language, and a constraint on the relatively vacuous theorizing that often attends speculation about the 'functions of speech'. We can then go on to ask how in other languages without such grammatical means, the same functions are achieved (if indeed they are). Such a way of proceeding has much to recommend it, but scant progress has been made in that direction.

To all such approaches to the uses of speech, a strong objection might be made along the following lines: rather than look for a series of static functions or contextual parameters, one should attend directly to the single most important dynamic context of language use, namely conversation, or face-to-face interaction. The centrality of this functional matrix for language use hardly needs arguing: face-to-face interaction is not only the context for language acquisition, but the only significant kind of language use in many of the world's

communities, and indeed until relatively recently in all of them. Those interested in functional explanations of linguistic phenomena ought then to have a considerable interest in the systematics of face-to-face interaction. The question is how best to approach the study of such interaction. There are perhaps two basic lines of attack: straightforward empirical analysis, and analysis-by-synthesis.

It is the first kind of approach that has so far yielded the most insight, but it is worth considering the possibility of the analysis of interaction by synthesis. Interaction, in the abstract sense intended here, can be understood as the sustained production of chains of mutually-dependent acts, constructed by two or more agents each monitoring and building on the actions of the other (in this sense the mathematical theory of games studies one kind of interaction; see Luce & Raiffa, 1957). Such an approach might begin by adopting Goffman's (1976) distinction between **systems-constraints** and **ritual-constraints**, where the first labels the ingredients essential to sustaining any kind of systematic interweaving of actions by more than one party, and the second those ingredients that, while not essential to the maintaining of interaction, are nevertheless typical of it – they are, if one likes, the social dimensions of interaction. Concentrating on systems-constraints, one may then ask what necessary and jointly sufficient conditions must be met in order for that highly co-ordinated kind of inter-dependent behaviour that we call interaction to 'come off'. Suppose, for example, we had as our task the programming of two robots in such a way that they could systematically aid one another in an open range of tasks: what properties beyond the specific abilities required for the tasks would they need to have? (It may be helpful to think in terms of some specific co-operative task, like the production of a building or a machine.) First, it is clear that they would need to be mutually oriented; they would each need to be aware of what the other was doing at any time. Secondly, they would need to be aware of the interactional domain (e.g. their scope for movement, and the properties of objects around them), and be constantly updating this as it was affected by their actions. Thirdly, they would need, in some sense, to be rational – to have an effective means-ends reasoning that told them how to implement each desired goal. Fourthly, each would need to be able to produce acts conditional on the other producing acts, thus securing the chains of inter-dependent acts typical of interaction. This would seem to require the

ability to reconstruct from each other's behaviour the probable goal that the behaviour was intended to achieve (otherwise, the inter-dependent actions would not be likely to culminate in the achievement of the joint task). Fifthly, there would need to be some specific relation between their overall goals (if interaction is agonistic, or in the terminology of the theory of games, *zero-sum*, then their goals must be inversely related; if interaction is co-operative then there must be some specific shared goals). Sixthly, each robot would have to know that the other had these properties, and know that each knew that, otherwise they could hardly rationally plan actions dependent on the other's plans. It is just possible that these properties would be sufficient, together with the abilities required by specific tasks, to engender a co-ordinated interdigitation of actions that would (remotely) resemble human interaction. The purpose of this thought-experiment is to draw attention to the fact that a number of pragmatic phenomena can be explicated by reference to just these sorts of features: for example, as we shall see, deixis can be thought of as based on the assumption of mutual orientation, presupposition on the assumption of shared knowledge of a domain and its updating, speech acts on the making explicit, for other participants, of one's interactional goals, conversational implicature on the assumption of interactional co-operation, and so on. Thus, if such an approach were developed, one might hope that all the essential concepts for the analysis of pragmatic phenomena would be traceable to the fundamentals of interaction (for an actual computer simulation of conversation along these lines see Power, 1979).

In fact, though, such an approach is still likely to be much too abstract to provide systematic functional accounts of the minutiae of linguistic structures. For a start, it would need to be complemented by the study of *ritual-constraints*, the social and cultural constraints on interaction. Amongst these, there are cross-situational constraints enjoining appropriate social decorum, while there are others appropriate just to specific interactional moments or specific kinds of cultural events. It might be thought that such social constraints would be likely, simply by being social, to be culturally variable, and thus of no great interest to a general (or universal) pragmatic theory. However, this does not seem necessarily to be the case. For example, there are clear pan-cultural principles governing the production of 'polite' or socially appropriate interaction, and these can be shown

to have systematic effects on the linguistic structure of many languages (Brown & Levinson, 1978; Leech, 1980). It is also clear that there are highly specific ritual constraints of a universal, or near universal nature: for example, nearly all cultures seem to have greeting and parting routines (see Ferguson, 1976). More speculatively, it is also likely that in all cultures there are social events demarcated as *formal events* (Irvine, 1979; J. M. Atkinson, 1982), and that some aspects of formality have universal linguistic realizations. Here again there has been very little systematic exploration, although such universal features of the organization of interaction are good candidates for potentially important functional pressures on linguistic structure. Whatever the attractions of universal features of interaction for the explanation of universal pragmatic phenomena, there are also clear language-specific pragmatic phenomena, as in the domain of **social deixis** and elsewhere, where functional accounts of language structure would need to relate these to culture-specific aspects of interaction. Finally, where there are important divisions between kinds of culture and society, one might well expect systematic differences between the associated languages – for example, it is likely that literacy has systematic effects on the lexical, syntactic and semantic structure of languages, even if these have never been spelt out (see Goody, 1977). Here it is evident that an interest in language usage motivated by functionalist approaches to linguistics would take us well beyond the confines of pragmatics (as sketched in the definitions above) into the domain of sociolinguistics and beyond. However, in so far as such social features are part of the meaning of utterances, they ought also to be treated in pragmatics; yet within pragmatics, these social constraints on language usage and their systematic effects on language structure, have been very much understudied, perhaps as a result of the philosophical and linguistic bias (no doubt reflected in this book) towards what Bühler (1934) called the *representational*, and Jakobson (1960) the *referential*, function of language.

The other more promising line of investigation is to explore directly the nature of conversational interaction. The basic concepts of conversation analysis, as employed in a branch of **ethnomethodology**, are the subject of Chapter 6. Here it will suffice to note that this kind of investigation, employing techniques quite alien to the dominant tradition in linguistics, has revealed that conversational interaction has an elaborate and detailed structure of which we

have very little conscious awareness. In this area, at least, the would-be functionalist is offered the kind of rich and intricate structure that may match the detailed organization of linguistic structure, and so can be claimed plausibly to stand in a causal relation to it. For example, the probable universal existence of tag-questions (under a functional definition) can perhaps be related to the universal operation of rules of turn-taking that allow as one option the ending of current speaker's turn by a selection of a next speaker. But, as yet, few linguists have applied the insights from conversation analysis to functionalist studies of linguistic structure.

Finally, there is another kind of empirical approach to the study of interaction, and its effects on linguistic structure, that might be claimed to have a distinct advantage. This is the study of the acquisition of language by children. During the early stages of acquisition, children establish an interactional matrix for language acquisition, and then slowly learn to utilize linguistic means for promoting learning. We are thus enabled to distinguish more easily the interaction. We are thus enabled to distinguish more easily the would-be functionalist's *explicans* (interactional structure) from the *explicandum* (language). A second advantage is that children's 'errors' or incompetences reveal to us what adult competences in verbal interaction must involve. Thirdly, just as cross-linguistic comparisons can reveal general functions of language by contrasts between what is encoded in one language and not in another, so comparisons across stages of acquisition can be revealing in the same way (Ochs, 1979a). There has been much recent work by psychologists and linguists on these early stages of acquisition that has direct relevance to pragmatics, but is not reviewed in this book (see e.g. Ervin-Tripp & Mitchell-Kernan, 1977; Snow & Ferguson, 1977; Ochs & Schieffelin, 1979; and the critical account of such work in M. Atkinson, 1982).

1.4 Computing context: an example

Abstract discussions about the scope of pragmatics like those we have reviewed above, may well leave the reader with little feeling for the nature of pragmatic phenomena. Here an extended example may help to clarify the kinds of facts with which pragmatic theories are concerned.²⁷ Let us take a simple three-sentence exchange

²⁷ The mode of explication, and a number of the points, are derived from Fillmore, 1973.

between two parties, and ask what information it provides us with above and beyond what might be given by the semantic content of the component sentences. More specifically, we can ask what implications are carried by the sentences about the contexts in which they are being used. The example is constructed – the reader is warned because good reasons for preferring naturally occurring conversational data will be given in Chapter 6. Here is the exchange:

- (32) (i) A: So can you please come over here again right now
(ii) B: Well, I have to go to Edinburgh today sir
(iii) A: Hmm. How about this Thursday?

It is not difficult to see that in understanding such an exchange we make a great number of detailed (pragmatic) inferences about the nature of the context in which (32) can be assumed to be taking place.²⁸ For example, we infer the facts in (33):

- (33) 1. It is not the end of the conversation (nor the beginning)
2. A is requesting B to come to A at (or soon after) the time of speaking; B implies he can't (or would rather not) comply; A repeats the request for some other time
3. In requesting, A must (a) want B to come now, (b) think it possible that B can come, (c) think B is not already there, (d) think B was not about to come anyway, (e) expect that B will respond with an acceptance or rejection, and if B accepts, then A will also expect B to come, (f) think that his (A's) asking may be a possible motive for B to come, (g) not be, or be pretending not to be, in a position to order B to come
4. A assumes that B knows where A is; A and B are not in the same place; neither A nor B are in Edinburgh; A thinks B has been to A's place before
5. The day on which the exchange is taking place is not Thursday, nor Wednesday (or, at least, so A believes)
6. A is male (or so B believes); A is acknowledged by B to have a higher social status than B (or to be playing the role of a superior)

²⁸ There may perhaps be some equivocation here between inferences that participants, i.e. A and B, might make, and inferences that observers or analysts – or readers of (32) – might make. For example, since A and B may well *presume* the facts in 4, 5, and 6, we might want to say that they didn't infer them; yet, from the fact that participants would be expected to correct errors in such presumptions, we can conclude that they must nevertheless make the inferences to check that their presumptions hold. Even the inference that the conversation is not about to terminate is, as we shall see

Obvious to the point of tediousness though some of these inferences may be,²⁹ they are not, on a reasonable circumscription of semantic theory, part of the semantic content of the three sentences. Rather, they reflect our ability to compute out of utterances in sequence the contextual assumptions they imply: the facts about the spatial, temporal and social relationships between participants, and their requisite beliefs and intentions in undertaking certain verbal exchanges. But if the inferences are not (or not all) part of the 'literal' meaning or conventional content of what is said, from what sources do they arise? One possibility is that the sentences simply invoke mental associations, in the way that hearing, say, the word *prognosis* might make one think of hospitals. But here that does not seem to be the case. The inferences are systematic, they are decodable by different interpreters in the same way, and without most of them the exchange cannot be understood; most of them must therefore be part of what is communicated, in Grice's strict sense of *meant-nn*. But above all, we can trace each of these inferences to the facts that trigger them, namely, aspects of the form and juxtaposition of the utterances themselves, and we can go on to specify the regular principles that, given such aspects of utterances, produce the inferences in question. The Chapters below will each be concerned with particular principles of this kind, but let us here just identify the aspects of the utterances that trigger each of the inferences.

First, we know ((33) 1) that it is not the end of the conversation because utterance (iii) is not a possible closing utterance: for one thing, it requires a response from B, and for another it is not a token of one of the regular closing forms that persons use in conversation (*Okay*, *see you later* or the like). That is, some turns at talking come in pairs, such that one part of the pair requires the second part in response, while conversations have overall structures with well-bounded beginnings and endings. In short we have strong expectations about the structure of conversation which warrant many different kinds of inference (see Chapter 6). We also incidentally know that it is not the beginning (although that, being known in advance by

in Chapter 5, potentially a constant consideration for participants. So at least most of these inferences are ones that A and B must calculate.

²⁹ Actually, all these inferences need further qualification of a rather tedious sort: e.g. the inference 6 in (33) that A is male, or at least that B believes that A is male, should have the additional qualification 'or at least, B is acting as if he or she thinks A is male', etc.

participants, is not part of what is communicated), because there is no token of a conversational opening (like *hello*), and the particle *so* with which utterance (i) begins has the function of tying the present utterance back to prior utterances.

We know the facts in (33) 2 in a rather more complicated way. Whereas the interrogative form of the first utterance might be claimed to encode a question, that is not all that is intended: it would be strikingly unco-operative if B were to say *yes* (meaning just 'yes I am able to come') and then not go to A. Somehow, the interrogative form can also convey a request, and this interpretation is strongly reinforced here by the presence of the word *please* (see Chapter 5). Much more difficult is to see how B's response in (ii) can be understood as a request refusal, for there is no overt relation at all between its semantic content and that function. The implication relies on some very general expectation of interactional co-operation, which allows one to assume that if one utterance calls for a response (and the request in (i) does so), then one may assume (other things being equal) that a following utterance is a relevant response (see Chapters 3 and 6). Such an assumption is strong enough that when one comes across a response that is apparently irrelevant (as (ii) overtly appears to be), an inference is triggered that would preserve the assumption of relevance. Here, in (ii), the utterance provides the clue: B has to go to Edinburgh; thus if A and B are both far from Edinburgh (and mutually know this), so that it will take the rest of the day to travel and do things there, then B is busy today; so B is indirectly producing a reason why he or she can't easily come to see A, and in so doing can be understood to be refusing A's request. In actual fact there is just one overt trigger for this inference: the particle *well* in English serves to warn the recipient that some inferencing must be done to preserve the assumption of relevance. It can be plausibly claimed that, like *so* and many other words, *well* has no semantic content, only pragmatic specifications for usage. (See Chapter 3; an alternative account of this inference and the role of *well* can be constructed using the notion of **dispreferred** response in Chapter 6.)

In (33) 2 we also have the inference that utterance (iii) counts as a repeat request. To account for this, we would need first of all to explain how the form *how about VERBing* is more or less restricted to usages in suggestions (again, this looks like a linguistic form that has pragmatic rather than semantic content, a problem discussed in Chapter 5). So A is suggesting that someone do something on

Thursday. Again, in order to preserve the assumption of relevance, an inference must be made about who is to do what: since the last mention of someone doing something involved B going to A, that is presumably what A intends, and may thus be taken to have meant. Here we seem to be implicitly relying on a further assumption, namely an assumption of topical coherence: if a second utterance can be interpreted as following on a first utterance, in the sense that they can be 'heard' as being concerned with the same topic, then such an interpretation of the second utterance is warranted unless there are overt indications to the contrary (again, see Chapters 3 and 6). Finally, the particle *hmm* is not dismissable as just a 'performance error' or a 'filled pause'; it has specific interactional functions, best explicated in terms of the system for taking turns at speaking in conversation, where it can be seen to be (amongst other things) a turn-holding device (see Chapter 6).

We come now to the inferences in (33) 3. What are the sources for these? We have already seen that indirectly the question in utterance (i) must be understood as a request. Now it simply follows that, if A is requesting B to come, and A is behaving rationally and sincerely, we may assume all the facts in (a)-(g). Why? Partly because if we explicate the concept of requesting, it will be found to be constituted of the very speaker beliefs and wishes listed in part in (a)-(g) (see Chapter 5). But of course it would be possible to go through the behavioural motions of requesting without having any of the requisite beliefs and intentions. Therefore one is warranted in making the inference from the behaviour to the beliefs and intentions of the speaker only by a general assumption of sincerity, or co-operativeness (see Chapter 3). If A knows in advance that B can't come, then he is being deceptive; but if he knows that B knows that he knows that B can't come, then he cannot be interpreted as requesting at all (utterance (i) might then be a joke, or if B is in an incapacitated (say, inebriated) state, perhaps a gibe).

The inferences in (33) 4 are easier to account for. The word *here* denotes the (pragmatically bounded) place where the speaker (A) is at the time of speaking; if B does not know (or cannot find out) where A is, *here* is uninterpretable in the sense that B cannot comply with the request to go there. So A would be less than fully co-operative or rational if he did not think that B knew (or could find out) where he was. We also know that A and B are not in the same place (or at least are at some distance from one another). We know this because

the word *come* (at least with the tense and aspect in (i)) denotes either motion towards the speaker at the time of speaking (as in *Come to breakfast, Johnny*) or motion towards the addressee's location at the time of speaking (as in *I'm coming, Mummy*). Note that as with *here*, the meaning of *come* can only be explicated by reference to pragmatic or contextual parameters (speakers, addressees, times and places of speaking). In utterance (i), *come* cannot denote movement towards the addressee, because the subject of *come* is *you*, and the addressee can hardly move to where the addressee already is. So it must denote movement towards the speaker; but again, the addressee can hardly move towards the speaker if there is no significant distance between them; therefore A and B are not in the same place. Here we might note that they are also not in Edinburgh: we know this for B because B claims to have to go to Edinburgh, and *go* here means movement away from the place of the speaker at the time of speaking; we know it for A also, because if A is in Edinburgh, then B's having to go to Edinburgh can hardly be an excuse for B not going to A today. We make all these inferences on the basis of the deictic words, *come*, *go* and *here* (not to mention *now*), together with reasonings about the nature of our physical world (see Chapter 2). A natural interpretation (by an observer or analyst) of this deictic set-up is that A and B are talking on the telephone. Finally, we know that (A believes that) B has been to A's present location before because of the word *again*: this can be claimed to be a pragmatic rather than a semantic implication just because, unlike semantic implications, those associated with *again* are not normally negated by the negation of the main verb. We are inclined therefore to say that *again* **presupposes**, rather than semantically entails, that some event referred to happened before as well (see Chapter 4).

The implication (33) 5, that the day of speaking is other than Wednesday or Thursday, is also due to deixis (explained in Chapter 2), for the word *Thursday* in utterance (iii) is used in a deictic way that invokes pragmatic parameters (there are other usages that do not, e.g. *Pay day is Thursday*). Here the modifier *this* picks out a particular Thursday in relation to the speaker's location in the week: *this Thursday* means the Thursday of the week in which the speaker is speaking.³⁰ But on Thursday, the Thursday of this week cannot, by pragmatic convention, be referred to as *this Thursday*; we must

³⁰ Or the Thursday in some week otherwise pragmatically identified, e.g. by gesture at a calendar.

instead say *today*. By the same token, we cannot say *this Thursday* on Wednesday, because we ought to say *tomorrow*. So the exchange in (32) takes place neither on Wednesday nor on Thursday. (There may be some different restrictions on usage here in different varieties of English, and there are also some interesting ambiguities; see Chapter 2 below and Fillmore, 1975.)

Finally, we have the inferences in (33) 6 that A is male, and of apparently higher social status than B. These are based most soundly on the vocative item *sir*, for that is what that word seems to mean. Again, on a truth-conditional theory of semantics, those meanings cannot be captured – we would not want to say that B's assertion in (ii) was false if B had simply misidentified A and assumed mistakenly that A was a male superior (that would make truths relative to whomsoever they are addressed).³¹ Further, in an intuitive way, the meanings of *sir* here are not part of the content of what is asserted; they are background assumptions about the context, specifically the kind of person B is addressing. We may therefore say that *sir* **conventionally implicates** that the addressee is male and socially higher in rank than the speaker (see Chapter 3).

There are no doubt many other pragmatic inferences that can be wrung from an exchange as short and insignificant as this. But these will serve to indicate the general nature of the phenomena that pragmatics is concerned with. The point is that we can compute out sequences of utterances, taken together with background assumptions about language usage, highly detailed inferences about the nature of the assumptions participants are making, and the purposes for which utterances are being used. In order to participate in ordinary language usage, one must be able to make such calculations, both in production and interpretation. This ability is independent of idiosyncratic beliefs, feelings and usages (although it may refer to those shared by participants), and is based for the most part on quite regular and relatively abstract principles. Pragmatics can be taken to be the description of this ability, as it operates both for particular languages and language in general. Such a description must certainly play a role in any general theory of linguistics.

³¹ But some assertions, e.g. those with *you* as argument of a predicate, do indeed have just such a relativity. The point here rests on the fact that the vocative item *sir* is not such an argument (e.g. subject or object of a verb); thus the meaning of *sir* seems not to be part of the proposition expressed by (ii), and thus not part of the truth conditions.