

1

Pragmatics and Linguistic Underdeterminacy

It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a terrestrial being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new.

(Frege 1923/77: 55)

all that is required is that the properties of the ostensive stimulus [utterance] should set the inferential process on the right track; to do this they need not represent or encode the communicator's informative intention in any great detail.

(Sperber and Wilson 1986a/95b: 254)

1.1 Saying and Meaning

It is widely observed that there is often a divergence between what a person says and what she means, between the meaning of the linguistic expression she uses and the meaning she seeks to communicate by using it. Some distinction or other of this sort is made by virtually everyone working in pragmatics and its reality is confirmed by our daily experience as speakers and hearers. I aim to do the following in this chapter: first, to chart the extent of this gap between the meaning of the linguistic forms we use and what we mean on occasions of our use of them; second, to examine why there should be such a gap and whether it is a contingent or necessary property of verbal communication; third, briefly to consider whether there is any analogous discrepancy between thought representations and their content.

Let us start with some phenomena which are obviously part of what is meant by the speaker but not part of what her linguistic string means, and move towards instances where the distinction is not so clear. The textbook case is irony and its standard characterization is that of saying one thing while meaning the opposite.

Though this is certainly an inadequate characterization, it is good enough for the immediate point. So a speaker may utter (1), when what she intends to communicate is that Joan has a very poor sense of direction, that she is bound to get lost and that it is laughable to expect her to arrive on time:

- (1) With her excellent spatial sense, Joan is sure to find a shortcut and be the first to arrive.

Tropes, or figurative uses of language, in general, tend to exemplify clearly the saying/meaning distinction. So metaphor, metonymy and hyperbole, for instance, all involve saying one thing in order to communicate something else. All of these can be, and standardly are, viewed as cases where what is said is not even a part of what is meant, but is merely a vehicle for conveying what is meant.

There is another class of cases, where what is said is included in what is meant, but constitutes only a small part of what is meant and is, at least in some instances, not the main point of the utterance. Similes, understatements, and indirect answers provide such examples, so the speaker of (2) or (3B) means what she says, but she means a great deal more as well:

- (2) Bill behaves like a three-year-old child whose teddy-bear has been taken away.
- (3) A: Did you enjoy the evening at Bob and Sue's?
B: I'm not much of a party person.

The speaker in each case intends her utterance to be taken literally, but she also intends her addressee to draw certain further implications from it: in (2), implications regarding Bill's behaviour and character, and, in (3B), a rather negative answer to A's question and other implications concerning her own preferences and dispositions.

A property that both of these classes of saying/meaning divergences may exhibit is a kind of open-endedness in what is meant, while what is said is usually felt to be determinate and singular. So in (4), a metaphorical case, where the speaker does not mean (at least part of) what she says, what she communicates is an impression of the sort of behaviour, demeanour and psychological state typical of Mary when she is crossed.

- (4) When she doesn't get her own way Mary becomes a raging inferno.

It would be difficult to formulate this in terms of a small definitive set of propositions and there is room for differences across hearers as to the specific implications they entertain as part of their understanding of the utterance. A similar point can be made about the more mundane example in (5), where the speaker does mean what she says, but would also standardly communicate a range of implications about her ability to function today, her readiness to get on with work, her improved state of mind, etc.

- (5) I'm feeling better today.

Utterances which employ a subsentential linguistic expression are another sort of case again. The utterance in (6) employs just a prepositional phrase and the one in (7) just an adjective, but what is meant by a speaker in both cases is something sentence-shaped (propositional), presumably quite obvious in the context.

- (6) On the top shelf.
 (7) Higher.

When (6) is uttered by a speaker who realizes that the hearer, making his breakfast, is looking for the marmalade, it communicates 'the marmalade is on the top shelf'. From there on, the example is just like those of the second set above in that it may well have various further intended implications: the marmalade does not belong on the bottom shelf, I have moved it to its proper place, I am not trying to hide it from you, etc.

What these examples demonstrate is that, in addition to a speaker standardly meaning more or other than she says, the 'what is said' of the utterance may itself involve more than the meaning of the linguistic expressions used. So it looks as if we have to distinguish two notions which, in these preliminary observations, have been run together: there is linguistic meaning, the information encoded in the particular lexical-syntactic form employed, and there is the thought or proposition which it is being used to express, that is, what is said. While there is a fair amount of variation in how the term 'what is said' is construed, it is generally agreed to be something fully propositional, that is, semantically complete, and so truth-evaluable.¹ It is this disparity, between linguistic meaning and the proposition expressed, that I want to concentrate on in what follows in this chapter. That other major symptom of the disparity between linguistic content and what a speaker means, the intended implications (the implicatures or implicit import) of an utterance, will be taken up again in chapter 2.

While subsentential utterances are typical of much ordinary conversation among familiars, most of those linguistic productions that have the status of discourses or texts are supersentential, that is, they generally consist of more than a single sentence. There is a range of relations which may be understood to hold between sequences of sentences uttered in a discourse and these too are frequently not encoded by the linguistic expressions used:

- (8) a. He mistook his wife for a hat-stand; he wasn't wearing his glasses.
 b. Her life was in a mess. Her lover had left her and her electric toothbrush wasn't working.

In an utterance of (8a), the second sentence would be understood as giving an explanation for the state of affairs described in the first. In an utterance of (8b), the second sentence would be understood as elaborating on or exemplifying the statement in the first. An utterance of a single sentence which consists of more than one clause

may also be understood as communicating a stronger relationship between the states of affairs described than is encoded by the element that connects the clauses:

- (9) a. He wasn't wearing his glasses and he mistook his wife for a hat-stand.
 b. When she saw Mrs Simpson coming down the aisle she hid behind the breakfast cereals.

In both of these cases, a cause–consequence relation is understood to hold between the states of affairs described, though neither of the clausal connectives, ‘and’ and ‘when’, nor any other linguistic element in the utterances, encodes this. Whether these communicated relationships are part of what the speaker has said (the proposition she has expressed) or are merely implications of the utterance will be considered in chapter 3.

Before looking more closely at ways in which encoded linguistic meaning falls short of determining the proposition expressed or ‘what is said’, there is another sort of case of the coming apart of speaker meaning and linguistic meaning which should be mentioned, if only to set it aside for the moment. This is the phenomenon of linguistic mistakes and misuses. Speakers believe that the linguistic forms they employ in an utterance have a meaning (or encode some information) and that knowledge of this form–meaning correlation is shared among competent users of the form. They generally intend these meanings encoded by their utterance to be recovered by their addressees and used by them, in conjunction with their pragmatic abilities, in the derivation of the intended content. This holds for all the examples discussed so far, including those where what is said is not part of what is meant. So, although the speaker of (1) did not intend her hearer to take her to be endorsing the view that Joan has excellent powers of spatial orientation, she nevertheless did intend her hearer to access those very concepts (encoded in the linguistic expression she used) in the process of arriving at the intended interpretation. When certain sorts of misuse occur, however, it seems that some specific element of the meaning encoded in the linguistic form the speaker employs falls outside any intention she has in producing the utterance. Consider the case of a speaker who uses the sentence in (10) with the intention of communicating that Mary is a member of the upper class, believing that ‘hoi polloi’ encodes the concept UPPER CLASS:

- (10) Mary is one of the hoi polloi.

The concept COMMON FOLK which is ‘actually’ encoded by ‘hoi polloi’ (in the public language system) falls under no intention the speaker has in uttering (10). She may, nonetheless, succeed in her communicative intention, if one or other of two special conditions pertain. Either the hearer is also ‘mistaken’, and in the same way as the speaker, in his understanding of the expression ‘hoi polloi’ (perhaps through a sound association with ‘hoity toity’); or the hearer, whose lexical form ‘hoi polloi’ maps ‘correctly’ to a conceptual address for [COMMON FOLK], recognizes the disparity between its meaning and the speaker’s intention and, charitably, makes the appropriate adjustment.² There are many types of mistake, each with its own particular

properties, to which these general remarks apply, including malapropisms, so-called Freudian slips, and various articulatory errors, such as spoonerisms, which are temporary malfunctions of the system, brought on by performance factors such as tiredness or emotional strain.

What is the proposition expressed by an utterance of (10) in this situation? What has been said? Some might say that, strictly and literally, it is that Mary is one of the common people, although this is quite different from what is meant, and a rather poor vehicle for (non-ironically) communicating what is meant. Others might say that the proposition expressed is that Mary is a member of the upper class, just as the speaker intended, although the concept decoded from the form ‘hoi polloi’ is quite different. Clearly, much depends on our conception of ‘what is said’ or ‘the proposition expressed’ by an utterance, how close it is taken to be to linguistic meaning, the extent to which, if at all, speaker intentions play a role in its determination. These issues are addressed in the next chapter.

One might reasonably feel that the very fact that these are errors, that what is encoded in a case like (10) (and, on some construals, therefore, ‘what is said’) falls right outside the speaker’s intentions, makes them special and marginal. Certainly, they bring an unclarity into the concept of what is said, since up to now we have been assuming that what a speaker says by an utterance is not at odds with what the words she uses mean, even if she doesn’t in fact intend to be taken as meaning what they mean. However, while such cases can be safely ignored by a semantic theory (concerned with explicating linguistic meaning), an adequate pragmatic theory, whose mission is to explain how utterances are interpreted, does have to attend to them; in particular, it has to account for how such encoding disparities between speakers and addressees can, sometimes at least, be cases of successful communication.

1.2 The Underdeterminacy Thesis

From the discussion above, three levels of utterance meaning have emerged, which, although quite distinct, remain in need of considerable clarification: linguistic meaning, what is said and what is meant. I started out by treating the first two as if they were the same, distinguishing them from what is meant, but it soon became clear that what is said has to be distinguished from linguistic meaning. As a result, we have three possible underdeterminacy theses:

- (a) Linguistic meaning underdetermines what is meant.
- (b) What is said underdetermines what is meant.
- (c) Linguistic meaning underdetermines what is said.

I do not think that anyone, apart, perhaps, from a rabid ‘language is all’ social semiotician, would dispute the first two. I want to examine the third one, which I will call the linguistic underdeterminacy thesis or the semantic underdeterminacy thesis or just *the* underdeterminacy thesis. What is meant by this is that the linguistic semantics of the utterance, that is, the meaning encoded in the linguistic expressions

used, the relatively stable meanings in a linguistic system, meanings which are widely shared across a community of users of the system, underdetermines the proposition expressed (what is said). The hearer has to undertake processes of pragmatic inference in order to work out not only what the speaker is implicating but also what proposition she is directly expressing. My purpose here is twofold: (a) to demonstrate the vast extent of this phenomenon, and (b) to prepare the way for an investigation in the next chapter of the various notions of explicitness found in the semantic and pragmatic literature, including 'saying' and 'what is said', 'making as if to say', 'proposition expressed', 'propositional form of the utterance', 'truth-conditional content', 'explicature' and, unlikely though it may sound in a discussion of explicitness, 'implicature'. All of these lie on one side of a divide, on the other side of which is 'implicature', the standard term for the implicit content of an utterance.

Before looking at some of the sources of linguistic underdeterminacy, I'll make a brief terminological digression. Perhaps it is sufficiently clear from what has been said so far that the two terms '*underdeterminacy*' and '*indeterminacy*', are not synonymous, but let's try to be explicit about this. '*Indeterminacy*' seems to be used with reference to several different phenomena. First, it is sometimes used in a contrast with ambiguity. Linguists tend to reserve the term 'ambiguity' for those random and arbitrary coincidences of bits of linguistic form which encode two or more distinct concepts, such as 'bank' and 'visiting relatives'. *Indeterminacy*, then, is used of some of the other sources of the linguistic underdeterminacy of propositional form, so we see '*referential indeterminacy*' used of indexicals and definite descriptions which require contextual considerations for the determination of their reference; regarded as formal types within a linguistic system, they have no determinate reference. The term could be similarly used in '*predicational indeterminacy*' or '*conceptual indeterminacy*', though here we more often find the expressions 'vagueness' and/or 'generality of sense', a matter of practice rather than principle.

Then, stepping outside the linguistic underdeterminacy thesis and considering the implicit content of utterances, there is the *indeterminacy* of implicature mentioned by Grice (1975/89b: 40) and given theoretical flesh by Sperber and Wilson's (1986a/95b: 195–200) concept of weak implicature, to be discussed in chapter 2. Examples (4) and (5) above are cases where the particular implicatures derived by an addressee may not have been specifically intended by the speaker; that is, there is *indeterminacy* regarding which implications within a range of possibilities fall within the speaker's informative intention. Finally, of course, there is the much touted '*indeterminacy of translation/interpretation*' thesis of Quine: according to this, there just is no fact of the matter concerning which of several hypotheses about the meaning of a linguistic expression or its translation into another language is correct; all of them may be compatible with the available evidence (the evidence allowed by Quine being restricted to observable features of the behaviour of the users of the linguistic expression). The common feature of these various uses of the term '*indeterminacy*' is, I think, captured by the phrase 'no fact of the matter'; no conclusion can be drawn because there is none to be drawn. Linguistic '*underdeterminacy*', by comparison, does not entail that there is no fact of the matter as regards the proposition expressed, but rather that it cannot be determined by

linguistic meaning alone. It may be that the proposition expressed by an utterance can also exhibit the property of ‘indeterminacy’ (a possibility considered in chapter 5), but that is a quite separate matter from the current focus on its linguistic underdeterminacy. That’s the end of this terminological digression.

1.2.1 Sources of linguistic underdeterminacy

Now to some of the ways in which a linguistic expression may underdetermine the proposition expressed, or, in other words, ways in which content is context-sensitive. First, there are linguistic ambiguities to be resolved and indexical expressions whose referents must be assigned. These two pragmatic processes are widely acknowledged, even by those who want to keep the gap between linguistic meaning and ‘what is said’ to a minimum; for instance, they are the two processes singled out by Grice (1975/89b: 25)³ as necessary additions to conventional content in identifying what is said. Those semanticists who aim to give natural language sentences a truth-conditional (hence propositional) semantics are, of course, not concerned with how ambiguities or referential indeterminacies are resolved but, nevertheless, have to accommodate both indexicality and ambiguity in their accounts. One semantic approach to indexicality is demonstrated in the two (roughly equivalent) versions of a truth-statement for the sentence ‘this is green’ in (11), where the statement quantifies over *utterances* of an indexical *sentence*, thereby abstracting away from particular contexts and so particular referential resolutions:

- (11) a. An utterance of ‘this is green’ is true just in case the entity that the speaker refers to with ‘this’ is green.
 b. $(u)(x)$ [If u is an utterance of ‘this is green’ and ‘this’ refers to x , then u is true just in case x is green]

This truth-conditional treatment of indexicality is taken up again briefly in section 1.5, where I consider the general feasibility and appropriateness of a truth-conditional approach to the semantics of natural language expressions.

The way ambiguity (lexical and syntactic) is reflected in truth-conditional theories highlights the difference between this sort of semantic theory and the cognitive processing account of utterance understanding that I am working towards. A semantics for an n -ways ambiguous natural language string is complete once it has provided n different T(ruth)-sentences in the metalanguage, one for each sense of the natural language string.⁴ This is obviously not a trivial undertaking, but the point is that the n different sentences are distinguished in advance of their treatment by the truth theory. What the pragmatic theory must confront is the very different issue of how the hearer recognizes (or ‘alights on’) the one (or, on the occasion of a pun, two) of these n possibilities the speaker intends on a particular occasion of use.

Although Grice acknowledges that reference assignment and disambiguation are necessary for a full identification of what the speaker has said, he does not say anything about how these processes are achieved. It seems reasonable to surmise from

the omission of any reference to conversational maxims at this point (in a lecture/article which is primarily focused on these maxims and the work they do in communication; Grice 1975/89b: 22–40) that he did not think they played a role in disambiguation and reference assignment, a point to be considered more thoroughly in chapter 2. On the more cognitively oriented approach of relevance theory, the communicative principle responsible for deriving conversational implicatures is also instrumental in identifying the intended sense of an ambiguous linguistic form and the intended referent of an indexical.

The third way in which linguistic content underdetermines what is said arises when the expression employed does not determine a full proposition even after all necessary reference assignments and disambiguations have taken place. Phrasal and lexical utterances, such as those in (6) and (7), are the obvious cases here. However, there are also fully sentential utterances whose encoded meaning does not seem to determine a fully propositional representation, that is, one which, in principle at least, could be assigned a truth value:

- (12) a. Paracetamol is better. [than what?]
 b. It's the same. [as what?]
 c. She's leaving. [from where?]
 d. He is too young. [for what?]
 e. It is raining. [where?]

As the bracketed questions indicate, these examples require completion before they can be judged as true or false of a state of affairs. What they determine (given reference assignment, etc.) has been described as a subpropositional logical form (Sperber and Wilson 1986a/95b: 188), or a propositional radical, or a fragment of a proposition (Bach 1994a: 269). The missing constituent which will bring them up to full propositionhood has to be supplied pragmatically. These sorts of cases were not described by Grice so we cannot know for sure what he would have said about them. Since he seems to have conceived of 'what is said' as fully propositional (the truth-conditional content of an utterance) he might have agreed that a completion process was necessary and that the missing material would be readily contextually determined, again most likely without any role for conversational maxims (but see discussion in section 2.2.2).

It is appropriate at this point to mention a principle that has been held fairly widely by philosophers, but which is questionable in the light of the previous considerations. This is the **Isomorphism Principle**. As Frege puts it (in the continuation of the quotation at the beginning of this chapter): 'we [are] able to distinguish parts in the thought corresponding to the parts of a sentence, so that the structure of the sentence serves as the image of the structure of the thought' (Frege 1923/77: 55). A more recent formulation, from Fodor and Lepore (1991), is: 'If a sentence S expresses the proposition P, then syntactic constituents of S express the constituents of P.' They describe this isomorphism as a perfectly universal feature of natural language, but acknowledge (in a footnote) that there is an issue to be addressed: 'Suppose, for example, that you hold that (in a null discourse) the sentence "it's raining" expresses the proposition that it's raining in the context of utterance. Then

either you must say that “it’s raining” has more constituents than appear on its surface or that the isomorphism principle can be violated by pragmatically carried information’ (Fodor and Lepore 1991: 333, note 2; and see Fodor 2001, for a fuller recognition of violations of the principle).

For those who want to preserve the principle, there are, in fact, two quite different directions that could be taken. The first is, as Fodor and Lepore note, to insist that, contrary to appearances, there are hidden unarticulated linguistic constituents, such as a covert location indexical in the sentence ‘it’s raining’, which, like overt indexicals, call for the contextual provision of a value. This position, currently taken by Stanley (2000) among others, is discussed in section 2.7 in the next chapter. The second approach, taken by Bach (1994a, 1994b), generally eschews hidden linguistic elements but involves such a strict and delimited sense of ‘what is said’ that the Isomorphism Principle is observed. For instance, what is said by a speaker who utters ‘it’s raining’ does not include anything concerning a location of the instance of raining. It follows from this stance that what is said in cases such as those in (12) is subpropositional and so cannot be evaluated for truth. The third approach to these examples, which I support, allows that pragmatic processes can supply constituents to what is said solely on communicative grounds, without any linguistic pointer, in which case the Isomorphism Principle does not hold. This is demonstrated in the next chapter (section 2.3.4).

Consider now some cases which are fully sentential and which, given reference assignment, seem to be fully propositional and so not to require any further contextual supplementation in identifying the proposition expressed:

- (13) a. Bob is well groomed.
 b. This fruit is green.
 c. That is difficult.
 d. It is serviceable.

Surely, one might think, Bob is either well groomed or he is not; the particular fruit in question is green or it is not; that is, these are truth-evaluable as they stand, once we know the referents of the subject terms. In fact, as pointed out by Gross (1998, chapter 1), from whom the examples are taken, the adjectival predicates here exhibit four different kinds of context-sensitivity, hence four further ways in which linguistic meaning may underdetermine the proposition a speaker expresses. In (13a), the adjective is ‘scalar’, by which he means it allows for comparison between things with respect to the degree with which they have the property concerned. For instance, Bob may be well groomed for your average graduate student, but not for a candidate for a job in a city bank; he may be well groomed for him (that is, compared with his usual appearance) but not for the sort of man Mary likes to be seen with. The colour predicate in (13b) is what Gross calls ‘part-dependent’, since on different occasions of use it may apply to different parts or aspects of the thing it is being predicated of. For instance, (13b) may be judged true in a particular context provided its peel is green even though its interior is white and its stem is brown, while in a different circumstance (say, fruits are being separated into the ripe and the unripe), the proposition expressed will be evaluated according to whether or

not the interior is green, the colour of the skin being irrelevant. The adjective ‘difficult’ in (13c) demonstrates the context-sensitive property of ‘relativity’, that is, its applicability is relative to something else; for instance, a problem might be difficult relative to my abilities but not relative to yours. Finally, there is the property of ‘vagueness’, where there is no clear boundary between things of which the adjective is true and things of which it is false, and the standards of precision may vary across contexts, so that in one context an object has the property, in another it doesn’t, and in yet another it is borderline. Probably all the examples in (13) have this property. Gross gives (13d) as a case where all four kinds of context-sensitivity are present, so that the proposition expressed by a particular utterance of it could look something like the following:

- (14) This program is serviceable, according to such and such a standard, compared to other programs for such and such a task, for beginners in computing.

These sources of context-sensitivity are not peculiar to adjectives but carry over to many other linguistic elements too (verbs, adverbs, etc.). Again, someone wanting to maintain the Isomorphism Principle has the options given above: (a) to posit hidden linguistic constituents corresponding to each of these pragmatically inferred elements, or (b) to deny that they are part of what is said (in which case ‘what is said’ is not truth-evaluable and the inferred elements arise at some other representational level). According to the ‘free pragmatic enrichment’ approach, which I will pursue in the next chapter, these elements do contribute to the proposition expressed (what is said) and their recovery is not only effected pragmatically but is also motivated pragmatically (rather than linguistically).

Most theorists, though not all (for instance, Bach), would agree that the processes discussed so far are necessary supplements to the linguistically encoded information for arriving at what the speaker has said (the proposition expressed). However, there are, at least, three more groups of cases about which there is little consensus. First, there are examples which raise some tricky questions about the semantics of particular elements of the language. Let’s briefly take the case of negation.

- (15) a. Everyone isn’t hungry.
 b. She didn’t butter the toast in the bathroom with a knife.
 c. The local witch didn’t put a spell on us.
 d. Bill didn’t eat some of the cakes; he ate all of them.

In (15a), there are two possible interpretations: ‘not everyone is hungry’ and ‘no one is hungry’, which are truth-conditionally distinct. This is usually described as a scope ambiguity: either the negation takes scope over the universal quantifier or vice versa. This may be taken as a linguistic ambiguity such that the grammar gives the sentence two logical forms and a truth-conditional semantics for the sentence would assign it two T-sentences. Or it may be that the linguistic system gives the negation operator wide scope over the whole of the rest of the

sentence and that, on occasion, a pragmatic process of logical strengthening eventuates in the stronger interpretation 'no one is hungry'; then there would be but one T-sentence specifying the truth-conditions of the *sentence*. Or the linguistic system may dictate nothing at all about the relation between the quantifier and the negation, so it is left to pragmatics to fix that relation; on such a conception it seems unlikely that any truth-conditional specification could be given for the sentence (an issue to be taken up below). At least these three positions have been supported at different times. On the first position, this example would simply present the pragmatic system with another ambiguity to resolve, a choice between two possible logical forms. On the third one, a pragmatic process of scope fixing would be obligatory; that is, there would be no fully propositional representation until that process took place. The second position gives rise to an interesting situation: once the intended domain of the quantifier has been inferred and reference fixed, we seem to have a fully propositional representation, say 'not everyone at the party is hungry', but in some contexts this would not be the proposition the speaker intended to express; it would be weaker than the truth-conditional content she intended the hearer to understand, namely 'no one at the party is hungry'. This is just one of a set of contentious cases where the proposition derived through the essential processes required to complete the encoded logical form is, arguably, not the proposition the speaker expresses.

Other ambiguities have been claimed for negation. Example (15b) has six or seven interpretations depending on which constituent the negation is taken to apply to (e.g. 'in the bathroom', 'the toast', etc.). In spoken utterances, these would typically be distinguished by the pattern of accentuation, but in written form, the ambiguity has to be resolved entirely pragmatically. Example (15c) has been taken by some to be ambiguous between an understanding on which the negation operator is presupposition-preserving (that is, the entailment of the corresponding positive sentence, 'there is a local witch', is maintained in the negative sentence) and an understanding (less immediately obvious) on which negation is presupposition-cancelling, since this sentence has a reading on which it is compatible with a following utterance denying that there is a local witch: 'the local witch didn't put a spell on us; there is no witch around here.' Finally, the negation in (15d) has been supposed by some to express a rejection of a previous utterance of the corresponding positive sentence, rather than to function in the logical truth-value-reversing way that it does in (15a): the use of the quantifier 'some' is rejected, not because it gives rise to falsehood, but because what it expresses is too weak, as the follow-up clause makes explicit. All, none or some of these various different interpretations of utterances of negative sentences might be a function of the language system itself, that is, different meanings encoded by the word 'not'. Those that are not encoded have to be accounted for pragmatically and would appear to be further strong candidates for pragmatic contributions to the proposition expressed ('what is said') by an utterance. The wide range of semantic and pragmatic analyses of negation that have been entertained are surveyed in chapter 4, and a particular account, within the precepts of a relevance-based cognitive view of pragmatics, is proposed and defended.

Some of the same issues arise for the analysis of ‘and’-conjunctions where it seems that a variety of relations between the conjuncts may be understood (including the cause–consequence connection in (9a) above). There are various rich semantic accounts which might be proposed to explain this, and there are more minimalist semantic accounts which leave it to pragmatics to supply stronger connections, thereby raising the question of whether these connections are aspects of the proposition expressed or distinct implicated assumptions. Other cases are scalar terms, which can have at least two different interpretations in a context – for instance, ‘some’ may be understood as ‘some and possibly all’ or ‘some but not all’ – and descriptions (definite and indefinite), which may have a range of interpretations, including the famous attributive or referential understandings. These cases (negation, conjunction, scalars, descriptions) are quite different from that of, say, pronouns, which patently do not encode either their referents or uniquely identifying descriptions of their referents, so that the role of pragmatics in determining their referents is indisputable.⁵ It is not obvious with the phenomena just surveyed whether they do or do not encode the interpretations they may have in different contexts; they may encode several senses, or a single strong sense, or a single weak sense. In such a situation, semantic and pragmatic analyses have to be developed together.

There is a second set of cases for which a pragmatic process is required to arrive at the proposition intended by the speaker, even though the representation recovered without this process is fully propositional and could, therefore, be argued to constitute what is said by the utterance.

- (16) a. Mending this fault will take time.
 b. The north island is some distance from the south island.
 c. Something has happened.
 d. I haven’t eaten lunch.
 e. I haven’t eaten frogs’ legs.
 f. There’s nothing on telly tonight.

Given reference fixing, each of (16a)–(16c) expresses a trivial obvious truth: any activity takes place over a period of time; there is some distance or other between any two islands; at any moment in time something or other has happened. The point is, of course, that these dull truisms are virtually never what a speaker has intended to express; there is hardly any context in which they will be relevant. So some pragmatic process of enriching or adding conceptual material is necessary in order to arrive at what the speaker intended to express: perhaps, ‘mending this fault will take a longer period of time than such fault-mendings standardly take’, ‘the north island is further from the south island than you think’, ‘something bad has happened on the day of utterance [to *x*]’. It’s worth noting the negative flavour of these enrichments; the relevance of these utterances lies in their alerting the hearer to a state of affairs that runs against his prevailing hopes or expectations.

Some sort of temporal span has to be assigned to (16d) and (16e) and the point of interest here is the difference that the object (‘lunch’, ‘frogs’ legs’) makes in each case to the understanding of the identical verbs, both with the perfect

aspect (have + en): in (16d), the most likely interpretation, across a wide range of contexts of utterances, is that the time-span to which the not-having-eaten-lunch applies is the day of utterance, while in (16e) it is probably the speaker's lifetime. Arguably, (16f) expresses a proposition (given reference fixing for 'tonight') and one which would be standardly false (there's always something on telly, however dire) and obviously so to the interlocutors. In order to arrive at the proposition intended by the speaker, the domain over which 'nothing' operates has to be narrowed down to something like 'programmes worth watching', and then it may well be true. If these assessments of the proposition expressed ('what is said') by the speaker of these examples are correct, then we have another group of cases where pragmatic inference must augment linguistic encoding, even though it is not strictly necessary for the derivation of a fully propositional form. However, for the examples involving a quantifier like 'something' or 'nothing', it might be claimed that there is an implicit (hidden) variable in the linguistic form marking the requirement that a quantifier domain be contextually supplied. If that is so, these examples do not express a proposition until that variable is contextually filled, so the pragmatic process does not, after all, take us from one proposition to another. As with all such claims, this one is contentious: Stanley and Szabo (2000a) favour it; Bach (2000a) and Neale (2000) oppose it (see brief discussion in the next chapter, section 2.7).

The third set of cases also demonstrates the enrichment of one proposition to give another, but differs from the previous sets in that the pragmatic process required seems to involve, not the adding of conceptual constituents, but rather adjustments to linguistically encoded concepts. Here are some possible cases of this:

- (17) a. I'm tired.
 b. Ann wants to meet a bachelor.
 c. The path is uneven.
- (18) a. Her face is oblong.
 b. The steak is raw.
 c. The room was silent.

The idea here is that, in certain contexts, utterances of the examples in (17) involve narrowings or strengthenings of the concepts encoded by 'tired', 'bachelor' and 'uneven'. For instance, in an utterance of (17a), the relevant degree of tiredness might vary from a mild form to a much stronger condition which prevents the speaker from doing a range of mundane household tasks; in an utterance of (17b), the sort of 'bachelor' Ann is understood as wanting to meet may belong just to a particular subset of the set of bachelors, the subset of those who are heterosexual, youngish, interested in marriage, etc. The examples in (18) are intended potentially to involve an opposite process of loosening or widening of a lexically encoded concept. For instance, an utterance of (18a) requires a relaxing of the concept 'oblong' since her face is not likely to be literally oblong; in (18b) the concept 'raw' (encoding 'not cooked') may be adjusted so as to be applicable to foods that have had some, but grossly insufficient, cooking. An interesting possibility to consider

with regard to these examples is that the lexically encoded concept in the logical form of the utterance is replaced by an *ad hoc* concept, pragmatically derived from the lexical one, and that this new non-lexicalized concept is a constituent of the proposition expressed by the speaker of the utterance. This is more controversial in the case of the examples in (18) than in those in (17), since there it effectively involves the loss of some linguistic content; for instance, the *ad hoc* concept that replaces the concept encoded by 'raw' does not analytically imply [uncookedness]. I will argue for this position in chapter 5.

These last three sets of cases all involve a pragmatic process whose result is not necessary in order to secure full propositionality, but seems to be required if we are interested in finding that proposition which it is rational to assume the speaker intended to express. A natural-language semanticist interested in giving a truth-conditional specification of these examples could feel quite justified in ignoring these pragmatic adjustments in a way that he cannot with ambiguity, indexicality and other features of a sentence that leave it subpropositional (semantically incomplete). These then are the most interesting and contentious cases when it comes to giving an account of 'what is said' (the proposition explicitly expressed) by a speaker, especially if this is equated with the truth-conditional content of the utterance, as it standardly is. Some would opt to rule these pragmatic developments out of any concept of 'what is said' and treat them at some other level of utterance understanding; others would prefer a concept of 'what is said' which incorporates the result of all these processes and so would have to argue that the concept of the truth-conditional content of an *utterance* is quite distinct from the concept of the minimal truth-conditional content which is to be assigned to a natural-language *sentence*.

It can be seen that considerable work in clarifying the concept of 'what is said', or the proposition expressed by an utterance, remains to be done; this is tackled in the next chapter. All I have tried to do in this section is to give examples of the range of ways in which encoded linguistic meaning may underdetermine the proposition a speaker expresses by her utterance of a particular linguistic string. These can be summarized in the following short taxonomy:

- 1 multiple encodings (i.e. ambiguities)
- 2 indexical references
- 3 missing constituents
- 4 unspecified scope of elements
- 5 underspecificity or weakness of encoded conceptual content
- 6 overspecificity or narrowness of encoded conceptual content

1.2.2 Underdeterminacy: essential or merely convenient?

The question which arises now and which the following sections will venture towards answering is: is the linguistic underdeterminacy of the proposition (or thought) expressed a necessary or contingent matter? It is presumably not *logically* necessary, since there seems to be no reason to suppose that there simply could not be a language system of some sort capable of fully encoding propositions (or Fregean

thoughts) including all those that a communicator could want to express. The question must be: does linguistic underdeterminacy follow inevitably from the sort of linguistic systems that human mind/brains naturally develop, or is it a feature of utterances which comes from some other source, say, a convention of linguistic usage or the outcome of some natural drive towards communicative efficiency? There are at least the following views on this question:

- 1 *The ‘convenient abbreviation’ view*: while the linguistic expression employed in an utterance does, more often than not, underdetermine the proposition or thought expressed, this is merely a matter of effort-saving convenience for speakers and another sentence which fully encodes the proposition/thought *could* always be supplied.
- 2 *The essentialist view*: underdeterminacy is an essential feature of the relation between linguistic expressions and the propositions (thoughts) they are used to express; generally, for any given proposition/thought, there is no sentence which fully encodes it. There are weaker and stronger versions of essentialism:
 - (a) Underdeterminacy is widespread, but there are some (few) sentences which do fully encode the propositions they are used to express.
 - (b) Underdeterminacy is universal and no sentence ever fully encodes the thought or proposition it is used to express.

On the ‘convenience’ view, for every underdetermining sentence (or subsentential expression), there is another sentence provided by the language system which does fully encode the proposition which the incomplete one, uttered in a particular context, was used to express. These proposition-determining, context-insensitive sentences are called ‘eternal’ sentences, a term which is fairly transparent, but, anyway, will be explained in the next section. So for any (non-eternal) linguistic string which expresses a proposition when uttered in a given context, there is an eternal (context-free) counterpart which expresses that proposition. For instance, for each of the (a) members of the examples in (19)–(21), which underdetermine the proposition expressed in one or more of the ways described in the previous section, there is a fully encoding counterpart. A possible candidate for this in each case is given in the corresponding (b) examples:

- (19) a. He went to the bank.
b. Simon Lewis went to a financial institution situated at 32 Tottenham Court Road in London between 2.00 and 2.30 on 18 May 1999.
- (20) a. It’s the same.
b. Ibuprofen is the same in chemical composition as Nurofen.
- (21) a. On the top shelf.
b. The thick-cut orange and ginger marmalade is on the top shelf of the cupboard facing the door in the kitchen of the attic flat at 57 Sunnyside Road, London N19.

Putative eternal sentences are usually longer and more complex than the underdetermining (non-eternal) linguistic expressions standardly used by speakers, as the (b) examples show; one might think here also of the terrible convolutions of many legal documents where the aim is full encoding. So speakers standardly choose to save themselves the mental effort of formulating (and the physical effort of articulating) eternal sentences, which fully encode the propositions they want to express. They know they can rely on the hearer's inferential powers to map the non-eternal sentence or phrase they have uttered on to a mental representation of the intended proposition, or on to the eternal sentence in the language which maps on to that mental representation (if there is a distinction between the two on this view of language and thought).

On the essentialist view, however, while the (b) examples come closer to encoding the proposition expressed by the speaker than the corresponding (a) examples, they are still underdetermining and, no matter how hard one tries to be fully explicit, by elaborating descriptions so that they may pick out unique entities and properties, one is doomed to failure, at least in the vast majority of cases, according to the weak essentialist view, and across the board, according to the strong view. On the weaker view, the following sentences, which are not context-sensitive in any obvious way, might be eternal sentences, but they are the exception rather than the rule in the linguistic system:

- (22) a. Lions are animals.
 b. The earth goes round the sun.
 c. Two plus two is four.

I shall argue against the 'underdeterminacy as convenience' view and in favour of the essentialist views, inclining toward the stronger of the two, despite the more immediate plausibility, perhaps, of the weaker one. I think that public-language systems are intrinsically underdetermining of complete (semantically evaluable) thoughts because they evolved on the back, as it were, of an already well-developed cognitive capacity for forming hypotheses about the thoughts and intentions of others on the basis of their behaviour. Formulating natural-language sentences of a progressively more determining sort may approach ever closer to a full encoding of propositions expressed, but the progression is asymptotic. Before trying to make this case, I'll consider the concept of eternal sentences and the 'convenience' view of non-eternal sentences in a little more detail.

1.3 Eternal Sentences and Effability

In the previous section, I implicitly equated thoughts and propositions, which seemed harmless on a Fregean notion of thought (as the sense of a proposition). However, there is a more fine-grained psychological notion of thought, on which thought and proposition may come apart. The propositional content of two thoughts, 'the meeting begins now' (occurring at 3 p.m.) and 'the meeting begins at

3 p.m.’, might be identical, but the former may move me to act while the latter does not (if I do not know that it is now 3 p.m.). Similarly, an utterance by X addressed to Y of ‘I would like to shake your hand’ might lead Y to extend his hand to the person speaking to him, while an utterance by the same speaker to the same addressee of ‘X would like to shake your hand’ might not (if Y doesn’t realize that the speaker is X), even though both express the same proposition. The thought expressed by a speaker of the indexical sentence and grasped by the addressee appears to have some property that distinguishes it from the proposition expressed (say, ‘X would like to shake Y’s hand at 2.30 a.m. on 27 June 2001’). (For extensive discussion, see Perry (1979, 1997), from whom these examples are borrowed.) In the next sections, talk of ‘thoughts’ will generally be of the individualist psychological sort (which, of course, includes a wide range for which thought expressed and proposition expressed are the same).

For the moment I am assuming that thoughts, whether expressed by utterances or not, are semantically complete, that is, that they are truth-evaluable in and of themselves, without any need of contextual completion or specification. In other words, whether or not there is a sentence in the public language that fully encodes them, they themselves are ‘eternal’: if a thought is true/false at this moment, it has always been true/false and always will be true/false. This assumption will be questioned in section 1.7, where I briefly consider the possibility of the underdeterminacy of thought.

1.3.1 Eternal sentences and Platonism

Now let us consider the concept of an ‘eternal sentence’, where by ‘sentence’ I mean a natural-language sentence, one of those syntactically complete entities which has phonological, syntactic and semantic properties, and which can be used by human beings to make their thoughts known to others. If we take thoughts to be semantically complete, then eternal sentences are ‘complete formulations’ of thoughts, as Wettstein (1979: 92) puts it; the truth value of an eternal sentence stays fixed through time and from speaker to speaker (see Hookway 1997). On the ‘convenience’ view of linguistic underdeterminacy, when a non-eternal linguistic expression is uttered (i.e. an incomplete formulation), what is expressed thereby can, in every case, be completely formulated by some eternal sentence. This position was held by Frege (1918a), and was more recently and fully presented by Quine (1960: 193–4). Katz (1972, 1977, 1978, 1981) endorses Quine’s view:

Quine’s idea is that a [non-eternal sentence] . . . can be expanded on the basis of the information in the context to provide another sentence that expresses a proposition that always makes the statement in question, no matter what the context of utterance. The expansion consists of replacing each indexical element by an expression that has the same reference as the indexical element it replaces but whose referent stays fixed with variations in time, place, speaker, etc. The usual indexical tense indicator will be replaced by such a referentially unique time designation, devised with respect to some

appropriate calendar and clock; indexical nominal elements like 'I', 'he', 'it' and 'John' will be replaced by precise specifications of the individuals or objects that include whatever information about their vital statistics is required to make the specifications resist changes in reference. (Katz 1972: 126)

In other words, the infinite set of sentences that a linguistic system generates can be partitioned into two infinite subsets, one consisting of the underdetermining non-eternal sentences, which speakers find a very convenient effort-saving means for communicating their thoughts, and the other consisting of the infinite set of fully determining (i.e. proposition-encoding) eternal sentences, which can be employed when total explicitness, leaving no room for interpretive manoeuvre, is called for. The relation between the shorthand-type sentences and the eternal ones must be a many-to-many mapping, a particular mapping in any given instance being determined by the context within which the convenient abbreviation is uttered. Katz concludes his discussion of the view that languages provide a large stock of eternal sentences as follows:

The only alternative to [this view] is . . . a form of mysticism that claims that some things to which we can refer by the use of indexical elements are, in principle, beyond the range of unique description. (Katz 1972: 127)

He has maintained this position over the years as he has developed his Platonist view of language (Katz 1981), within which it sits comfortably. According to the Platonist view, languages (hence sentences) are abstract objects whose properties can and should be investigated independently of their instantiations in human minds. The analogy is with systems of mathematics (hence numbers) and logic (hence propositions), whose properties have been extensively investigated independently of the mathematical or logical knowledge represented in human mind/brains. The contrast is with Chomsky's conceptualist (mentalist) view, according to which language is a natural object, the only reality it has being in the form of a, largely genetically programmed, system of knowledge, one component of a much more complex natural object, the human mind/brain.⁶ Katz (1981) presents some interesting arguments for the Platonist view; however, it would be too great a digression for me here to review these or the various counterattacks that have been launched (see Fodor 1981a; Chomsky 1986; and, for a recent sustained defence of a 'naturalistic' and 'internalist' approach to language, Chomsky 1995). Suffice it to say at this point, it probably makes better sense within a Platonist conception of language than within any other to posit the view that sentences, one kind of abstract object, map directly on to another kind of abstract object, namely propositions.⁷

1.3.2 Effability principles

The property of 'effability' or 'expressibility' concerns the extent to which it is possible, through the use of a public language system, for us to make our thoughts

available to others. The most general formulation of a principle of effability is along the following lines: ‘each proposition or thought can be expressed by some sentence in any natural language’. Much hangs on what is meant by ‘expressed’ here. In the previous sections, when I have talked of a proposition or thought expressed, I have not assumed this meant that it was ‘encoded’, or fully formulated, by a linguistic expression, quite the contrary in fact. But, as used by Katz (1978, 1981), ‘can be expressed by some sentence’ would seem to mean ‘can be encoded by some sentence’. So there are at least the following two, very different, possible principles to be considered:

First Principle of Effability: ‘Each proposition or thought can be expressed (= conveyed) by some utterance of some sentence in any natural language.’

Second Principle of Effability: ‘Each proposition or thought can be expressed (= encoded) by some sentence in any natural language.’

Note that while the second of these entails the existence of eternal sentences in the language (one for each proposition or thought), the first does not. The first principle is quite weak and seems largely unobjectionable. It refers to utterances in context rather than to abstract sentences and makes no stipulations about what the linguistic expression used must encode; it leaves open the possibility that much of the determining of the precise conceptual content of the thought is effected by means other than linguistic coding. Given no arbitrary limits on the richness of contexts or on ways in which contextually available material can be used to supplement encoded material, this effability principle does not raise too many problems. Of course, an individual speaker may not have the ability to express verbally a particular thought she has, but that does not touch on the claim, provided the thought *could* be expressed (by a more able speaker) in some context.

The second principle is considerably stronger and open to the following objection:

It seems plausible that in our internal language we often fix time and space references not in terms of universal coordinates, but in terms of a private logbook and an ego-centred map; furthermore, most kinds of reference – to people and events for instance – can be fixed in terms of these private time and space coordinates. Thoughts which contain such private references could not be *encoded* in natural language but could only be incompletely represented. (Sperber and Wilson 1986a/95b: 192)

Clearly, the conception of ‘thought’ here is finer than that of ‘proposition expressed’ or Fregean thought; it is the individualist psychological notion of thought mentioned at the beginning of this section and it, indeed, does not seem amenable to natural language encoding. The force of this point is perhaps most vividly felt by considering thoughts one has about oneself; how I represent myself to myself must inevitably be quite different from the way you or anyone else represents me, and so it must be for all of us. The same holds for the way I mentally represent my spatial and temporal location at any given instant, that which I might express by the words ‘here’ and ‘now’; your representation of my here and now is likely to be very dif-

ferent from mine and my representation of your here and now is likely to be very different from yours. This is a function of the ‘ego-centred map’ referred to in the quote and it extends far beyond these self-references. My mental representation of the woman who is my mother is doubtless a private one, probably not even shared with my siblings. This can be extended step by step to all of the people I have encountered in my life, and to all the activities and events I have taken part in or observed. My mental representation of the cup of coffee in front of me on my desk at this moment is determined by its relation to me, and it is that representation that enters into my current thoughts about it; it would be differently represented by another person sitting elsewhere in the room and, to that extent at least, the thoughts he might have about it would differ from mine.

Recanati (1993, 1994) endorses this line of argument. In his terminology, a *de re* thought (that is, a thought about a particular object) involves a particular ‘mode of presentation’ of that object and that particular mode of presentation may be entirely private, that is, peculiar to a given individual.⁸ So the *de re* thoughts of two people predicating the same property of the same object are generally distinct from each other and may also be entirely private.⁹ These sorts of differences in representations of an object are not, and cannot be, encoded in natural-language sentences. Imagine for a moment that they could be, perhaps by rapidly making up new audible or visible signals for each different psychological mode of presentation (assuming we have sufficient awareness of these). Such a process would be totally counterproductive since these signals could not be used for the very communicative purpose for which they were supposedly being invented; my public sign for the cup of coffee (which would have to change as my mental representations of it changed) would be meaningless to you; your sounds expressive of your mental representations of the people around you would be meaningless to me. These ‘signals’ could not be used in the way that natural-language referring expressions are used; they present an unsolvable coordination problem. So they would not in fact be ‘linguistic’ symbols, properly speaking; they could not acquire that status as they would have no hope of settling into the language system.

Recanati (1994: 157) considers (but does not endorse) another effability principle:

Third Principle of Effability: ‘For every statement that can be made using a context-sensitive sentence in a given context, there is an eternal sentence that can be used to make the same statement in any context.’

For the purposes of the discussion to follow, I shall assume that there is no crucial difference between ‘statement made’ and ‘proposition expressed’, and, therefore, that this principle could be reformulated using the latter term. Though it is clearly stronger than the first principle in which expressibility is equated with communicability, this third principle is not as strong as the second one. The shift from ‘thought’ (on an individualist interpretation) to ‘statement made’ (or ‘proposition expressed’) secures this principle against the objection just considered. There are *de re* statements (and propositions), that is, statements (or propositions) which predi-

cate a property of an entity, but these are different from *de re* thoughts in at least one essential respect: they do not contain private (psychological) modes of presentation of the object they are about. As Recanati (1994: 157–8) puts it, ‘a *de re* statement corresponds to a class of [*de re*] thoughts, each involving a particular (and, perhaps, private) mode of presentation of the object referred to. . . . statements are public objects at a more abstract level than thoughts, and as such do not contain private modes of presentation.’ So the possibility of eternalization of statements made (propositions expressed) needs to be given separate consideration.

This third effability principle seems to be implicit in the quote in the previous section from Katz (1972: 126), where the focus is on eternal sentences which encode *statements* made. This carries through to his subsequent Platonist position on which ‘Every proposition is a sense of some sentence in each natural language’ (Katz 1978: 216), where propositions are to be understood as the abstract entities themselves, rather than as mental representations of them. If this principle is correct, then it is not out of necessity that speakers use non-eternal (context-sensitive) means of expressing propositions, but it must be for some other reason. It comes as no surprise, then, to find that Katz takes the convenience view of linguistic underdeterminacy: ‘it [underdeterminacy] allows speakers to make use of contextual features to speak far more concisely than otherwise. . . . Pragmatics saves us from . . . wasteful verbosity’ (Katz 1977: 19–20). Thus our capacity for pragmatic inference, on this view, is a useful add-on to our language capacity, not strictly essential in making possible the sort of expressive and communicative powers we have.

In the rest of this section, I present evidence against this third effability position, drawing on arguments from Wettstein (1979) and Recanati (1987b, 1994). Then, in section 1.4, I will outline a view of pragmatic interpretation and of the cognitive capacities that underpin it (the relevance-theoretic view), from which, I believe, the essential nature of linguistic underdeterminacy follows. The picture is the opposite of Katz’s in every respect: only the first principle holds (on which effability = communicability), there are no eternal sentences, and pragmatic inference is fundamental.

Wettstein (1979) specifically addresses Katz (1972) and argues against his view that, for any statement made (proposition expressed) by the use of a non-eternal sentence in a particular context, there is an eternal sentence that can be used to make the same statement (express the same proposition) across contexts (or in the hypothetical absence of any context). A crucial part of the process of ‘eternalizing’ a non-eternal sentence is the replacement of each indexical expression by a non-indexical expression which picks out, in all contexts, for all time, the object that was referred to by the given use of the indexical in a particular context. According to Frege, Quine and Katz, what achieves this is a uniquely denoting description. Wettstein points out that the object concerned can be picked out by a range of *non-synonymous* descriptions. For instance, the pronoun ‘she’ in an utterance of (23a) can be replaced by a variety of descriptions, including those given in (23b)–(23d). Assuming, as Wettstein does, that there are such things as uniquely denoting descriptions and that these are likely candidates, each of them may denote the woman in question. (I haven’t attempted to eternalize the predicates.)

- (23) a. She left in a hurry.
 b. The woman who spoke to Tony Blair at t_1 left in a hurry.
 c. The woman in the red velvet dress who was in the Islington Town Hall between t_1 and t_2 left in a hurry.
 d. The middle-aged lady who organized the anti-junk-food campaign in June 1999 left in a hurry.

But these cannot all be part of the complete formulation of the statement the speaker made, and it is not clear that any one of them rather than any other, or indeed that any of them at all, does in fact figure in the statement made by the speaker of the original indexical sentence. As Wettstein (1979: 94) says, ‘Since these descriptions are not synonymous, it would seem that each of the resulting eternal sentences formulates a *different* proposition. The genuine eternal sentence counterpart will be the one that actually formulates the proposition the speaker asserted. But is there clearly one of these eternal sentences that, as opposed to the others, actually formulates what was asserted?’ He argues that there is no basis for thinking that any one among the several non-synonymous eternal sentences is the one that corresponds to the proposition expressed by the indexical utterance, and that this is because, in fact, none of them does.¹⁰

The point carries over equally to attempts to eternalize patently incomplete definite descriptions like ‘the table’, ‘the child’, ‘the government’, etc., whether used attributively or referentially. There are many possible completions all of which may uniquely denote the entity picked out in context by the original non-uniquely denoting description, but no one of which is *the* eternalized version of the incomplete description (see Wettstein 1981). The point also extends in an obvious way to certain subsentential utterances. Suppose two people are talking at a party and one of them, looking in the direction of a man near the door, says ‘Tom’s father’. As Stainton (1994) has shown, a speaker can use ‘Tom’s father’ to assert of a certain man that he is, say, the father of Tom Adams. But, of course, the search for the ‘right’ eternal sentence to encode what is asserted here is fairly unconstrained; either of the following, or any number of others, might do:

- (24) a. The man wearing a pink tie and drinking a martini at t_1 is the father of Tom Adams.
 b. The nervous-looking man standing in the doorway of the kitchen of Tom Adams’s flat at t_1 is the father of Tom Adams.

As Wettstein would doubtless say about this example, it is not clear that any particular one of these eternal sentences, as opposed to any other, ‘actually formulates what was asserted’. The issue cannot be resolved by referring to the speaker’s intention, since very often she will have no determinate intention; if asked which of the various possible eternal sentences correctly formulates the statement she intended to make she will be unable to answer. Indexical reference (and unarticulated reference, as in this last case) appears to be irrevocably context-bound; it is not reformulable in terms of a uniquely denoting description, but depends on the addressee’s

capacity to identify the intended entity by some means which is non-linguistic or, at least, not wholly linguistic.

This argument effectively undermines the third effability principle, held by Frege, Quine and Katz, and so supports an essentialist view of linguistic underdeterminacy. The absence from the language of any eternal sentence which encodes the proposition expressed by an indexical sentence (uttered in a context) is sufficient to achieve this, even if there are, in fact, some eternal sentences in the language. Wettstein seems to assume that there *are* such things as eternal sentences and that the elaborated sentences given above are candidates; similarly, Sayward (1968), who also shows that there is no eternal sentence equivalent for a range of indexical utterances, assumes that, nonetheless, there are some eternal sentences in the language. On this view, there is generally no mapping between indexical sentences (uttered in a context) and eternal sentences; rather, they are complementary in their expressive powers.

So the essentialist position established is the weaker one, which is good enough for my purposes. However, the head-on attack against the strong effability principles (the second and third) would be to deny the existence of eternal sentences altogether, as Recanati (1994) does. His discussion extends beyond the usual concentration on referring expressions to some consideration of predicates, quantifiers and tense. I will take a brief look in the next two subsections just at his arguments against the existence of eternal referring expressions and of eternal predicates.

1.3.3 Eternal reference?

Although Wettstein establishes that any process of substitution of an indexical by a complete description leads to a difference in the proposition expressed (and statement made), he doesn't question the existence of uniquely denoting descriptions. In fact, he speaks of a continuum of descriptions, from the patently incomplete 'the table', to more complete but still indefinite definite descriptions, and, finally, to a uniquely denoting description, usually containing proper names and precise temporal specifications, such as 'the table in room 209 of Camden Hall at t_1 ', where t_1 is, presumably, to be replaced by some eternal natural-language phrase (Wettstein 1981: 253–4).

Recanati (1987b, 1994, 1995) suggests that there simply cannot be reference (or unique denotation) without a context, that reference is always a pragmatic context-dependent matter. To establish this very broad claim, the case has to be made for proper names and for complete definite descriptions (whether used attributively or in the strongly referential way discussed by Donnellan 1966/91). It would take some considerable time and space to do this in anything approaching a conclusive fashion, given the vast range of work in the philosophy of language on proper names and descriptions. I shall not attempt that here. As regards proper names, the approach which treats them as a variety of indexical is intuitively appealing: just like pronouns, one and the same proper name may refer to different individuals in different contexts. Note that Katz (following Frege and Quine) included them in the class of cases which are to be replaced by a uniquely denoting description in the process

of formulating an eternal sentence. In my view, the analysis of the linguistic meaning of a proper name as some sort of rule requiring that it refers to a bearer of the name seems to be on the right track. This position is defended by Bach (1987) and Recanati (1993, chapters 8–9). Whether it turns out to be correct or not, proper names are not eternal, any more than any other indexical, so nor are descriptions that contain them.

Recanati (1987b) goes on to consider apparently complete descriptions like ‘the prime minister of Britain in 1999’. Of course, the names of countries and nations like ‘Britain’ are not essentially different in their semantics from the names of persons (although multiple use of these names is more studiously avoided for practical reasons), so they are also indexical and any description containing them is non-eternal. However, Recanati takes up a different line of argument: he sets out to show that the reference of a definite description, even one with no indexical element, always depends on the ‘domain of discourse’, what Fauconnier (1985) terms a ‘mental space’ and Barwise and Perry (1983) call a ‘resource situation’. The domain of discourse is ‘that with respect to which the speaker presents his or her utterance as true’ (Recanati 1987b: 62). Possible domains of discourse are the actual world, a fragment of the actual world (say, the current political situation in Britain), someone’s belief-world, a fictional world, a fragment of some counterfactual world.

He discusses the following sort of case: you and I know that Lucinda wrongly believes that Peter Mandelson is the prime minister of Britain in 1999. Knowing that Mandelson is in the next room, I utter (25) to you (in 1999):

- (25) Lucinda will be delighted to find that the current prime minister of Britain is in the next room.

I am here using the definite description to refer to Mandelson rather than the actual prime minister (Tony Blair), because I intend it to be interpreted with respect to Lucinda’s belief-world, within which Mandelson is the British prime minister. Of course, this example has a metarepresentational flavour, involving as it does an implicit attribution to Lucinda of a particular belief, and it could be argued that this somehow takes it outside the realm of the cases under discussion. I don’t think any such objection can be sustained, since this sort of use is utterly commonplace and merely reflects the fundamental point that the linguistic system is a tool, which does not have fixed communicative content but can be used in a range of ways by human communicators. The proposition expressed by the given utterance of (25) includes as a constituent (a concept of) someone who, on a different utterance of the same sentence in 1999, would not be picked out by the description ‘the current prime minister of Britain’; the same holds for the arguably more complete description, ‘the prime minister of Britain in 1999’.¹¹

This relativity of reference to the domain of discourse extends to those complete descriptions which have seemed the least likely to yield to the general context-dependence thesis. These are cases of so-called rigid descriptions where the semantic value of the description is the same across all possible worlds, for instance ‘the cube root of 27’. While 3 is the cube root of 27 in all possible worlds, it is not so in all domains of discourse, since Lucinda might believe that 9 is the cube-root of

27, and, as we have seen, Lucinda's belief-world is a possible domain of discourse. Recanati (1987b: 64–5) goes on to show that when a domain-indicator like 'Lucinda thinks that . . .' or 'According to Lucinda . . .' is given explicitly, as a constituent of the sentence uttered, its own interpretation is relative to a domain of discourse, which may or may not be the actual world, and which has to be pragmatically inferred, for instance, 'In Paul's story, Lucinda believes that . ..'. So there is an unavoidably context-dependent dimension in understanding any utterance. Even the 'normal' (some would say 'literal') understanding, where the description is interpreted with respect to the actual world, is a pragmatic matter, as it is dependent on the identification of the domain of discourse as the actual world. On this basis, even the best candidates for eternal sentences, for instance those in (22), are, in fact, context-dependent.

Not all subject noun phrases involve reference to particular entities; quantified noun phrase subjects like 'everybody', 'some students', 'most films', etc., are not referential. Of course, to be correctly understood these require a domain of quantification (e.g. 'some students doing the BA Linguistics at University College London in 1998/9') which is supplied contextually. Could the domain be explicitly described, that is, could the quantified phrase be made eternal? Again, this seems highly improbable because of the indexical nature of elements in the domain specification. Recanati (1994) argues further that, even if the quantificational domain could be given fully explicitly, the correct interpretation of it could only be made relative to the intended domain of discourse (mental space). Consider in this regard the phrase 'some of the spin doctors of the prime minister of Britain in 1999', interpreted relative to Lucinda's belief-world as described above. So it looks as if no quantificational sentence is eternal either.

If natural language predicates can also be shown to be non-eternal, then that will be yet another source of underdeterminacy in the descriptions that are supposed to replace indexicals and achieve a timelessly fixed specification.

1.3.4 Eternal predication?

The *ineffability* of reference provides perhaps the clearest argument against eternal sentences and those principles of effability which entail the complete linguistic encodability of the proposition expressed by an utterance. However, the case can be extended to include the predication function of language, the assigning of properties and relations to the entities picked out by referring expressions. That this is so has been shown in some detail by Travis (1981, 1985, 1991), Lahav (1989, 1993) and Gross (1998). Travis considers simple examples like the following:

- (26) a. The kettle is black.
 b. The table was covered with butter.
 c. Hugo is a sailor.

Discussing what is meant by the predicate 'black' in (26a), he outlines a range of possible circumstances in each of which it has a different interpretation:

Suppose the kettle is normal aluminum, but soot covered; normal aluminum but painted; cast iron, but glowing from heat; cast iron but enamelled white on the inside; on the outside; cast iron with a lot of brown grease stains on the outside; etc. (Compare a postage stamp, black on one side – a black stamp?, a ‘yellow’ labrador retriever painted to look like a black one – is the dog black? a ‘black’ narcissus, with a green stem; the North Sea [look at it from the deck on a normal North Sea day, then pull up a bucket of it and look at that].) (Travis 1985: 197)

His point is that the sentence in (26a), like virtually all sentences, may be used to say any of indefinitely many distinct things, each of which is true under different conditions. The bearer of truth is not the sentence but the proposition the speaker uses the sentence to express on the given occasion of utterance. One of the sources of these propositional differences in (26a) is the property communicated by the adjectival predicate ‘black’, both what property that is (clearly visible black, a wider colour spectrum taking in various dark browns, invisible black) and what exactly it is taken to apply to (the whole kettle, the inside, the outside, most of the outside, etc.). The same issues arise for (26b) and (26c); Travis spins numerous possible ways in which butter might be conceived of as covering a table and Hugo might be conceived of as a sailor (that is, different sets of truth conditions for different occasions of use of the sentence). This has interesting implications for the project of formulating a truth-conditional semantics for natural-language *sentences*, some of which are discussed in section 1.5.1; the context-sensitivity of predicates, in fact, proves to be a greater problem for the T-sentence approach than the context-sensitivity of reference.

Consider the following example, discussed by Sperber and Wilson (1997/98a: 192–5), where the concept communicated by a predicate is more specific than the concept it encodes:

- (27) A: Do you want to go to the cinema?
 B: I'm tired.

Most of us are tired to some degree or other most of the time; what B communicates by the predicate ‘tired’ in this context is something much more specific than that, something stronger perhaps than the sort of tiredness that makes an undemanding evening out quite agreeable, but not as strong as the sort of tiredness that necessitates going to bed immediately. Just how narrowed down this *ad hoc* concept of tiredness is will depend on other contextually available information, perhaps concerning B’s general energy levels, her liking for the cinema, etc. The prospects for finding another lexical item or even a lengthy description which fully encodes the concept of tiredness communicated here, and still others that encode the innumerable other concepts of tiredness that may be communicated in other contexts, look dim. In other words, as well as not uniquely determining the objects they can be used to refer to, natural language expressions seem to be intrinsically underdetermining of the properties and relations they may be used to predicate of an object.

Following the well-known example of Austin (1962), there is a range of examples which can be described as cases where the concept communicated (as opposed

to the one encoded) depends on the standard of precision relevant in the particular context:

- (28) a. France is hexagonal.
 b. This steak is raw.
 c. The fridge is empty.

These are cases where, arguably, the predicate is clearly defined and the definition is part of the native speaker's knowledge of the language; for instance, 'a hexagon is a geometric figure with six equal sides', 'something is raw if and only if it has received no cooking', etc. However, the proposition expressed on particular occasions of use might vary considerably depending on the degree of looseness the context allows or calls for.¹² The addition of modifiers such as 'approximately', 'to some extent', 'more or less' will make it explicit that the following predicate is not being used strictly but, clearly, they will not effect a full encoding of the intended concept of hexagonality, rawness, etc. Eternalization does not look like a possibility here.¹³ In fact, as I will suggest in chapter 5, the proposition expressed (what is said) by these cases may in some instances be indeterminate; there may be no absolute fact of the matter; speaker and hearer may diverge somewhat in the propositional form they entertain, though in ways that are quite innocuous as regards the success of the communicative interaction.

It is worth noting that Recanati's concept of a 'domain of discourse' needs to be applied with care in these predication cases. Consider a situation in which we know that Lucinda believes that cats are primates. On an application parallel to the case of the denotation of 'the prime minister of Britain' in the previous section, I could say to you, of my ancient tabby, 'Dear old Fleabag is a primate', where the domain of discourse pertaining to my use of the predicate 'primate' is Lucinda's belief world. Or consider the case of a young child who calls all four-legged animals of a certain size 'dogs', leading me to say to you 'Fleabag is a nice dog', where 'dog' must be understood relative to the child's belief world. But these seem to be different in an important way from the reference cases in that both Lucinda and the child are talking a different language from the one that you and I know; they assign a different ('wrong') sense to the words 'primate' and 'dog'. As far as I can see, the interpretations you would recover from these utterances are 'Lucinda thinks Fleabag is a primate', 'The child thinks Fleabag is a dog', where you would understand 'primate' and 'dog' in accordance with your knowledge of the language and not as communicating some different (wider) concept. In the reference cases, recognition of the relevant domain of discourse plays a crucial role in picking out an intended referent despite the fact that the description used for the purpose does not apply to him/her in the actual world. In that case, the belief world called upon involves idiosyncratic beliefs about entities in the world (e.g. that Mandelson is the prime minister of Britain) rather than idiosyncratic linguistic encodings. As I understand it, the idea of varying 'domains of discourse' concerns distinct belief worlds in which, nonetheless, the language has the same meaning. So there is an important difference here between reference and predication, no doubt just one manifestation of a more general asymmetry between the two phenomena.

To conclude, I should note that although I've concentrated on utterances of declarative sentences, the points made carry over to all the other sentence types: interrogatives, imperatives, optatives, hortatives, subjunctives, etc. On the whole, the advocates of eternal sentences, generally logicians or philosophers of language who put a premium on the statement-making function of language, especially the stating of scientific truths, have not been much interested in interrogatives and even less in sentences in the other moods. Katz, as a semanticist of natural language, does not confine his interest in this way; it is clear from Katz (1977), where he explores the semantics of sentences used to express a range of illocutionary forces, that he intends his effability principle to apply across the board. A belief that indexicals can be replaced by uniquely denoting non-indexicals must apply to sentences of all mood types, and so too a belief that vague or open-textured predicates can be replaced by well-defined ones. It follows then that the arguments against the existence of eternal declarative sentences apply equally against the existence of eternal imperatives and interrogatives.¹⁴

Summing up: the position I've been arguing for is that there are no eternal sentences in natural languages (that is, no sentences which encode a proposition or thought which is constant across all contexts), from which it follows that the linguistic underdeterminacy of the proposition expressed by an utterance is an essential feature of natural language. Neither of the strong principles of effability (the second or the third) applies to the semantic structures provided by natural languages. This is not to say, of course, that linguistic expressions, though inevitably non-eternal, cannot be used in appropriate contexts to communicate most, if not all, the propositions which humans are capable of instantiating in thought. This may well be so. In the next section, I will try to sketch a view of the human mind from which both this unbounded communicability of thought *and* the absence of eternal sentences in natural language follow, if not necessarily, very naturally.

1.4 Metarepresentation, Relevance and Pragmatic Inference

1.4.1 Mind-reading and ostension

Dan Sperber has emphasized the following fact about human cognition: while observed behaviour can, in principle, be conceptualized both in purely physical terms and in mentalistic (intentional) terms, we almost inevitably go for the latter (Sperber 1994a, 1996). Imagine observing a scene in which a man slowly lowers himself, head and arms first, down into a hole in the ground while another man holds on to his legs. Very few observers will represent this scene to themselves as I have just described it and leave it at that; most of us will look for some plausible beliefs, desires and/or intentions that we can attribute to these two men, some set of mental states which will explain their behaviour. For instance, we may attribute to both men a *belief* that there is something worth retrieving down in that hole, to the first man an *intention* to retrieve it, to the second man a *belief* that the first may fall into the hole and hurt himself if his legs are not held, etc.

We can't help doing this sort of thing, that is, we can't help attributing beliefs, desires and intentions, with quite specific content, to others; it seems to be built in to our cognitive system for interpreting the behaviour of our fellow humans. This capacity is more intelligent than one that simply assumes that every observed outcome of human action is an intended outcome; it is able to consider the sorts of beliefs and intentions people are likely to have and those they are not likely to have. So if the second man loses his grip on the first man's legs and the first man emerges some time later covered in slime, we will recognize these physical happenings as undesirable to the men and so not to be explained in terms of any intentions they had. Or, if there is some desirable outcome to their behaviour but which the men could not reasonably be expected to have foreseen, we will not try to explain it in terms of their beliefs or intentions. This capacity is not confined to one or two levels of attribution nor to attributions which involve but a single cognizer: you can attribute to *me* an *intention* to get *you* to *believe* that some *third person* does not *want* to go to a party. The mental faculty responsible for this is known as the 'theory of mind' or 'mind-reading' capacity and there is now a huge psychological literature on its nature, its place in our overall cognitive architecture, how it develops in infancy, its impairment in certain pathological conditions, such as autism, and its rudimentary manifestation in some other primates (see, for example, Astington, Harris and Olson 1988; Premack 1988; Baron-Cohen 1995; Smith and Tsimpli 1995, 1996; Scholl and Leslie 1999).

On the representational/computational view of the mind (which is assumed throughout this book), having a cognitive mental state like a belief or an intention involves being in a relation to a mental representation (a conceptual sentence in the language of thought, perhaps; see Fodor 1978, 1980). For instance, my believing that Felix is a cat involves my being in a belief relation to the representation 'FELIX IS A CAT' (or, as it is sometimes put, having a token of this representation in my belief box). Given this general picture, it can be seen that a crucial feature enabling the attribution of a mental state to someone, which may itself involve the attribution of a mental state, is a capacity for *metarepresentation*, that is, an ability to represent not just states of the external physical world but also other representations, and representations of still further representations, etc., up to several orders of complexity. This capacity makes it possible for us to reflect on our own mental states; for instance, to recall our former cognitive selves as consisting of beliefs, desires and hopes that may have been superseded. It makes it possible for us to hold reflective beliefs, that is, beliefs embedded in attributive phrases (such as 'The wise elders say that . . .'), which may be at odds with basic factual beliefs which we hold, or which may duplicate the content of factual beliefs; for instance, certain religious or mythical beliefs, on the one hand, and certain scientific or theory-embedded beliefs, on the other (see Sperber 1982/85, 1997a). Most important in the current context, the metarepresentational capacity makes possible the kind of communication which appears to be unique to humans: ostensive-inferential communication. Sperber (1994a) claims that adult communicators employ at least fourth-order meta-representations, and that the interpretation of utterances and other ostensive behaviours requires inferential processes involving premises of several meta-representational levels, which hearers perform with ease.

Continuing the scenario from above: suppose the second man, who is holding the legs of the first, swivels his eyes leftwards in our direction and starts to jerk his head quite violently from left to right. It is likely that we'll take him to be communicating something to us, that we'll take the head movement to be, not some involuntary tic he developed upon seeing us, but rather a movement designed to make it evident to us that he wants our attention and has something to tell us. We might even hazard a guess at (infer) what the intended message is, something like 'I want you to help me' perhaps. Note that this is achieved *without any element of encoding* whatsoever; the same type of head movement would be interpreted in quite different ways in different situations. Ostensive behaviour of this sort is explained by the attribution to their originators of a particular sort of intention, which Sperber and Wilson (1986a/95b: 50–64) call a 'communicative intention'. This is an intrinsically higher-order mental state, hence requires metarepresentation, as it is an intention to make manifest (or evident) an intention to inform someone of something (that is, to say, tell, ask, make known something). (See appendix 1 for a stricter definition of a communicative intention.) For instance, in the given example, the man who is making the ostensive head movements intends thereby to make it manifest to us that he intends us to recognize that he wants us to help him.¹⁵

Sperber (2000) argues in favour of a comprehension system whose domain is specifically utterances and other ostensive stimuli. This is a metarepresentational system and may be a submodule of the theory of mind (or 'metapsychological') system, to which it is clearly intimately related. The main argument for the claim that it is a distinct mental module hinges on the fact that the comprehension process requires a particular pattern of inference which distinguishes it from the inferential processes involved in interpreting non-ostensive behaviour. Someone observing the activities of the two men described above can impute to them certain intentions on the basis of an observed desirable outcome of their behaviour (e.g. the retrieval of a diamond ring). But in interpreting an instance of ostensive behaviour, the desirable effect (which is that the addressee grasp the communicator's meaning) cannot be achieved without the addressee's prior recognition of the communicator's intention to achieve that effect. That is, the standard pattern of inference from behaviour to identification of desirable outcome and then to intention is not available to the process of understanding acts of overt communication. Relevance theory makes a specific proposal about the particular computational strategy employed by a dedicated ostension comprehension module.¹⁶ I turn to that in the next section.

1.4.2 Relevance and utterance understanding

According to Sperber and Wilson's (1986a/95b) framework, relevance is a property of the inputs to cognitive processes (whether perceptual or higher-level conceptual); it is a positive function of cognitive effects and a negative function of the processing effort expended in deriving those effects. Cognitive effects (or contextual effects) include the strengthening of existing assumptions of the system, by providing further evidence for them, the elimination of assumptions that appear to be false, in the light of the new evidence, and the derivation of new assumptions ('contextual impli-

cations') through the interaction of the new information with existing assumptions. A basic principle of the framework is the **Cognitive Principle of Relevance** according to which the human cognitive system as a whole is oriented towards the maximization of relevance. That is, the various subsystems involved, in effect, conspire together in a bid to achieve the greatest number of cognitive effects for the least processing effort overall. The perceptual input systems have evolved in such a way that they generally respond automatically to stimuli which are very likely to have cognitive effects, quickly converting them into the sort of representational formats that are appropriate inputs to the conceptual inferential systems; these systems then integrate them, as efficiently as possible, with some accessible subset of existing representations to achieve as many cognitive effects as possible.

What distinguishes ostensive behaviour (including verbal utterances) from nonostensive behaviour (and, all the more so, from observed events that are not the result of volitional behaviour at all) is that it raises an expectation of a particular level of relevance in the relevance-seeking cognitive system of the addressee. A speaker (or more generally, an ostensive communicator) overtly requests an expenditure of mental effort from an addressee (an outlay of attentional and inferential resources) and that licenses an expectation of a worthwhile yield of cognitive effects with no gratuitous expenditure of effort. This is captured by the **Communicative Principle of Relevance**: every act of ostension communicates a presumption of its own *optimal relevance*; that is, a presumption that it will be at least relevant enough to warrant the addressee's attention and, moreover, as relevant as is compatible with the communicator's competence and her personal goals and preferences. The specific procedure employed by the comprehension system, on the basis of the presumption of optimal relevance, is given in (29):

- (29) Check interpretive hypotheses in order of their accessibility, that is, follow a path of least effort, until an interpretation which satisfies the expectation of relevance is found; then stop.

The least-effort strategy follows from the presumption of optimal relevance in that the speaker is expected to have found a vehicle for the communication of her thoughts which minimizes the hearer's effort (within the parameters set by the speaker's own abilities and goals/preferences); the justification for the addressee stopping processing as soon as an interpretation satisfies his expectation of relevance follows similarly, in that any other interpretation that might also achieve the requisite level of effects will be less accessible and so incur greater processing costs.¹⁷

The operation of this procedure, peculiar to the processing of ostensive behaviour, provides a solution to the apparent problem, mentioned in the previous section, that the desirable outcome (the grasping of the communicator's meaning) is dependent on a prior recognition of the communicator's intention. Processing by the addressee's pragmatic system employing the strategy in (29) is automatically triggered by an ostensive stimulus, irrespective of the actual intentions of the producer of the stimulus, and this strategy provides a reliable, though by no means foolproof, means of inferring a speaker's meaning. In fact, there are three distinct ways the

general strategy may be implemented, depending on the precise nature of the 'expectation of relevance' which an interpretive hypothesis must satisfy if it is to be accepted. With each of these three sorts of expectation there come differences in the complexity of the (meta)representational manipulations involved in the comprehension process.

The simplest case is an expectation of actual optimal relevance, which depends on assuming complete competence and good will on the part of the speaker; a more sophisticated expectation allows for speaker fallibility (for instance, gaps in her knowledge which make her unable to assess accurately what is relevant to the addressee at the particular time) and the most advanced expectation drops the assumption of inevitable speaker benevolence (after all, it might not be in the speaker's own interests to give relevant information to the addressee). Assuming an utterance with content P and assuming communicative success, on the most straightforward strategy (probably the default case), the result of comprehension is a second-order metarepresentation: she intends me to believe that P. In its most sophisticated version, the expectation is merely one of purported relevance so that the strategy is geared to finding, not the first sufficiently relevant interpretation, but the first one that the speaker could have thought would *seem* relevant enough to the addressee. This results in a fourth-order metarepresentation: 'She intends me to believe that she intends me to believe that P.' (For much more detailed exposition with worked examples, see Sperber 1994a.)

Expectations of relevance may also vary in their specificity. For instance, in the case of a conversation initiated by a stranger who happens to be sitting beside you on an aeroplane, you may have only quite a general expectation, while, in the case of an utterance made by a close friend in response to a question you've just asked, your expectation will be quite specific (that is, that she will give the information you asked for). The role of expectations of relevance of different degrees of specificity is explored in chapter 2 (section 2.3.4), where the relevance-based procedure is seen at work in the interpretation of some cases of verbal utterances.

Recall that, although Katz (1977) recognizes that, as a matter of fact, speakers rely on hearers' capacities to infer their informative intentions, his view is that this is just to save speakers the effort of (a) finding the natural-language sentence that fully encodes the proposition they want to express, and of (b) having to work their articulatory apparatus unnecessarily. According to this view, there is nothing fundamental about pragmatic inference. However, in the previous section, I made the case that the language system does not have the resources to encode the propositions speakers succeed in expressing, and what I am suggesting here is that there is a very good reason for this: that sort of expressive power is redundant. A powerful 'mind-reading' capacity is employed in the interpretation of human behaviour quite generally, with a specialized subsystem for dealing with ostensive behaviour including instances involving no coded element (such as various ostensive movements of parts of the face or body). The wide application of this capacity in human cognitive activity and its presence in a rudimentary form in apes, who lack a linguistic system, make it reasonable to suppose that the linguistic code evolved later than the general capacity to attribute mental states and also later

than, or in step with, the more specific capacity to attribute communicative intentions. If it had been in place in the absence of an ability to attribute beliefs and intentions, it would have been largely functionally inert, at least for communicative purposes.¹⁸

If these somewhat speculative thoughts are on the right track, it seems that nature designed (speaking metaphorically, as we do) a linguistic code that has just the expressive resources that are needed to supplement an already pretty effective interpretive system. A linguistic system is undeniably enabling; it allows us to achieve a degree of explicitness, clarity and abstractness not possible in non-verbal communication (try communicating the proposition just expressed without using a language), but it is not essential for the basic function of referring, and the predicates it offers are but a tiny subset of the properties and relations that humans can think about and communicate.¹⁹

Within the bounds of her competence and her personal concerns and preferences, a speaker's choice of linguistic form takes account of the hearer's immediately accessible assumptions, encoding just what seems to be necessary to direct the hearer's inferential processes to the intended interpretation; what the coded bits of an utterance do is 'set the inferential process on the right track'. A (very) rough non-psychological analogy might be with a constructed system of banks and trenches, which channel the inevitable downhill flow of a river in certain directions, diverting it from others it might go in if left to its own devices. Verbal communication, on this view, is not a means of thought duplication; the thought(s) that the speaker seeks to communicate are seldom, if ever, perfectly replicated in the mind of the audience; communication is deemed successful (that is, good enough) when the interpretation derived by the addressee sufficiently resembles the thoughts the speaker intended to communicate. That communication is often successful in this sense is due partly to the channelling provided by the linguistic code, partly to our innate ability to attribute beliefs and intentions to each other and, crucially, given our constant cognitive bid for relevance, to the prevailing presumption of optimal relevance carried by utterances (and ostensive stimuli generally).

Recall Katz's charge, quoted in the previous section, that the only alternative to the availability of 'eternal' sentences is some kind of mysticism. I believe the cognitive picture just outlined provides an alternative of a non-mystical sort; it is a solid empirical hypothesis about the nature of human communication, one that is rather well supported by what is currently known about human cognitive capacities. Part of the motivation for Katz's formulation of an effability principle was to try to capture the essence of natural language, that which distinguishes it from artificial languages, on the one hand, and from animal communication systems, on the other. There are, doubtless, properties of natural human language, such as the complexity and recursiveness of its syntax, which distinguish it from animal signal systems, but the sort of uniqueness that concerned Katz, the 'expressive power of natural languages', seems to follow not from the linguistic system itself, but from its recruitment by the distinct capacity for ostensive communicative behaviour, itself dependent on a highly developed metarepresentational ability for instantiating and attributing intentional states of several orders of complexity.²⁰

1.5 Underdeterminacy, Truth Conditions and the Semantics/Pragmatics Distinction

if you just take a bunch of sentences . . . impeccably formulated in some language or other, there can be no question of sorting them out into those that are true and those that are false; for . . . the question of truth and falsehood does not turn only on what a sentence *is*, nor yet on what it *means*, but on, speaking very broadly, the circumstances in which it is uttered. Sentences are not *as such* either true or false. (Austin 1962: 110–11)

This quote from one of the most famous of the ordinary-language philosophers is, in effect, a statement of the underdeterminacy thesis. The designation ‘ordinary-language’ is to distinguish the focus of these philosophers from that of preceding generations of philosophers whose energies had been engaged in the construction of ‘formal’ logical languages (Austin’s specific targets here are Carnap and Ayer). The aim of the formalists, I think, was to achieve the precise expression of scientific truths in a representation system which would wear its logic on its face, its logical implications following transparently from its form, as is the case with the predicate calculus, for instance. Non-natural formal languages are set up so as to be free from all those features that make natural languages underdetermining of propositions: they contain no indexicals, no ambiguous or vague predicates, no operators with unspecified scope, and no unarticulated or hidden constituents. On a given interpretation of the language, each of its well-formed formulas has context-free truth conditions and is either true or false, regardless of the context in which it appears.

However, it seems quite clear that if these formal languages were for some reason pressed into use by normally functioning human beings, in ordinary communicative situations, they would soon appear to lose their well-definedness. They would be used to communicate all sorts of propositions that, as originally constructed at least, they did not encode, that lay beyond their intrinsic (context-free) power of expression. Then, over time, if employed in this way, they would probably lose their original characteristic of being uniquely denoting: names would become indexical and many predicates would lose their univocality or would become very much more general in meaning. Why should this be so? It is communicatively convenient (that is, it makes for savings in the overall cognitive economy) to have forms that cause several senses to spring to mind or which have a quite general and open-textured sense. At least, it is convenient given the pragmatic inferential capacities of humans, which are acutely responsive to contextual considerations, enabling them easily and rapidly to choose among senses and to home in on more specific interpretations of a general sense.²¹ Note that this answer is, in effect, the reverse of the ‘pragmatic inference as convenient shortcut’, which I have rejected in the previous sections. Even if the integrity of the original, precisely defined system could be somehow (artificially) preserved, the basic pragmatic fact about natural languages would carry over to the use of the formal language: communicators would succeed in expressing, and being understood as expressing, propositions that this language did not encode.

The linguistic underdeterminacy view is now fairly widely endorsed by philosophers of language and pragmatists, in some form or other, to varying degrees and for different reasons. The outstanding exception to this is Grice, who appears to want to keep ‘what is said’ as close as possible to conventional (encoded) sentence meaning (conceding only indexicality). His reasons for this are explored in the next chapter. As we have seen, Travis (1981, 1985, 1991, 1997) and Recanati (1989b/91, 1993, 1994, 1995) hold strong essentialist positions on underdeterminacy, as do relevance theorists. Searle (1978/79, 1983, 1992) holds a radical underdeterminacy thesis which applies not only to language but to all intentional mental states (beliefs, intentions, thoughts); I will look at this in section 1.6.

Jay Atlas (1977, 1984, 1989) makes remarks that indicate support for the underdeterminacy view, albeit from the fairly restricted perspective of particular semantic analyses:

Meanings [of sentences] might be identified with mappings from points of reference into propositions, but not with the propositions themselves. If semantic representations represent meaning, they are not propositions or [Russellian] logical forms, though which propositions can be literally expressed by a sentence is determined by its semantic representation. (Atlas 1977: 332–3)

The sense-generality of a sentence radically underdetermines (independently of indexicality) the truth-conditional contents of its utterances. (Atlas 1989: 31)

His specific interest is in establishing his ‘sense generality’ thesis, as against appeals to ambiguity (multiplicity of sense), in the semantic analysis of several key areas of language, in particular negation (this is discussed in chapter 4 of this book). Kempson and Cormack (1981, 1982) and Bach (1982) have taken a similar line on sentences with several quantifiers (that is, that sentence meaning is neutral as regards their relative scopes). It follows from this semantic nonspecificity that sentence meaning is not fully propositional and there is an obligatory pragmatic process of scope fixing.

Within more linguistically oriented pragmatics, underdeterminacy has been noted by several authors (for instance, Fauconnier (1975, 1978, 1985), Dascal (1981), Green and Morgan (1981), Levinson (1988, 2000)), though not in the strong form I am advocating, that is, that linguistically encoded meaning *never* fully determines the intended proposition expressed. Levinson presents a huge range of data illustrating what he calls ‘pragmatic intrusion’ into the truth-conditional content of utterances. He finds this problematic since, in his view, it undermines the autonomy of semantics, which, following Gazdar (1979), he equates with truth conditions. It gives rise to an allegedly problematic circle because a crucial input to pragmatic inference is the semantic representation of the utterance, but pragmatic inference is necessary in order to establish that very semantic (truth-conditional) content (for disambiguation, reference assignment, supplying of unarticulated constituents, enrichment, etc.). I shall suggest below and in the next chapter that the problem comes from an implicit equation of sentence meaning with ‘what is said’ (= truth-conditional content) and an accompanying assumption that all pragmatic inference gives rise to implicatures.

1.5.1 A truth-conditional semantics for natural language?

The semantics of a formal logical language is typically given in terms of a truth theory for the language, which assigns to each sentential formula conditions on its truth; the propositionality and context-independence of the sentences of the language are important factors in making this feasible. The context-sensitivity and non-propositionality of natural-language sentences raises the question of how their semantics is to be characterized. A dominant view is that, notwithstanding indexicality, vagueness and incompleteness, sentence meaning *must* be given in terms of truth conditions. As is well known, the Tarskian ‘normal form’ for the statements of truth conditions (that is, the T-sentences, which are theorems of the truth theory for the language), alternative formulations of which are shown in (30), cannot be applied directly to the context-sensitive sentences of natural language.

- (30) a. S is true if and only if p
 b. ‘S’ is true if and only if S

If the p in (30a) is replaced by a context-insensitive clause, as it is in (31a) below, then the resulting T-sentence is false for the vast bulk of the occurrences of the object sentence, for instance, ‘I am happy’ when uttered by Bill Clinton, or by Madonna, or by you. If, on the other hand, the transparently disquotational form in (30b) is instantiated, as in (31b), then the whole T-sentence is context-sensitive. Either we have to find some way of capturing the different interpretations it has in different contexts, or we have to take it as relativized to one particular context, in which case, again, the sentence ‘I am happy’ will be assigned non-variable truth conditions which will be false on many occasions of use. The widely adopted solution is to employ a different kind of schema, one which quantifies over *utterances* and contextual features. An instance of this is given in (31c), where a single contextual factor, that of the speaker of the utterance, is incorporated.

- (31) a. ‘I am happy’ is true iff Robyn Carston is happy.
 b. ‘I am happy’ is true iff I am happy.
 c. If u is an utterance of ‘I am happy’, and x is the speaker of u , then [u is true iff x is happy].

So the ‘modified normal form’ for the statement of the truth conditions of natural language sentences is the conditional schema shown in (32). The truth-statement itself is given in the consequent (shown in square brackets in (32)), but is conditional on the various contextual parameters enumerated in the antecedent. This schema provides a means of meeting the requirement of giving context-variable truth conditions for a sentence without getting embroiled in the details of particular contexts:

- (32) If u is an utterance of sentence S, and $uR(x, y \dots z)$, then [u is true iff $F(x, y, \dots z)$]

The questions to be addressed briefly here are: (a) Can an account in terms of conditional truth-conditional statements capture the full range of encoded linguistic meaning? (b) Do these statements provide an appropriate input to the pragmatic processes involved in the recovery of full utterance meaning (that is, what is communicated)? In considering these matters, I refer principally to the framework of Higginbotham (1986, 1988, 1994), who advocates the application of the conditional truth-conditional approach, initiated by Davidson (1967) and Burge (1974), to the project of describing the semantic knowledge (competence) of native speakers. This work is taken to complement accounts of the phonological and syntactic knowledge of native speakers, each of these being distinct components within the generative grammar enterprise of characterizing native speakers' linguistic competence.

Higginbotham (1988) discusses the application of the approach to demonstratives and other elements used in referring. Consider, for instance, the accommodation of the context-sensitive expressions 'this', 'she' and 'that woman' by the following conditional T-statements:

- (33) a. If u is an utterance of 'this is green', and the speaker of u refers with 'this' to x , then [u is true iff x is green].
 b. If u is an utterance of 'she is happy', and the speaker of u refers with 'she' to x , and x is female, then [u is true iff x is happy].
 c. If u is an utterance of 'that woman is happy', and the speaker of u refers with 'that woman' to x , and x is a woman, then [u is true iff x is happy].

The linguistic meaning of 'I' in (31c), 'this', 'she' and 'that woman' in (33), which consists of constraints on what they can be used to refer to, occurs in the antecedent among the conditions that must be fulfilled for the T-statement in the consequent to hold. This sort of truth-conditional account thereby succeeds in capturing the semantic contribution of these linguistic elements whose meaning does not enter into the truth conditions themselves.

The approach is extended to 'incomplete' definite descriptions, such as 'the ginger cat', 'the murderer', on both their attributive and referential uses.²² On the attributive use, the descriptive content occurs in the consequent of the conditional statement as an essential element of the truth conditions. However, a second-order variable is introduced into the antecedent in order to indicate the contextual requirement that the domain of the description be appropriately restricted (Higginbotham 1988: 39):

- (34) a. The murderer (whoever s/he is) is insane.
 b. If u is an utterance of (34a), and the speaker of u confines the domain of 'murderer' to entities x such that X_x , then [u is true iff (ix : murderer(x) and X_x) (x is insane)].

The requirement of a contextually supplied domain restriction extends to quantifiers quite generally, e.g. 'Every student [in such and such a domain] must write a

dissertation’, and to tense and time-reference, e.g. ‘We worked hard [at times *t* within such and such a period of time]’. So the underdeterminacy (context-sensitivity) of a range of referring and quantificational expressions is represented as part of speakers’ semantic knowledge; certain pragmatic requirements (to find particular contextual instantiations on particular occasions of use) are marked out in the antecedent of these conditional semantic statements.

This approach raises a number of interesting issues. First, Higginbotham himself (1993b: 2), in discussing demonstrative cases such as (33a), says ‘the speaker’s perspective on what she speaks about – the object *x* – is wholly left out of account, so if we cannot be said to have understood the utterance without knowing that perspective then there will be aspects of understanding not covered by semantic rules’. In the current context of the strong underdeterminacy view, this comes as no surprise and does not strike home as a major worry; on this view, there are numerous aspects of the understanding of an utterance that are not encoded (either as truth-conditional content or as constraints on truth-conditional content) in the linguistic system, but are supplied on wholly pragmatic grounds.

Of more concern here is the fact that there are linguistic elements whose semantic contribution does not seem to be captured by this approach. Higginbotham (1994: 98) discusses the case of the ‘specific indefinite’ description, ‘a certain *F*’, which he compares with the simple indefinite description, ‘an *F*’, in examples such as the following:

- (35) a. A politician rang me yesterday.
 b. A certain politician rang me yesterday.

While there seems to be no truth-conditional difference between (35a) and (35b), hence no difference in their T-sentence specifications, there clearly is one when the simple sentence is embedded:

- (36) a. If a politician rings today, tell him I’m out.
 b. If a certain politician rings us today, tell him I’m out.

Similarly, as observed by Gazdar (1979: 166–7) and Seuren (2000), different patterns of what is called the topic-comment, or presupposition-focus, structure of a sentence are truth-conditionally inert in simple sentences, such as those in (37) and (38), but make themselves felt when embedded in certain other structures, such as those in (39) and (40):

- (37) a. Jane gave ME the tickets.
 b. Jane gave me the TICKETS.
 (38) a. It was Sam who won the champagne.
 b. It was the champagne that Sam won.
 (39) a. Jane gave ME the tickets by mistake.
 b. Jane gave me the TICKETS by mistake.

- (40) a. Mary was annoyed that it was Sam who won the champagne.
 b. Mary was annoyed that it was the champagne that Sam won.

The adverbial ‘by mistake’ in (39) is understood as pertaining just to the element highlighted by contrastive stress in each case, and the attitudinal verb in (40) as applying just to the element in the focus of the cleft construction.

Finally, there is a category of elements, noted by Segal (1994: 112) and Larson and Segal (1995: 44), whose encoded content seems to make no contribution to truth conditions or to constraints on truth conditions, in either simple or embedded sentences:

- (41) a. Luke likes Sam and Hank loves Rob.
 b. Luke likes Sam but Hank loves Rob.
 c. I am surprised that Luke likes Sam and Hank loves Rob.
 d. I am surprised that Luke likes Sam but Hank loves Rob.

The crucial element here is the much discussed ‘but’, whose truth-conditional contribution seems to be identical to that of ‘and’, although its inherent meaning clearly incorporates another feature (of ‘contrast’, roughly speaking). There are numerous other cases of this sort (often labelled devices of ‘conventional implicature’), for instance, ‘although’, ‘however’, ‘nevertheless’, ‘moreover’, ‘anyway’, ‘whereas’, ‘after all’, ‘even’, ‘yet’, ‘still’, ‘besides’, and on certain uses, ‘so’, ‘therefore’, ‘since’, and ‘while’.

In his discussion of the specific indefinite description, ‘a certain F’, Higginbotham (1994: 99–100) suggests the following adjustment to the conception of a semantic theory: ‘Suppose that the theory of knowledge of meaning gives us, not quite the truth conditions (or conditional truth conditions) of an utterance, but rather what a person who used the utterance to make an assertion would represent himself as believing.’ Very often, truth conditions and what a person represents herself as believing coincide, but on occasion they do not. In the case of an assertion of (35a), the speaker represents herself as believing that some politician or other rang her yesterday, that is, $(\exists x)[F(x) \ \& \ G(x)]$, while in the case of an assertion of (35b), she represents herself as having in mind a particular person and believing of him that he is a politician and that he rang her yesterday, that is, $F(\alpha) \ \& \ G(\alpha)$, where α is some individual concept or sense. Although these are clearly distinct beliefs, the hearer gets the same information in the two cases about how things must be if the world is as the speaker says it is. Higginbotham believes that a move of this sort will mop up most, perhaps all, cases of implicature and presupposition which are directly triggered by linguistic form. Although he does not specifically address the stress and focus cases in (37)–(40), I see no reason, in principle, why the idea shouldn’t extend to them as well. But what conclusion are we to draw from this about the nature of semantic description? According to Higginbotham, despite this ‘concession’ concerning the limits of truth-conditional semantics, truth conditions (and reference more generally) remain fundamental, since it is only they that have the properties needed for a systematic compositional semantic theory (Higginbotham 1994: 97,100).

Finally, let us consider the adequacy of the conditional truth-conditional approach in accommodating cases of predicate context-sensitivity. Here I closely follow Gross (1998, chapter 3) and focus on what he calls the ‘part’ context-sensitivity of adjectival predicates, as exemplified in the following:

- (42) a. The book is green.
 b. This fruit is smooth.
 c. Mr Jones is hairy.

As discussed earlier, the truth conditions of particular utterances of these sentences vary according to that part of the subject referent that is contextually relevant; in other words, the proposition expressed is linguistically underdetermined as regards the part to which the property is attributed. In the case of (42a), it could be the cover of the book which is the relevant part and so that to which greenness is being attributed, or it could be just a dominant part of a design on the cover, or just the spine (when one is scanning a shelf of books) or, in a less typical circumstance, the pages inside, and so on.²³ As a first try at capturing this kind of underdeterminacy, we might take the earlier treatments of referring expressions as a model and quantify over contextually relevant parts of the subject of predication, thereby registering that the pragmatic determination of the relevant part is an essential contribution to the truth conditions of an utterance of the sentence:

- (43) If u is an utterance of ‘ a is green’, and x is the contextually relevant part of a , then [u is true iff a is green-at- x]

As Gross points out, this raises a pressing question: what is it to be green-at- x ? Among the answers he considers are the following:

- (44) a. a is green-at- x iff x is green
 b. a is green-at- x iff x is wholly green

The problem with (44a) is obvious: the right-hand side is context-sensitive, and in just the same way as the original sentence, ‘ a is green’, whose truth conditions we are trying to specify. For instance, suppose ‘green’ is predicated of a book and the contextually salient part is its cover, then it is the cover that must be green. But what part of the cover – the whole of it, or all of it except the white lettering, or the background only, against which there is a gold and black design, or . . . ? The next suggestion is (44b), where, let us suppose, ‘wholly’ is not context-sensitive, so the right-hand side yields context-free truth conditions for each occurrence of the sentence (given a contextual value for x). The problem now, though, is that the truth conditions are wrong for a number of cases: suppose the utterance is ‘the apple is green’, and the contextually relevant part is the peel of the apple, which is indeed green except for a tiny patch of yellow near the stem. In such a case, the claim that the apple is green might be true even though it does not meet the requirement of being ‘wholly green’.

Gross considers a number of further variants of this general approach of treating the part context-sensitivity of predicates as a matter of referential indetermi-

nacy. Then, since these are unsuccessful, he tries locating the context-sensitivity elsewhere, that is, directly in the property expressed by the predicate. So the conditional T-sentence features a variable for the different properties the predicate can contribute to the truth conditions of different utterances of the sentence:

- (45) If u is an utterance of ‘ a is green’, and ‘green’ expresses property G , then
 [u is true iff a is G]

The trouble with this is that, as it stands, replacement of ‘green’ by ‘red’, or by any other colour term, will result in the same T-statement. In fact, not only will the truth conditions provided for all colour terms be the same, but they will be the same for all predicates which are part context-sensitive (e.g. ‘smooth’, ‘hairy’, ‘spotted’, ‘soft’, etc.). This is clearly unsatisfactory for a theory which aims at a comprehensive description of speakers’ semantic knowledge. One might look to impose some sort of constraints on the property that a given predicate can express, comparable to the constraints imposed by referring expressions like ‘she’ (the referent is female) or ‘that man’ (the referent is a man). The problem becomes one of finding sufficiently restrictive constraints coined in a vocabulary which does not introduce yet further context-sensitivity. For instance, in the case of ‘green’, the restriction might be something like ‘having to do with the colour of a certain part and not being red or blue, or . . .’, but ‘not red’ is just as context-sensitive as ‘green’ itself.

Gross concludes the discussion by considering the possible stance that it is not, after all, up to the truth-theoretic account to distinguish these predicates, that its purpose is just to characterize the type of semantic value assigned to linguistic categories (indexicals, demonstratives, names, classes of predicates) and to the different modes of combination allowed by the language. However, if limited in this way, the truth-conditional approach provides, at best, a very incomplete description of speakers’ semantic competence, and has to be complemented by some other account which captures speakers’ knowledge of the distinct meanings of different predicates.

Let’s return now to the two questions that were asked at the beginning of this section: (a) Can the truth-conditional approach account for the native speaker’s knowledge of linguistic meaning? (b) Do T-statements provide an appropriate input to the processes of pragmatic comprehension? In order to give a definitive answer to the first question, much more assessment is needed (including investigation of other truth-conditional frameworks), but the considerations gathered together in this section seem to indicate a negative conclusion. Several distinct aspects of encoded linguistic meaning looked at here do not seem to be amenable to the truth-conditional treatment. And, as the discussion in the following chapters will show, there are numerous other ways in which the proposition expressed by an utterance is linguistically underdetermined, some of which provide further challenges to the truth-conditional approach.

Clearly, the apparent shortfall in accounting for speakers’ semantic knowledge bears directly on the second question too; the elements not accounted for (for instance, the focus cases and the so-called conventional implicature examples), play a significant role in shaping the interpretation of the utterance that the hearer will recover, not only in the complex cases, where they may affect truth conditions, but

also in the simple cases where they do not. Their semantics is a crucial input to the processes of pragmatic interpretation. Recall Higginbotham's suggestion prompted by these cases: perhaps the account of semantic knowledge gives us, rather than truth conditions, what a person who assertively utters a sentence represents herself as believing. He sets the idea aside, since it doesn't seem to have the properties that a semantic theory needs (systematicity, amenability to combinatorial rules, etc.). I have to set it aside, because it seems quite remote from my concern for a linguistic semantics which can serve as appropriate input to an inferential processing account of utterance interpretation. What his suggestion highlights for me is how distinct the two overall approaches to linguistic meaning are; the difference is at least one of explanatory levels – the one concerned with characterizing a system of knowledge (a *competence*) in quite abstract terms, the other with finding a representational level (or levels) which enters into an account of a particular kind of mental *performance*, that of comprehending the intended meaning of a linguistic utterance. At the least, there would seem to have to be the mediating work of another performance system, the *parser*, which might access and deploy elements of both the syntactic and semantic competence systems in building a representation usable by the comprehension performance system. The relationship between the account of the on-line processes of utterance interpretation and the theory of knowledge of meaning is probably not a simple one; the components of neither one of them may translate into or correlate with, in any direct way, the components of the other.

Although they acknowledge the insufficiency of the truth-conditional approach, Higginbotham (1994) and Segal (1994) maintain that 'knowledge of conditions on reference and truth is the backbone of a theory of meaning'. Higginbotham makes no claim for the incorporation of the T-statements of semantic theory into an account of utterance processing, but Segal (1994: 112) and Larson and Segal (1995: 20–2) appear to do so. They locate the T-theory within a wider (modular) view of interacting mental systems and see certain of these systems as taking the output of the T-theory, that is, the T-sentences, and employing them (together with other sources of information) in such linguistic performances as understanding utterances and making judgements of meaning. So, unlike many truth-conditionalists, they are claiming that the truth theory for a language plays a direct role in the performance theory of verbal comprehension. The statements of the T-theory seem to me to be at one (or several) remove(s) from an account which will actually run, as it were, an account in terms of representations and processes (computations). However, even if they do play a more direct role in these processes than I envisage, they will have to be supplemented by a distinct account of the semantic contribution of the non-truth-conditional elements in the sentence (or phrase) uttered.

1.5.2 A translational semantics for natural language?

I turn now to the relevance-theoretic conception of linguistic semantics and its role in the cognitive account of utterance interpretation. The distinction between linguistic semantics and pragmatics is seen in performance terms, as closely tied to

a distinction between two types of processes involved in understanding utterances: linguistic decoding and pragmatic inference (see Sperber and Wilson 1986a/95b). The decoding process is performed by an autonomous linguistic system, the parser or language perception module. Having identified a particular acoustic (or visual) stimulus as linguistic, this system executes a series of deterministic grammatical computations, or mappings, resulting in an output representation, which is the semantic representation, or logical form, of the sentence or phrase employed in the utterance. It is a structured string of concepts, with certain logical and causal properties, but it is seldom, if ever, fully propositional. It is a kind of template or schema for a range of possible propositions, rather than itself being a particular proposition. As a 'schema', it is a formula that contains slots to be filled; what may go into a given slot may be partially constrained by a procedure which specifies how to go about filling it. Such a formula necessarily requires the second type of cognitive process, pragmatic inference, to develop it into the proposition the speaker intended to express.

On this construal, a characterization of an individual's linguistic semantics consists of a systematic description of the elements of meaning that the linguistic forms making up that individual's language encode. It will give an account of which conceptual representations, in what syntactic configurations, are activated in the mind, by which bits of phonological (or graphological) form. It will also include entries for all those linguistic forms that have been set aside by truth-conditionalists as cases of (non-truth-conditional) conventional implicature or presupposition. Interestingly, in the relevance-theoretic account, the distinction between the truth-conditional and the non-truth-conditional is largely reflected in a distinction between two types of cognitive information that a linguistic form may map on to: concepts, which are constituents of mental representations, and procedures, which do not enter into representations, but rather constrain the processes of pragmatic inference involved in deriving particular conceptual constituents of representations. This distinction within linguistic semantics is discussed in the next chapter (section 2.3.7).

Now, in the view of many truth-conditional semanticists, we are simply not going to be saying anything of semantic interest if our characterization of the meaning of sentences is not in terms of truth and reference. The familiar charge is that accounts which give the semantics of linguistic expressions in terms of a logical, or some other, notation are merely translating one kind of representation into another kind of representation, and you can go on doing that sort of thing *ad infinitum* without ever getting any closer to the essence of semantics, which is that it concerns a relation between representations and the non-representational external world they represent. The classic statement of this position is by David Lewis, in a critique of the approach of Katz and Postal (1964), who gave 'semantic' representations of natural-language sentences in terms of a system of what they called 'semantic markers'. Lewis says: 'But we can know the Markerese translation of an English sentence without knowing the first thing about the meaning of the English sentence: namely, the conditions under which it would be true. . . . Markerese semantics [does not deal] with the relations between symbols and the world of non-symbols – that is, with genuinely semantic relations' (Lewis 1970: 18).²⁴ If we substitute 'Mentalese' (the conceptual thought representation system) for 'Markerese', Lewis's objection

would appear to carry over directly to the position being advocated here, so what response do we have?

We can agree that, by giving an account of the encoded meaning of natural language sentences, we are essentially performing a translation into another system of representations, but resist the charge that this is a vacuous enterprise by insisting that it is this latter representational system which is given a 'real' semantics, that is, is related to the objects and states of affairs in the world which it represents. This, in essence, is the view of Jerry Fodor:

It's entirely natural to run a computational story about the attitudes [beliefs, intentions and other kinds of thought] together with a translation story about language comprehension; and there's no reason to doubt, so far at least, that the sort of translation that's required is an exhaustively syntactic operation.

... Syntax is about what's in your head, but semantics is about how your head is connected to the world. Syntax is part of the story about the mental representation of sentences, but semantics isn't. (Fodor 1989/90: 187)

See also Fodor (1975, 1981b, 1990). This position has been expressed often in the relevance-theoretic literature in talk of 'two types of semantics': (1) a translational linguistic semantics, which could be described in statements of the form '*abc* means (= encodes) '*ijk*', where '*abc*' is a public-language form and '*ijk*' is a Mentalese form (most likely an incomplete, schematic Mentalese form); (2) a 'real' semantics, which explicates the relation between our mental representations and that which they represent (so it must be 'disquotational') and whose statements may take the form '*hijk* means (= is true iff) *such-and-such*'. There are, no doubt, various qualifications called for here. For instance, the translational schema given in (1) involves the quoting (that is, the mention rather than use) of forms in the two distinct representational systems, and if this were the format of the mentally represented translations of each element of linguistic meaning, a process of disquoting would be essential at some stage in the compositional construction of the conceptual representation of the proposition expressed. As manifest in the minds of speakers/hearers, the translation process is better characterized as a direct mapping from a natural-language form to a Mentalese form, such that with the recognition of a particular linguistic element, for instance, the morpheme 'cat', comes the activation of a particular concept, say CAT.²⁵ What the translation process does is effect an interface between two distinct processing systems: the linguistic and the pragmatic.

Although this is the merest sketch of a picture, it would appear, at least in principle, to meet Lewis's objection: translation stops, truth conditions enter the account and the connection with the world is made. In fact, on an account of this sort, we could construe at least some natural-language expressions as having a 'real' truth-conditional semantics by inheritance; that is, given that they map on to parts of propositional thought representations, they can be thought of as having the truth-conditional (referential) content that those parts of the thought representations have. For instance, assuming the word 'cat' maps to the concept CAT and the concept CAT refers to (is true of) cats, then 'cat' inherits this referential semantics from CAT. Fodor (1975, 1998) argues that there is largely a one-to-one mapping of this sort between words (or morphemes) and concepts. To the extent that this is the case, the con-

struction of a truth-conditional semantics for natural language (really a (partial) truth-conditional semantics for thought) could proceed in parallel with the translational account advocated here. After all, we have a much better grip on the components of natural language than we have on those of Mentalese (assuming there is such a thing) so, given the putative one-to-one mapping, we might as well take the former as the domain of the truth-theory. Then the translational 'semantics' can be thought of as the complementary enterprise of showing how phonetic forms are transformed (translated) into a representational format usable by the utterance comprehension system. The computational processes of this system have to integrate information from a range of sources – language, immediate perception, memory – so they must all be in, or be translatable into, a single common language that these processes can operate over.

However, given that natural language sentences (understood as expression *types*) quite generally underdetermine propositional thoughts, there won't be any wholesale straightforward inheritance by *sentences* of natural language of a truth-based semantics of thought. Recall, for instance, the example 'The book is green', discussed in the previous section; this sentence can express myriad different propositions in different contexts and, despite employing the conditional T-schema, it seemed impossible to give it a context-free truth-conditional statement, so that inheritance by this English sentence of a truth-conditional semantics from that assigned to Mentalese sentences looks highly unlikely. Perhaps, then, the 'semantics by inheritance' idea works just at the *word* level, at least for those words that encode concepts (as opposed to inferential procedures), like 'cat', 'eat', 'clever'. After some detailed discussion of particular cases in the intervening chapters, I will suggest in the final chapter (section 5.4) that even this quite modest claim for the truth-based properties of natural language is doubtful. Word meaning may be of such an abstract and schematic nature that even at this most basic level of linguistic expressions there is no direct inheritance of the referential semantics of components of thoughts, that is, concepts.

I'll move on now to say a little more about the sort of linguistic semantics (internalist, translational) that a cognitive account of utterance interpretation such as relevance theory requires. Here I follow Sperber and Wilson (1986a) and Wilson and Sperber (1993a). Take a simple sentence: 'She hasn't called.' On any normal occasion of use, this will be understood as expressing a complete proposition in which it is predicated of a particular female that it is not the case that she has called (in some specific sense of 'call') some other particular person within some relevantly delimited time span up to the time of utterance. However, the sentence form itself encodes something much less specific, a non-propositional (non-truth-evaluable) logico-conceptual structure, an 'assumption schema', which functions as a template for the construction of fully propositional (truth-evaluable) logico-conceptual structures. It is this schematic logical form that the initial (purely linguistic) phase of understanding delivers and which is the input to the pragmatic processes aimed at constructing the propositional form intended by the speaker, or one similar enough to it to have the intended effects.

Exactly how to represent the encoded logical form (or assumption schema) remains an open question with a number of subsidiary issues to be resolved, includ-

ing how to represent the encoded meaning of indexicals like ‘she’ and what the syntactic structure of a logical form looks like.²⁶ Here is one, undoubtedly wrong, possibility, which might at least give the flavour of the idea (upper case indicating concepts):

(46) NOT [t_i {past} [CALL₁ (X{singular, female})]]

Much explanation is in order: (a) ‘CALL₁’ is one of the conceptual encodings of the ambiguous verb form ‘call’; strictly speaking, then, the surface sentence form ‘she hasn’t called’ may map on to several logical forms depending on how many lexical items of the form ‘call’ there are; (b) there may be an unarticulated ‘object-of-the-calling’ constituent here as well, which would be recovered pragmatically; this may be marked in the logical form as a further open slot, Y, in the second argument position if the concept ‘CALL₁’; (c) although the scope of the negation operator is shown as maximally wide, it might be that all scope specifications are left to pragmatics and the operator is scope-neutral in logical form (these possibilities are discussed in chapter 4); (d) the somewhat mysterious ‘{singular, female}’ is intended to indicate that the encoded linguistic meaning, or character, of the pronoun ‘she’ does not enter into the logical form as a conceptual constituent but is instead a procedural indication to the pragmatic processor of the sort of entity being referred to; (e) the same goes for ‘{past}’, which is a constraint, contributed by the tense marking on the verb. Once the hearer has accessed the referents involved (a concept of an individual and a temporal-span concept), those concepts fill the slots marked in the logical form by X and t_i , and the procedural features disappear, having served their purpose. As argued by Wilson and Sperber (1993a), this distinction between procedural and conceptual encodings is a cognitive processing correlate of the character/content distinction made by Kaplan (1977/89a) (though there are also some differences, which I won’t go into here). I hope this example gives at least some idea of the schematic nature of the semantic representation envisaged as the result of linguistic encoding. This is not a truth-conditional entity. It is fully propositional conceptual representations, rather than sentences, or even utterances of sentences, that are the primary bearers of truth conditions. Intuitions about the truth conditions of utterances are intuitions about the truth conditions of the proposition(s) the speaker intended to express.

The objection that decoding or translation into Mentalese is not semantics is widely made, not only by natural-language semanticists, but also from within pragmatics. Levinson (1988, 1989) contends that the position of relevance-theorists (and of those generative grammarians who investigate a syntactically determined level of logical form (LF), and Katz and Jackendoff, whose ‘semantic’ representations are similar in the crucial syntactic respect) is one of ‘semantic retreat’ and constitutes ‘throwing in the sponge’ (Levinson 1988: 59). He starts with a (perfectly accurate) characterization of the relevance-theoretic conception of linguistic semantics as follows: it consists of the algorithmic extraction of a semantic representation from a syntactic representation; the result of this is an extremely impoverished level of representation with scope of operators undecided, metavariables for pronouns, etc. His concern seems to be that, if this is how semantics is construed, most of the well-established aims of a semantic theory cannot be realized: specifi-

cally, traditional sense relations, such as entailment, synonymy, contradiction cannot be captured at this level of semantic representation (or logical form) and truth conditions cannot be assigned.²⁷ This is right, but merely indicates that if a statement of sense properties and relations is considered desirable, it will have to be formulated over some other level (conceptual representations, perhaps). At most, there might be a case for a terminological change here. The term ‘semantics’ could be reserved for the account of the relation between fully propositional forms and the states of affairs they represent (‘real’ semantics). The output of decoding, which is the input to pragmatic inference, could be called something else: logical form, the level of conceptual interface, the semantically relevant level of syntax, a linguistically determined partial Mentalese representation, or whatever. Once we know what we are talking about this is a matter of little interest; as is sometimes said of this sort of wrangle, it is merely a matter of semantics (in yet another use of the term).

I would like to consider a final objection to the relevance-theoretic stance on linguistic semantics, specifically to its conception of the logical form of sentences as non-propositional. According to this objection, it is simply false that the vast array of natural language sentences do not encode propositions. This is a disagreement from within the overall project, since it is held by people who are generally sympathetic with the decoding (translational) approach to linguistic semantics, and to the role it is assumed to play within the wider cognitive account of utterance comprehension. Someone holding this propositional view (let’s call him Leon)²⁸ grants that a sentence seldom, if ever, encodes, or fully determines, the very proposition that a speaker expresses by uttering it on a particular occasion, and that its logical form functions as a template for the pragmatic construction of the more specific proposition expressed by the speaker. As Leon points out, quite rightly, it does not follow from this that the sentence does not encode any proposition at all. He then goes on to argue that, in fact, sentences do encode propositions and that a proposition (or assumption) schema must itself be propositional because a genuinely inferential process must proceed from one propositional form to another. It is over these latter claims that we differ.

Leon’s propositional view comes in two versions. According to the first version, the indexical sentences in (47) encode generic propositions, which are (inevitably, roughly) represented by the natural-language sentences in (48):

- (47) a. She carried it in her hand.
 b. Paul’s book is there.
 c. It’s raining.
- (48) a. Some female entity carried at some past time something in some female entity’s hand.
 b. Some unique book that is in some relation to somebody called Paul is somewhere.
 c. Raining occurs somewhere at some time.

The procedure which takes us from a sentence in (47) to the corresponding one in (48) is simple enough: wherever you find an indexical you put in a phrase which

spells out the encoded constraint (the character of the indexical) in conceptual terms and wherever you spot an unarticulated constituent you use an appropriate member of the family of ‘some’-indefinites (*something, somewhere, sometime*, etc.) to make it visible. The first problem with this is the assumption that the pairs are truth-conditionally synonymous; in the case of (47a) and (48a), this seems clearly false. Any of the well-developed truth-conditional theories around, such as that of Higginbotham discussed above, would certainly assign quite different T-sentences to these. So (47a) would receive a conditional T-sentence along the lines of (49a), omitting several details, such as those to do with tense, and the likely anaphoric nature of ‘her’, while (48a) would receive a simple (non-conditional) T-sentence since it has no referential elements:

- (49) a. $(u)(x)(y)(z)$ [If u is an utterance of ‘she held it in her hand’, and the speaker of u refers to x with her use of ‘she’, and to y with her use of ‘it’ and to z with her use of ‘her’ and, x is female and z is female, then (u is true iff x held y in z ’s hand)]
 b. ‘Some female entity carried something in some female entity’s hand’ is true iff $\exists x \exists y \exists z$ (x held y in z ’s hand)

The statement in (49a) makes it clear that no truth condition (hence no proposition) is forthcoming until contextual values are supplied to the variables x , y and z ; the statement in (49b), on the other hand, makes it clear that no such fixing of referential values is required before the sentence/utterance can be assigned a determinate truth condition.

Second, each of the propositional representations in (48) cannot but be true (I don’t mean that they are necessary truths, but given the way the world is they are bound to be true). Take any book you like. It will be in some relation or other with some person by the name of Paul in the world (in fact, it will be in *some* relation with *every* person by the name of Paul); for instance, the book may be in the relation of being 62.34 miles north-north-east of my friend Paul Johnson, and it will, inevitably, be located somewhere. Since (48b) is allegedly encoded by (or synonymous with) (47b), we would expect (47b) to have the same property, that is, to be inevitably true. So, let me ask you, is the sentence ‘Paul’s book is there’ true (or is it false)? I’m afraid you may not seek clarification as to who is meant by Paul, or which book, or place, is being referred to. What’s under discussion here is a *sentence*, not a use of the sentence. The answer to the question surely is that the sentence is neither true nor false, that it is not the sort of thing that can be true or false, it is not truth-evaluable, it doesn’t encode a proposition, not even a very general one.

The second version of Leon’s view claims that, while the sentences in (47a)–(47c) may not encode the respective propositions represented in (48a)–(48c), they do entail them. For instance, the semantic representation of the sentence ‘she carried it in her hand’ entails the propositional form ‘some female individual carried at some past time something in some female individual’s hand’, and, since entailment is a relation between propositional forms, it follows that the sentence itself must encode a proposition. This line of reasoning has a plausible look to it (the sentences/

propositions in (48) do seem to follow from the sentences in (47)), although note that it is now not at all clear just what proposition it is that is encoded by the indexical sentence.

In fact, much depends here on how one understands the logical relation of ‘entailment’. One line of resistance would be to deny that the entailment relation holds only between propositional entities; I won’t try that here (but see Sperber and Wilson 1986a/95b: 73, on the logical properties of non-propositional logical forms). Another is to deny that, contrary to appearances, there really is an entailment relation here. I don’t believe that the relation here is one of entailment as that notion is standardly understood, which is as a relation between two truth-evaluable (propositional) entities such that from the truth of the first the truth of the second inevitably (necessarily) follows. First, as argued above, the indexical sentence (abstracted from any occasion of utterance) is not truth-evaluable (hence is not propositional). Second, the undoubted inference to the ‘some female entity carried . . .’ proposition can be explained as arising from consideration of the incomplete conceptual representation together with the referential constraints encoded by the indexicals. That is, on the basis of knowing the meaning of all the parts of the sentence in (47a), we are able to infer that, for any given *utterance* of the sentence, if/when the required contextual/pragmatic work of finding referents for the indexicals has been done, the resulting proposition will entail (48a). If this is right, the second version of the propositional view does not succeed in establishing that indexical sentences (or sentences whose use involves the recovery of an unarticulated constituent) encode propositions.^{29,30}

A final consideration here, that crosscuts both versions of Leon’s propositional view, concerns utterances such as ‘Tom’s father’, or ‘on the top shelf’. These phrasal utterances encode a semantic representation (a logical form), which provides the pragmatic processor with a template for the construction of the proposition expressed by the utterance. On some occurrences, these will be ellipsed versions of complete sentences (for instance, when they are produced as answers to questions, such as ‘Who’s the guy by the door?’, etc.), but they need not be (see Stainton 1994, 1997; Elugardo and Stainton forthcoming b; and the discussion in the next chapter – section 2.3.5). There is a genuine and rather commonplace phenomenon of non-elliptical subsentential utterances, and what they encode simply cannot be propositional. Leon could respond to this that he is only concerned with sentential utterances. But recall that one of the main factors motivating his propositional view is his assumption that all inferential processes must operate over fully propositional forms. From this, it follows that the input to the (non-demonstrative) inferential process of developing the logical form of an utterance into that proposition which the speaker actually expressed must itself be propositional. But this consequence of the assumption, and so the assumption itself, is utterly undermined by the subsentential cases, which show that the pragmatic processor *can* take as input a non-propositional logical form and enrich it into a fully propositional form. There is no compelling reason to suppose that the story goes differently for the sentential cases. I conclude that the claim that sentences encode propositions has not been upheld and that there is not even any particular reason to view propositionality as a desirable property for sentences to have.

Summing up this section, the first stage of interpreting an utterance involves decoding the linguistic expression employed. The result of this is a (generally non-propositional) logical form which is in a format usable by the pragmatic processor. This logical form is a syntactically structured string of concepts with some slots, or free variables, indicating where certain contextual values (in the form of concepts) must be supplied; the nature of the contextual value may be constrained by procedural information (as in the case of certain indexicals). So, on this sort of account, linguistic semantics specifies mappings (translations) between lexical forms and concepts or procedures, and between surface syntactic structures and their underlying logical structure.

1.6 Radical Underdeterminacy and the Background

Perhaps what is inexpressible (what I find mysterious and am not able to express) is the background against which whatever I could express has its meaning.

(Wittgenstein 1931/80: 16e)

1.6.1 The Background

John Searle (1978, 1980, 1983, 1992, 1996) is a strong advocate of the underdeterminacy thesis:

The literal meaning of a sentence only determines a set of truth conditions given a set of background practices and assumptions. Relative to one set of practices and assumptions, a sentence may determine one set of truth conditions; relative to another set of practices and assumptions, another set; and if some sets of assumptions and practices are given, the literal meaning of a sentence may not determine a definite set of truth conditions at all. (Searle 1980: 227)

Searle believes that existing discussions barely scratch the surface of the phenomenon of the linguistic underdeterminacy of propositions expressed (see Searle 1992: 181; 1996: 131). In his view, the issue that needs highlighting is that of the Background (with a capital 'B'); this is fundamental to meaning and understanding, there is no literal meaning without it, there is nothing truth-evaluable without it. The force of his position is best appreciated by considering some of the examples he discusses, although it should be borne in mind that the kind of underdeterminacy that is a result of the Background is not a characteristic of just some groups of linguistic expressions, but is a property of linguistic meaning quite generally.

He shows how the literal use of common verbs, like 'cut', 'open', 'close', 'clean', 'mend', 'play', 'bring', 'take' and 'make', determines different truth conditions as a result of different relevant parts of the Background coming into play. For instance, let's take several substitution instances of 'X opened Y', restricting ourselves to human agents, concrete objects and a strictly literal use of 'open':

- (50) a. Jane opened the window.
 b. Bill opened his mouth.
 c. Sally opened her book to page 56.
 d. Mike opened his briefcase.
 e. Pat opened the curtains.
 f. The child opened the package.
 g. The carpenter opened the wall.
 h. The surgeon opened the wound.

(adapted from Searle 1983: 145)

Though the semantics (the encoded meaning) of the word ‘open’ is the same in these examples, it is understood differently in each case; the contribution it makes to the truth conditions of quite literal utterances varies with the sentential context it occurs in. What constitutes opening a book is very different from what constitutes opening one’s mouth, which is quite different again from what constitutes opening a package, etc. Importantly, although it looks as if it is the meaning of the expressions we substitute for ‘X’ and ‘Y’, particularly ‘Y’, that determines the interpretation of ‘open’, this is only so given a Background of assumptions concerning what is involved in an X opening a Y. That is, we could imagine a situation in which, say, a mouth had been stitched closed for some reason (to prevent some disease, or to stop an obese person from eating, etc.), and then the process of opening would be quite different from the one we assume given our standard assumptions about people and their mouths.

Furthermore, given our current Background we are unable to understand (assign truth-conditional content to) the following:

- (51) a. Bob opened the grass.
 b. Chris opened the fork.
 c. Jane opened a hair.

This is because, although we understand each word in these sentences and their syntax is unproblematic, our Background does not supply us with any know-how concerning opening grass or forks or hairs.

Searle also considers some cases which are among the most likely counterexamples to the generality of the thesis of the Background, that is, plausible candidates for eternal sentencehood, such as the following:

- (52) a. Four plus five is nine.
 b. Snow is white.

The claim is that even these only determine a set of truth conditions against a Background of practices and assumptions, which are so deeply entrenched that we do not notice them. If arithmetic practices changed, in the one case, or if some fundamental alteration to the course of nature took place and our Background assumptions with it, in the other, these sentences might determine different truth conditions from those they currently do, or perhaps determine none at all (Searle 1980: 229–31).

Before locating the Background within Searle's wider account of intentionality, I want to mention two points which arise within the present narrower context of utterance interpretation. First, it seems that for Searle there is no distinction between utterance meaning and sentence meaning when the speaker means what she says; he takes it that the speaker's knowing of literal meaning that it applies only against a particular Background is part of her linguistic (semantic) competence. So there is no sharp distinction between a speaker's semantic competence and her (background) knowledge of the world (Searle 1978/79: 134). This may ultimately prove to be the case (it has some formidable supporters (see note 7 of this chapter)), but I don't think we have to accept it as a foregone conclusion on the basis of the facts about interpretation just considered. I find an equivocation in the way Searle talks of 'literal meaning'. On the one hand, he says 'the *same literal meaning* will determine different truth conditions given different Backgrounds' (Searle 1983: 145), and, on the other, he talks of the dependence of literal (sentence) meaning on context/background (Searle 1978/79: 134–5). That is, on the one hand, 'open' has the same literal meaning in each of the examples above, and, on the other, its literal meaning depends on context and so is different in each case. The first way of talking of 'literal meaning' leaves open the possibility that it can be given its own Background-free characterization, though it cannot be applied (used, understood) except in relation to a body of Background assumptions/practices. In short, a principled semantic/pragmatic distinction is not ruled out by accepting, as I do, the fundamental and pervasive role of the Background.

Second, for anyone interested in an account of utterance interpretation, that is, an account of the hearer's interpretive processes and the representations he recovers, there is a pressing question here. I have suggested above that hearers recover unarticulated constituents of the proposition expressed by an utterance. So the question is: how much of this great mass of contextual/background material goes into the proposition expressed? How do we distinguish between what the hearer must infer and build into his representation of the speaker's informative intention and what is left in the Background?

Recanati (1993: 260) quotes one of Searle's more grotesque scenarios in which he discusses the example 'I have had breakfast' from Sperber and Wilson (1986a/95b: 189–90). In addition to the narrowing down of the temporal specification (in most instances, to the day of utterance), which they give as an example of a pragmatic contribution to the proposition expressed, Searle points out that 'having breakfast' is interpreted as '*eating*' breakfast and 'eating' breakfast is taken to mean putting it in one's mouth, chewing it, swallowing it, etc., as opposed to stuffing it into one's ears or digesting it through the soles of one's feet, though none of this is encoded in the phrase itself. Recanati agrees with Searle about this and concludes from it 'that "what is said" – the situation our utterance intends to describe – necessarily involves unarticulated constituents. No proposition could be expressed without *some* unarticulated constituents being contextually provided' (Recanati 1993: 260). This statement is surely right, but it is a rather lame conclusion in this context, since it runs together the deepest taken-for-granted unrepresented aspects of meaning (e.g. what's involved in eating) and other much less

general, context-particular, features of understanding that a hearer will have to infer and represent if he is to recover the speaker's intended message. I will return briefly to this question in chapter 3 (section 3.7.1), in the context of a discussion of the pragmatics of 'and'-conjunctions, where I consider whether we can distinguish which among the many unarticulated aspects of meaning must be recovered and mentally represented.

Searle's thesis of the Background should be situated within his overall view of the mind, which I will now indicate, in a very sketchy but, I hope, accurate way. We have intentional (representational) states, such as beliefs, desires, intentions; these are real, they are properties of human biology; that is, they are not convenient fictions and they are not reducible to something else, although they are, of course, caused by neurological processes in the brain; linguistic meaning is (one instantiation of) derived intentionality, that is, it is grounded in the more basic intentionality which is an intrinsic property of the mind/brain; consciousness (like intentionality) is an intrinsic and ineliminable feature of the human mind/brain; some, but not all, conscious states are intentional in that they represent something beyond themselves; some, but not all, of one's intentional states are conscious at any given moment, but all are capable, in principle, of being brought to consciousness; without consciousness there is no intentionality (in fact, for a state to count as mental it must be potentially conscious);³¹ the workings of the (holistic) system of intentional states are wholly dependent on a massive set of capabilities, dispositions, know-how, that are not themselves intentional, that is, the Background.

Searle makes a distinction between features of the Background that are common to all human beings and features that are culture-specific. So, for instance, the basic know-how involved in walking and eating are aspects of the 'deep Background' while specific aspects of, say, appropriate conduct at meals, or the sitting and standing conventions at public gatherings would be elements of the 'local Background'. He makes a further distinction between knowing *how to do* things and knowing *how things are*; for example, we know how to walk and one aspect of our knowing how things are with the world is our taking it for granted that the ground won't shift around beneath our feet.

Searle admits that there is some obscurity in the concept of the Background, but he finds its existence an inescapable matter of fact and has a range of arguments for it, some of which, pertaining specifically to language understanding, have been alluded to above. I do not doubt the existence of some such body of capacities and assumptions, but I would like to try to think about it in a way that will mesh better with the relevance-theoretic approach to utterance understanding that I am employing. It is crucial to Searle that the Background set of assumptions/practices is not itself 'intentional', since it is what makes intentionality (the 'aboutness' of our beliefs and of the meanings of our utterances, etc.) possible; it is that without which intentional states would be indeterminate. The Background is, however, not to be construed as actual objects or states of affairs in the world; it is mental ('in the head' – Searle 1991: 291). I suggest that the concept of 'manifestness' as characterized by Sperber and Wilson (1986a/95b: 38–42) is helpful here: an assumption is manifest to an individual at a given time if and only if he or she is capable at that time of

representing it mentally and accepting its representation as true or probably true. They then define the concept of an individual's 'cognitive environment' as the set of assumptions that are manifest to him or her at a particular moment. Thinking of the Background as a set of (merely) manifest assumptions seems to answer to Searle's broad requirements that the Background does not consist of states of affairs in the world, on the one hand, nor of intentional states (representing the world), on the other, but is mental insofar as it is dependent on – would not exist without – the mind/brain.

There are several features of the concepts of manifestness and cognitive environment that are relevant in applying them to the Background: (a) manifestness is a matter of degree, and the degree of manifestness of a given assumption in an individual's cognitive environment may shift from moment to moment depending on features of the external physical environment and on his/her internal cognitive states (for instance, where attention is focused); (b) some assumptions are such that once they are manifest to an individual they remain so thereafter, as stable elements of his or her cognitive environment, and others are temporary and may be very short-lived, a function of where one happens to be, who one happens to be with, etc., at a particular moment; (c) an individual's cognitive environment overlaps to a greater or lesser extent with every other individual's cognitive environment: assumptions that are common to human existence (e.g. 'objects are solid and permanent', 'the earth does not recede beneath one's feet', etc.) are part of everyone's cognitive environment; other assumptions are shared by largish subsets of cognitive environments, such as those pertaining to practices and conventions taken for granted in a particular culture; close friends share a huge further range derived from their shared experiences, verbal and non-verbal. Some shared manifest assumptions are in fact mutually manifest, that is, it is manifest to the sharers that they share those assumptions and with whom they share them.

We might usefully think of the Background as a set of assumptions and practices that maintain a fairly steady degree of not very high manifestness, across time, in an individual's cognitive environment. A subset of the Background consists in assumptions/practices which make up the mutual cognitive environment of all (non-pathological) human beings – the deep Background; other subsets are the mutual cognitive environments of what can be loosely termed culturally defined groups of human beings – local Backgrounds). Most, perhaps all, Background assumptions and practices figure in some mutual cognitive environment or other. Some of these assumptions might be occasionally actually represented by the individual when he or she is confronted with 'strange' situations, situations with features which contradict some feature of the Background and so make that feature highly manifest to him or her, situations of shock. For the most part, though, the Background keeps its place, in the background, unrepresented, an essential foundation for thinking and understanding and, though not discussed here, action.³²

This way of thinking of the Background may or may not do full justice to Searle's conception; I suspect that I have not succeeded in fully capturing its bedrock nature. Some of the capacities, the know-how, the *savoir-faire*, that Searle refers to, do not seem appropriately thought of as sets of *assumptions*. However, I do not see any glaring problem in extending the concept of manifestness to a broader array of

dispositional sorts of structures, like procedures, action schemas and processing schemas; these seem better suited than assumptions to be the vehicles of such aspects of the Background as how to eat, how to walk, what's involved in behaving normally in a shop/classroom/someone else's house/at home, etc. Constituents of these schemas could, in principle at least, be represented by an individual and held as true, for instance, 'in order to eat x , one places x in one's mouth'. Anyway, given the rather limited use to which I will be putting the Background in what follows, I think a characterization in terms of manifest assumptions and procedures will suffice.

1.6.2 Radical underdeterminacy and 'expressibility'

Let's return briefly to the issue of linguistic (in)effability. Recall that Searle was mentioned, along with Katz and Frege, as one of those espousing a strong principle of effability, though his is labelled the principle of 'expressibility': 'whatever can be meant can be said' (Searle 1969: 19–21). It might seem that anyone holding as strong an underdeterminacy thesis as is entailed by the thesis of the Background must be a supporter of the anti-effability view I took in the previous section. So is Searle being inconsistent? Did his development of the idea of the Background put paid to his earlier advocacy of Expressibility? The answer to the second question is certainly 'no':

there is nothing in the thesis of the relativity of literal meaning [to the Background] which is inconsistent with the Principle of Expressibility, the principle that whatever can be meant can be said. It is not part of, nor a consequence of, my argument for the relativity of literal meaning [to the Background] that there are meanings that are inherently inexpressible. (Searle 1978/79: 134)

This emphatic denial is all we get from Searle on the issue. One way of trying to convince oneself that the two theses are consistent would be to reason as follows. The principle Searle endorses must be quite different from that of Katz and Frege who claim that, for each indexical sentence paired with a context, natural language affords an eternal sentence, which fully encodes the proposition expressed by the first. Searle's principle must be the innocuous truth that the human interpretive capacities are set up so that for any propositional content that a person might want to communicate, there is some linguistic expression or other (perhaps several) that she could use to express that propositional content in a given context. An essential part of those interpretive capacities is, of course, the Background; another is a pragmatic inferential system. This position would be equivalent to the unobjectionable first principle of effability, given in section 1.3.2. However, a look at the actual formulation of the Principle of Expressibility shows that this is not the way out:

For any meaning X and any speaker S whenever S means (intends to convey, wishes to communicate in an utterance, etc.) X then it is possible that there is some expression E such that E is an exact expression of or formulation of X . Symbolically: $(S)(X)(S \text{ means } X \rightarrow P(\exists E)(E \text{ is an expression of } X))$. (Searle 1969: 20)

Clearly, by 'said' he intends 'encoded' (that is, 'an exact expression of or formulation of meaning X'), making this principle more like the second principle of effability discussed earlier. It seems, then, that the explanation for the apparent coexistence of this principle with the Background must lie with the concept of 'what is meant'. Just as what we say is relative to the Background, so is what we mean, what we intend to convey, what we think, etc. Essential though they are to interpretation, I take it that elements of the Background do not fall within the scope of our communicative intentions (they cannot, since they are non-intentional). Then the fact that they do not enter into what is said (= 'encoded' here) does not threaten Expressibility. The non-representational Background underlies both the said and the meant, and so is not the source of any disparity between them. Thus there is no inconsistency in maintaining both the Principle of Expressibility and the thesis of the Background. This does not, of course, touch the arguments in section 1.3, on the basis of which Searle's Principle of Expressibility, like Katz's Principle of Effability, must be wrong.³³

1.6.3 Radical underdeterminacy and semantic compositionality

Having considered the Isomorphism Principle briefly and the Effability Principle lengthily, I come now to a third principle, perhaps the most cherished of all by philosophers of language. This is the **Principle of Semantic Compositionality**: 'The semantic value (meaning) of an expression is determined by the semantic value (meaning) of its constituents and the manner in which they are combined.' This is usually attributed to Frege but is maintained, in variant forms, by most semanticists. The principle is taken to hold for every type or level of semantics, so, for Frege, who distinguished between sense/content and reference/extension, the principle holds for both these types of entity,³⁴ and, for those who make further distinctions between linguistic meaning or character, and propositional (truth-conditional) content, it holds also at this third (or fourth) level: the linguistic meaning encoded by an expression is determined by the encoded meaning of its constituent parts and their manner of combination.

Compositionality is generally deemed essential in explaining the learnability, productivity and systematicity of the human language capacity, that is, in explaining how it is that the finite human mind is able to produce and understand indefinitely many novel sentences, which is itself a factor in accounting for the unbounded expressive and communicative power of human utterances (see, for example, Fodor 1987a; Fodor and Pylyshyn 1988, section 3; Butler 1995a). But the Principle of Semantic Compositionality imposes a stronger constraint than is entailed by these considerations; it embodies a determinacy requirement: the semantic value of a complex expression E must be 'completely determined by the constituent expressions $e_1 \dots e_i \dots e_n$ of which it is composed. That is, any semantic property with which the complex expression is endowed, must be traceable to one of the constituent elements, or to the construction itself' (Welsh 1986: 553–4).

As Searle (1980) points out, given the role of the Background, the Compositionality Principle (CP) does not hold for the view that the meaning of a sentence is to be given as a specification of its truth conditions. The truth conditions of a completely literal use of a sentence, such as ‘I have had breakfast’ or the various examples in (50) above, involving the verb ‘open’, seem not to be fully determined by the contribution of their constituent parts and their mode of combination, but to be profoundly affected by a body of Background assumptions that inevitably permeates meaning/understanding.

Some of the other sources of linguistic underdeterminacy, discussed in sections 1.2 and 1.3, also make trouble for the CP if the meaning of a sentence is taken to be the proposition it expresses:

- (53) a. It’s raining.
 b. The book is green.
 c. He handed her the key and she opened the door.
 d. She insulted him and he left the room.

In understanding each of these, we seem to supply a component of meaning that is not determined by any of the lexical constituents or introduced by the syntactic composition process: a location component for (53a), a salient-part component for (53b), an instrument component for (53c) (she opened the door *with the key he gave her*), and a cause-consequence component for (53d).

Pelletier (1994) discusses another set of cases, all of which involve an ambiguity that does not seem to be traceable to any syntactic or lexical feature of the sentence:

- (54) a. Every linguist knows two languages.
 b. The philosophers lifted the piano.
 c. When Alice rode a bicycle, she went to school.

Each of these sentences can be understood in two ways (that is, has two distinct propositional meanings, hence two distinct semantic values): in (54a), the two readings involve different scope relations between the two quantifiers; in (54b), there is a distributive/collective ambiguity concerning whether it was all of the philosophers together who lifted the piano or whether they each did it separately; in (54c), there is an understanding on which each of the instances of bicycle riding is an instance of going to school, and an understanding on which Alice’s bike-riding days were back when she was a schoolgirl. I agree with Pelletier that it is not very plausible that any of these originates from a lexical ambiguity or from two distinct *syntactic* analyses; rather, there are some additional factors (of a pragmatic inferential sort) which interact with the single linguistic meaning to give the two distinct understandings for each case.

Jackendoff (1997: 51–67) discusses a range of examples that he considers problematic for any ‘syntactically transparent’ notion of semantic composition, including the following:

- (55) a. Mary finished the novel.
 b. Mary finished the beer.
 c. The ham sandwich in the corner wants some coffee.
 d. I'm parked out back.

The truth-conditional content of (55a) and (55b) seems to involve the finishing of quite different processes in the two cases, content that arises from general, rather than linguistic, knowledge: one reads (or writes) a novel, one drinks a beer. The next two examples are cases of 'reference transfer': the subject of (55c) is understood as a customer who has ordered, or who is eating, a ham sandwich, and the subject of (55d) is understood as the speaker's car. Jackendoff points out that the inferred referent in these cases can enter into grammatical binding relations, so that it seems impossible to argue that the result of the pragmatic transfer does not fall within the compositional process.

Each of the three authors has a different response to these problems (and others that they discuss) for the Principle of Semantic Compositionality. Pelletier opts for the abandonment of the principle. Jackendoff accepts that 'one cannot "do semantic composition first and pragmatics later"' and moves in the direction of developing a concept of 'enriched compositionality'; this allows the compositional process to incorporate a measure of pragmatic inference together with the linguistically encoded meaning of the constituent parts, in order to arrive at the meaning of sentences such as those in (55) (see Jackendoff 1997: 49–55). Searle's (1980) solution to the relativity of truth conditions to the Background is to drop the assumption that the literal meaning of a sentence is a set of truth conditions, a move that has been advocated throughout this chapter.

However, neither the examples just cited, nor Searle's points about the Background, affect the validity of the CP as applied to linguistic semantics construed as an account of encoded meaning (the translational, as opposed to truth-conditional, approach). As discussed in section 1.5, on this approach, the semantic representation of a sentence is an incomplete conceptual structure, with some empty slots marking the need for further conceptual material, that is, it is a schema for proposition construction. The arrangement of conceptual constituents and slots reflects the relations encoded by the syntax (the combinatorics) of the natural-language expression used, or in some instances, such as the scope of negation and quantifiers perhaps, leaves those relations to be fixed pragmatically. That the CP holds at this level is perhaps trivially true. The 'semantic' representation (logical form) of a sentence cannot but conform to the CP, since it just is a product of the linguistic semantics of each word/morpheme making it up and the semantic relations imposed by the syntactic structure of the sentence, and of nothing else. That simply follows from the modular view of linguistic processing. It should be noted, in passing, that this 'semantic' entity may never be mentally represented as such by a hearer processing an utterance. By the time the last word of an utterance is processed, the earlier parts of the string are already buried within the pragmatically supplied flesh of a fully propositional thought. As the utterance is processed millisecond by millisecond (left to right, as it were), pragmatic processes come into play; indexical references are resolved, disambiguations performed, unarticulated constituents supplied and

decoded concepts enriched, as soon as they can be. Standardly, that is before the whole utterance is heard and, certainly, before the decoding processes have delivered the whole sentential semantic representation.

Paul Horwich, who has written extensively on semantic compositionality (see Horwich 1997, forthcoming), says the following:

the content of SEM [= the meaning of a complex linguistic expression] is determined by the contents of its parts . . . but this can be explained trivially – in a way that has nothing to do with truth conditions. It suffices to suppose that *understanding* a complex expression (i.e. implicitly knowing its meaning) is, by definition, nothing over and above understanding its parts and appreciating how they have been combined. If this is so, then the property in virtue of which a sentence possesses its particular content is simply the property of its being constructed in a certain manner from primitives with certain meanings. (Horwich forthcoming: 20–1)

Although Horwich's own approach to linguistic meaning is quite different from those discussed in this chapter (it is in terms of conceptual roles), this statement reflects rather closely the way in which compositionality applies to the translational approach to linguistic semantics.

It could be objected at this point that, while the CP clearly does hold of sentence meanings (logical forms), as construed here, if this is all there is to be said for it, the interest of the CP is greatly reduced. The linguistically determined logical form is, some might argue, essentially pre-semantic; it is not communicated, but is merely a vehicle for what is communicated, it is not knowingly 'grasped' by addressees, it is not phenomenologically salient. What we are really interested in is the propositional thought expressed, perhaps communicated, by an uttered sentence, and this is what the CP has been generally supposed to apply to. I have some sympathy with this dissatisfaction. The conclusion reached above is a bit of a comedown for the CP, although its applicability just to abstract encoded linguistic meaning is probably quite sufficient to answer to the observations concerning the productivity and systematicity of language. There just is no escaping the fact that the propositions that may be expressed by sentences in use are a function, not only of linguistic meaning, but also of pragmatic inference. Perhaps this marks the demise of an *interesting* principle of semantic compositionality, or perhaps it points to the possible development of a different sort of compositionality principle, one that can accommodate an interaction of decoded and pragmatically inferred meaning in the determination of the proposition expressed (a principle of semantic/pragmatic compositionality).³⁵

Finally, Searle's observations raise the further question of whether a truth-conditional semantics of sentences in *Mentalese* conforms with the CP. Assuming, as I do, that the question is meaningful (that there are such things as structured sentences of *Mentalese*, and that it makes sense to assign them truth conditions), it seems that the thesis of the Background must lead to a negative answer here too. The syntactic (compositional) nature of sentences in the language of thought might explain their productivity and systematicity, but the truth conditions of a thought are dependent, not only on the referential properties of basic constituents (concepts)

and their compositional relations to one another, but also on the Background. Whether they are also subject to (some of) the other sources of truth-conditional underdeterminacy which are typical of natural language sentences is briefly considered in the next section.

1.7 Underdeterminacy of Thought?

I follow Jerry Fodor's view that having a thought with a particular content *P* involves the occurrence (the mental 'tokening', as it is often put) of a sentence of the language of thought (Mentalese) which means that *P*. I take it that Mentalese is distinct from any particular natural language, that it has a compositional syntax and semantics, but no phonology, and its syntax consists of a single level (as opposed to the multiple levels often assumed in syntactic theories of natural language). A Mentalese sentence has truth conditions, that is, there is a specifiable state of affairs which, if it is the case, makes the sentence true. The basic constituents of Mentalese sentences are concepts, such as CAT and SMILE, and there are no dummy elements, that is, elements with no semantics (like expletive 'it', or the copula 'is', in some of its manifestations in English). These representations stand in certain causal relations with each other, relations that constitute inference of a sort; for instance, the belief that Hugo is a spaniel causes the belief that Hugo is a dog. See Fodor (1975, 1978, 1987a) for detailed arguments supporting the existence of such a 'language of thought' and Lycan (1990) for succinct characterizations of various versions of the Mentalese view.

In the previous sections of this chapter, much has been made of the context-sensitivity of linguistic utterances, that is, of their having content (expressing a proposition) which is not fully encoded, content which is underdetermined by the natural-language sentence employed in the utterance. So the question here, by analogy, is 'Are there aspects of the (truth-conditional) content of a thought that are not encoded in the Mentalese representation which occurs in having the thought; that is, are Mentalese sentences context-sensitive?' I'll do no more than skim the surface of this large (and, probably, insufficiently articulated) question in the next two subsections.

1.7.1 Mentalese, pragmatics and compositional semantics

Mentalese is the medium of thought. Given the computational view of mental states and processes, it follows that, for every feature of content that a thought process is sensitive to, there must be a formal element present in the Mentalese representation of that content. So, for instance, if my train of thought moves from recalling that the Smith family have two cats to the conclusion that some of my neighbours have pets, then the relation between the properties *the Smith family* and *neighbours of mine* and between *having cats* and *having pets* must be reflected in the form (the symbols) of those Mentalese representations which enter into the (computational) thinking process. A consequence of this 'formality condition' is that once the sensory

perception systems have delivered their representations of the external world to central thought processes, the external world has no further influence on the thinking that follows. That is then a matter of what is present in the form of those representations and of other representations within the system that the processes have access to. This is reflected in the well-known assumption in cognitive psychology of a strategy of ‘methodological solipsism’, that is, of proceeding in the study of mental states and processes as if nothing outside the mind existed (see Fodor 1980).

This view seems to indicate an immediate answer to our underdeterminacy question with regard to thought: Mentalese sentences do not underdetermine their truth-conditional content, they are not context-sensitive, they are eternal. Given the formality constraint, Mentalese sentences and phrases must be unambiguous; that is, there are no elements of form which have several distinct semantic values, as do ‘bank’ and ‘visiting relatives’ in English. Otherwise, a ‘bank’ thought, for instance, would be simultaneously a thought about a financial institution and a thought about the side of a river; you could never have one without the other and every ‘bank’ thought would send you off on two wildly divergent inferential tracks. Similar considerations would seem to apply to indexicals. For instance, take the thought of a person who, if she wanted to express it, might utter the English sentence ‘That one is better than that one’, while pointing first at one object, then at another (say, two cars). The English sentence contains identical elements for referring to the two cars, but the Mentalese sentence must contain distinct components. If it did not, that thought would be indistinguishable from a nonsensical one in which a person predicates of an object that it is better than itself. In fact, the issue is wider than this example makes it seem; it doesn’t arise just for utterances involving *multiple* uses of an indexical. A thought about a particular object, say a vase I am looking at, which could be expressed by an utterance of ‘That’s beautiful’, must surely involve a formal component that distinguishes it from those thoughts I have about other objects which could also be expressed by the utterance of a ‘that’ sentence. Otherwise, we could not keep track (to the quite considerable extent that we do) of distinct objects in our environment. The question of what form these object-distinguishing constituents of Mentalese sentences might take is discussed in the next section.

What about the other symptoms of linguistic underdeterminacy, such as subsentential representation, unarticulated constituents, generality of sense, vagueness and polysemy, which are so characteristic of natural-language utterances? Surely these are not possible features of thoughts. After all, the process of understanding *utterances* with these properties is one of, as it were, recovering the missing bits, so as to discern the *thoughts* being expressed. Recall the following sorts of example:

- (56) a. On the top shelf.
 b. Nurofen is better. [than what, and for what purpose?]
 c. It’s raining. [where?]

While (56a) is a subsentential element of English and the other two are fully sentential, their Mentalese counterparts are all subsentential elements of that representational system; that is, they are all semantically incomplete, and so

not truth-evaluable. For any utterance of one of the strings in (56), the thought expressed is complete and so incorporates a constituent whose content is not encoded in the linguistic expression, but is pragmatically derived.

Another sort of case, canvassed earlier, involves the use of a word to communicate a concept which differs from that encoded by the word:

- (57) a. I am happy.
 b. The house was silent.

In the case of (57a), the idea here is that different concepts of ‘happy’ are communicated in different contexts, from a mild sense of acceptance of life to a feeling of joyous delight; supposing the word meaning is general enough for all these different senses to fall within its domain, then any more specific concept has to be pragmatically inferred. The example in (57b) might go the other way, that is, from the absolute concept of noiselessness putatively encoded, looser notions of ‘silent’ might have to be inferred in understanding the utterance (for instance, the inhabitants of the house may be asleep, so there is the sound of their breathing, as well as the various creakings, crackings and drippings of the fabric of the house). The particular concept is a component of the thought the speaker seeks to communicate (a ‘word’ of the Mentalese sentence tokened) and, if communication is successful, of the thought recovered by the addressee. On this view, Mentalese must have a large stock of concepts that are not encoded by any element of natural-language form. This idea is pursued in more detail in chapter 5.

What bridges the gap between the underdetermining encoding of a natural-language utterance and the thought(s) expressed is a powerful pragmatic inferential mechanism, whose job it is to figure out the informative intention behind a linguistic utterance (or any other act of ostension). Having thoughts is a strikingly different kind of mental activity from comprehending utterances. We do not have to undergo a process of *comprehending* occurrences (tokenings) of our thoughts (hence Mentalese sentences), as we must comprehend occurrences of linguistic utterances. Thoughts are, in effect, an end-product: they are the *result* of utterance comprehension processes, or are prompted by the deliverances of our sensory-perceptual systems, or are generated from within the central thought system itself (in ways that remain largely mysterious). A thought token is not ostensive, it is not a communicative act, it doesn’t involve an interaction of speaker and addressee, it doesn’t come with a presumption of optimal relevance, which warrants a particular expectation of effect and effort and so a particular pattern of processing geared to satisfying that expectation. There surely cannot be a ‘pragmatics of thought’, in the sense of pragmatics at issue here.³⁶ Furthermore, as Butler (1995b: 14) says, ‘Since mental representation is presupposed by communication, it cannot, and obviously does not, involve communication itself (on pain of regress).’

In accordance with these observations about the nature of thoughts, and setting aside issues of the Background for the moment, it looks as if a Principle of Semantic Compositionality should hold for Mentalese: the semantic value of a sentence is fully determined by the semantic value of its constituents and their syntactic combination. Assuming that the ‘semantics’ at issue here is truth-conditional, what the

principle says is that the truth conditions of a Mentalese sentence are fully determined by the semantic entities that its constituents refer to and the way in which they are related according to the structure of the sentence. For instance, the thought that it is raining encodes a location constituent, LOC_x , so that it is true iff it is raining at LOC_x (at such and such a time); the thought that Mary is happy encodes a particular concept of happiness, $HAPPY^*$, so that it is true iff the property denoted by $HAPPY^*$ applies to Mary (at such and such a time). It's worth noting differences in the application of this principle to natural languages and to thought: while the semantics of thought is compositional, this is not a semantics that thinkers know (in the way in which truth-conditional semanticists envisage speakers knowing the semantics of their natural language system) and there is no compositional *process* of understanding it as there is in understanding linguistic utterances, because there is no process of understanding it at all. (For relevant discussion of the compositionality of language and of thought, and of the issue of understanding sentences in the two kinds of representational system, see Fodor and Lepore 1999.)

Butler (1995a, 1995b) discusses a range of cases of alleged context-sensitivity of mental representation, which have been taken to threaten the compositionality of thought. The cases in question are extrapolated from the *linguistic* context-sensitivity of examples such as the following:

- (58) a. The apple is red. [on its peel]
 b. The watermelon is red. [in its flesh]
 c. Mary finished the novel. [finished reading or writing]
 d. Mary finished the beer. [finished drinking]

He concludes: 'What we find . . . is *not* a context-sensitivity of mental *semantics*, but a context-sensitivity of the *deployment* of mental representations' (Butler 1995b: 13). What he means by this is that the occurrence of the mental representation that corresponds (roughly) to 'red on most of the peel', or 'red on the inside flesh', in the cases of (58a) and (58b), is sensitive to other elements of the sentential representation of which it is a constituent, that is, whether redness is being predicated of an apple, a watermelon, or something else. He confines his attention to sentence-internal context-sensitivity: 'red' predicated of 'the apple', 'finished' predicated of 'the novel', etc., but this statement about the deployment (occurrence) of a particular conceptual representation applies to a wider notion of context-sensitivity. Although the interpretation he assumes for the examples given in (58a) and (58b) is a kind of default, predicating a colour of either of those (or of any other object) could give rise to a quite different mental representation. Suppose, for instance, that a chemical test for the presence of a particular bacterium in apples turns their flesh red when there is a positive result; then, obviously, in such a context, an occurrence of a thought that a particular apple is red will involve the deployment of a mental representation with a different part constituent from the default case.

The old idea that communication is a matter of a speaker encoding her thoughts and a hearer decoding them has been completely undermined by, on the one hand, the linguistic underdeterminacy facts, and, on the other, the presence in human cognitive architecture of an inferential device dedicated to recognizing the informative

intentions that lie behind ostensive communicative acts. These considerations do not seem to carry over to thoughts; a conception of them as fully encoding their truth-conditional content seems to make very good sense. However, consider the following remark by Kent Bach: ‘Since the contents of thoughts cannot be conceptually incomplete, the conceptual representations that comprise them cannot be semantically underdeterminate *in the way* that sentences can be’ (Bach 1994b: 157; my emphasis). There is an implication here that thoughts can perhaps be semantically underdeterminate *in some way*, even if not in the way that natural language utterances are. What Bach is allowing for here is a kind of indexicality within thought, specifically within what are known as singular or *de re* thoughts (that is, thoughts about particular objects). And, in fact, the philosophical literature abounds with papers on indexical (egocentric or demonstrative) thought. (See, for instance, Castañeda 1966, 1967, 1990; Burge 1977; Perry 1977, 1979, 1993; Kaplan 1977/89a; Stalnaker 1981; McGinn 1982; and Recanati 1990, 1993). This needs some consideration since, if it is right, the view of thoughts as involving occurrences of Mentalese sentences that fully encode their content, and so are eternal, has to be modified. Just as indexicality in natural language ensures the context-dependence of the truth-conditional content of linguistic utterances, so too for thoughts: if they can be indexical, then their truth conditions are context-dependent. I move now to a short discussion of mental indexicals.

1.7.2 Mental indexicals and the mind–world connection

In order to understand someone’s use of a natural-language singular term (such as a name, a pronoun, a demonstrative or a definite description), we have to figure out (by considering the context) what it is he is referring to. Our *thoughts* about objects are, clearly, not like that; there is no such process of ‘understanding’ to be undertaken. Suppose you have a thought which, if you chose to, you might express by uttering the English sentence in (59):

(59) That’s beautiful.

The Mentalese sentence that you token in having the thought does not involve you in figuring out the referent; it is part and parcel of having the thought that you think of the referent and, in accordance with the ‘formality’ constraint, the thought constituent which corresponds to the utterance of the word ‘that’ should be distinct from your mental representations of other objects if your thinking about the world is not to go wrong. So, if there are mental indexicals, they must be different in some fundamental way from natural language indexicals. Let’s consider some views on the nature of the mental symbol that occupies the position in your thought that ‘that’ occupies in your utterance of (59).

An obvious first idea about this is that Mentalese (unlike natural languages) provides uniquely denoting descriptions and that, in thinking the thought that you might express by (59), you token a Mentalese sentence with such a descriptive singular term in subject position; for instance, a conceptually complete, hence

eternal, version of ‘The curving glass vase in Nawab’s gift shop . . .’. However, a number of authors have made it clear that this cannot be right (see, for instance, Bach 1986, 1987; Levine 1988; Elugardo and Stainton forthcoming a). Levine (1988: 226) points out various problems for the description view. First, a person may be entirely unaware of any definite description under which she thinks of the object and although she might acknowledge, if asked, that the object she has proclaimed beautiful is indeed the curving glass vase in . . . , still there is no clear sense in which she had that description in mind when she uttered (59). Second, for any description of the form ‘The F’ that might be proposed as the one that represents the object, the sentence ‘That is the F’ seems to express something informative, that is, not something equivalent to ‘The F is the F’. Third, for any description under which a person might think of the object, she could turn out to be mistaken in thinking the object satisfies the description (for instance, the vase is not made of glass but of perspex) and yet she has, nonetheless, had a thought about that very object.

Bach (1986: 188) makes a more general point against any wholesale descriptivist account of thought:

If *all* your thoughts about things could only be descriptive, your total conception of the world would be merely qualitative. You would never be related in thought to anything in particular. . . . Since the object of a descriptive thought is determined *satisfactorily*, the fact that the thought is of that object does not require any connection between thought and object. However, the object of a *de re* thought [i.e. a thought of that object] is determined *relationally*. For something to be the object of a *de re* thought, it must stand in a certain kind of [real, natural] relation to that very thought.

The clearest case of this sort of directly world-connected thought is one that involves an object currently perceived, as in the thought about the vase occasioned by seeing it in a shop. The relation between the object and the thought about it is a causal one: the object causes a pattern of activation on a sensory receptor, which, in turn, causes the formation of a certain percept, which eventuates in the tokening of a symbol in a Mentalese sentence, which is a thought about the object.

At this stage, we can approach the symbol question with some constraints in place. One is that the mental demonstrative element does not specify (and is not an abbreviation for) an individuating description of the object in question; rather, it must be ‘directly referential’, that is, it must reflect the causal relation between the object and the occurrence of the mental symbol. Another constraint follows from the fact that it is possible to think of the same object in different ways without realizing that it is the same object. For instance, you might think of an object that you can see, that it is a candlestick, and of an object that you can hear, that it is a flute (so not a candlestick), without realizing that they are one and the same object, so the mental symbols representing the object in the two thoughts must be distinct. One of the attractions of the description view is that it can meet this second requirement (‘the long wooden object with a flute-like shape is . . .’; ‘the instrument making a sweet breathy sound is . . .’), but this possibility is no longer in the running.

Simplifying somewhat, the proposal is that the mental symbol concerned is either a percept of the object, that is, a way in which a physical object can appear to a

perceiver (Bach 1986, 1987), or a mental pointer to a percept (Levine 1988).³⁷ So your thought about the vase that it is beautiful is represented as 'BEAUTIFUL ($[\rightarrow \alpha]$)', where ' $[\rightarrow \alpha]$ ' is a Mentalese pointer and ' α ' is your percept of the vase. My thought about the two cars, expressed verbally as 'this one is better than this one', is represented as 'BETTER THAN ($[\rightarrow \beta]$, $[\rightarrow \gamma]$)', where ' β ' is my percept of the first car referred to and ' γ ' is my percept of the second one, and the two thoughts about the flute are 'CANDLESTICK ($[\rightarrow \Omega]$)' and 'FLUTE ($[\rightarrow \Delta]$)', where ' Ω ' and ' Δ ' are different percepts of the one object. So Mentalese has at its disposal a category of symbols not present in natural language, symbols that denote an object without being descriptive of it, while, nonetheless, reflecting the uniqueness of that object for the thinker at that moment.

The perceptual component in the Mentalese sentences given in the previous paragraph functions as a mental indexical: its semantic value (its referent) is determined, not via the 'satisfaction' of any conceptual conditions, but by an appropriate perceptual causal relation between an object in the context and that mental token. So thoughts in which these symbols occur are not eternal propositions; that is, their truth conditions are context-dependent. This is not a problem for the Compositionality Principle which requires only that the semantic value of the sentence be determined by the semantic value of its parts (and the combinatorial structure). Indexicals are parts of Mentalese sentences and their semantic values (the entities in the external world that caused their occurrence), together with the semantic values of the descriptive conceptual content of the sentences, play the role required by the principle.

There is nothing in these Mentalese symbols that makes them about one particular object as opposed to another except for their causal histories; that is, one and the same indexical symbol can pick out different objects in different contexts. This last point could, perhaps, do with some elaboration (see, especially, Burge 1977; Bach 1986, 1987), but the following simple illustration, adapted from Bach (1987: 22), may suffice here. I am looking at a tomato and I have a thought of the form 'RIPE ($[\rightarrow \#]$)', where ' $\#$ ' is my percept of the tomato; I look away for a moment, then glance back and think again 'RIPE ($[\rightarrow \#]$)', and perhaps resolve to buy the tomato (= ' $[\rightarrow \#]$ '). However, in the brief moment that my attention was elsewhere, the tomato I had seen first was removed and another, very similar, put in its place. As regards the object mentally demonstrated the two thought tokens are identical in form, and so, by the formality constraint, must play one and the same role in my thought processes, but they are truth-conditionally distinct: the first one is true iff the first tomato is ripe and the second one is true iff the second tomato is ripe. Another way of seeing the point is to envisage the possibility of two individuals, one in the context just given, the other in some other tomato-involving context, each of whom tokens 'RIPE ($[\rightarrow \#]$)'. On a narrow (psychological) construal of thought content, the thoughts of these individuals are the same because the percepts are type-identical, but on a 'wide' (truth-conditional) construal of content, their thoughts are different because they are about different things (the token percepts have different causal histories).³⁸

Bach (1986, 1987) extends his account of *de re* thought, from perception-based thoughts about objects, to memory-based and communication-based ones. Each of

these involves a different, more complex, relation between a thinker and the object thought about than in the perceptual case, but, in both cases, a crucial feature of the account is the ultimate causal connection between the object in the world and the representation(s) of it in the mind. Levine (1988; 239, note 20) sees no reason to limit the type of representation mentally demonstrated to percepts; so, for instance, recalling a day last week when I quarrelled with a friend and then missed my train, I may have a thought which could be verbally expressed by 'That was a bad day', where the representation mentally pointed to is some complex of memory traces. The details and full range of possibilities do not matter for the limited point I want to make here, which concerns the indexical (context-dependent) nature of thought.

There are, clearly, many more mental indexicals than there are linguistic indexicals (any percept is potentially a mental indexical), and the two types of indexical belong to quite different sorts of representational systems. Linguistic indexicals are communicative vehicles *par excellence*; they encode a procedure, or rule for use (in the case of 'that', something like 'is used to refer to a salient entity'). Mental indexicals can be thought of as possible cognitive correlates of particular uses of a linguistic indexical on particular occasions, though, of course, many occurrences of mental indexicals are never verbally expressed. What these two kinds of element have in common is their intrinsic context-sensitivity; their semantic value (a Mentalese symbol in the case of a linguistic indexical, an object in the world in the case of a mental indexical) is determined contextually. However, just as there are important differences in the functioning of the two kinds of representational systems, public language and Mentalese, and in their respective semantics, translational for linguistic sentences, truth-conditional for Mentalese sentences, so the appropriate notion of context is different in the two cases. The processes involved in understanding a verbal utterance lie within the solipsistic psychological system of representations and computations, so the 'context' that addressees employ in understanding an utterance (in disambiguating, in assigning referents to indexicals, in recovering unarticulated constituents, etc.) is a set of mental representations. The context of thought, on the other hand, is external to the system, it is the world of objects and states of affairs within which the thinker is situated and parts of which he mentally represents. The context for utterance interpretation (a set of mentally represented assumptions) can be called an I-context ('I' for internal, individual, interpretative), as opposed to the E-context of the semantics of thought ('E' for external, environmental). As McGinn (1982: 209) puts it, the object which a mental indexical is about 'is determined by the occurrence of a representation *in* a context, not by way of a representation *of* the context.' The relation between objects in the E-context and mental representations of those objects is not itself mentally represented and so is not a matter for a (solipsistic) cognitive psychology; rather, it is the domain of a 'real' semantics, whose concern is to articulate the relation between mental representations and the features of the outside world that are represented.

This brief discussion of mental indexicals has focused on thoughts about individuals and objects other than oneself. Thoughts about oneself (first-person thoughts), such as the thought that one might express verbally by 'I am tired', seem to be special in certain ways; they issue 'from within', as it were, they are not

perception-based in quite the way that thoughts about others are, and, while one can be wrong about the identity of the referent of one's thoughts about other individuals and objects, one's 'I'-thoughts are immune to that sort of error (I am acquainted with the object of my thought 'I am tired' in a way that I am not acquainted with any other object). Similar points apply to thinking of a place as 'here', or of a time as 'now'. Each of these involves a special source of information: information about a person acquired through being that person, information about a place or a moment in time gained from occupying the space and time that one does. These and many other facets of such thoughts, which I cannot pursue here, are investigated in detail by Recanati (1993). Some similar observations regarding tensed thoughts (as opposed to tenseless ones that specify a time descriptively) are made by Higginbotham (1995); such thoughts involve a mental indexical element of a temporal sort (for the present, the past, the future).

Consider the following pairs of thoughts, which, let us assume, are truth-conditionally equivalent, so that the first two refer to the same person and the second two (taken from Higginbotham 1995: 22) refer to the same moment in time:

- (60) a. My child is crying.
 b. The child in the next room is crying.
- (61) a. My root canal operation is over.
 b. My root canal operation is over as of 4 p.m., 31 October 1994.

The thought represented by (60a) may set in train certain further thoughts and cause me to take action (for instance, to go and pick the child up) in a way that the thought represented by (60b) might not if, for some reason, I am unaware that the child in the next room is, in fact, my child. Similarly, the thought represented by (61a) might bring with it a feeling of deep relief, while the thought in (61b) would not occasion a sense of relief, unless I know of the time at which I have the thought that it is, or is later than, 4 p.m., 31 October 1994. In short, indexical thoughts have a range of causal properties that distinguish them from any truth-conditionally equivalent descriptive thoughts.

Summing up, while occurrences of Mentalese sentences are free from many of the sources of truth-conditional underdeterminacy that are typical of natural-language utterances, they are not, in general, eternal either. Recall Katz's (1972, 1981) position that for every context-dependent, non-eternal sentence there is an eternal, hence context-independent, sentence, so for every occurrence of an indexical, there is a descriptive expression 'that has the same reference as the indexical element it replaces but whose referent stays fixed with variations in time, place, etc.' As we've seen, this is not true for natural-language sentences. As regards Mentalese sentences, even if mental indexicals can be replaced by mental representations whose reference is eternal, a thought which involves tokening a mental indexical and a thought with the corresponding eternal representation would be different thoughts, playing different roles in thought and feeling, and leading to different behaviours on the part of the thinker.

As discussed in section 1.6, according to Searle, the content of thoughts (and intentional states generally) determines conditions of truth only against a Back-

ground. If this is correct (see note 32), then thoughts share the property of underdeterminacy with natural-language utterances to the extent that it is an effect of the Background. In addition, as we've seen, indexicality of a sort is a common feature of thoughts. But, whatever properties of context-dependence and truth-conditional incompleteness our thoughts may have, the encoded content of natural-language sentences takes underdeterminacy to a different level, in both quantity and type.³⁹

1.8 Summary

To end this long chapter I will briefly and baldly state its main claims, which are assumed in the rest of the book:

- 1 Linguistically encoded meaning underdetermines the proposition expressed by an utterance (its truth-conditional content).
- 2 Linguistic underdeterminacy is an essential feature of natural languages because there are no eternal sentences in natural languages.
- 3 The primary mental capacity underlying the communicative and interpretive powers of humans is the capacity to infer the mental states of others, and this runs to several orders of attribution.
- 4 The pragmatic inferential capacity, whose specific domain is utterances and other communicative acts, employs a particular interpretive strategy, distinct from that of the more general capacity of mental state attribution, and warranted by the presumption of optimal relevance that is automatically conveyed by such stimuli.
- 5 All intentionality (mental and linguistic) is dependent on a massive Background of weakly manifest (taken-for-granted) unrepresented assumptions and practices.
- 6 Linguistically encoded meaning is, then, doubly underdetermining of utterance meaning; as well as its own inherent underdeterminacy it inherits the underdeterminacy of the representational states it is used to express.

Finally, I should mention that in the chapters that follow I will sometimes try to give some sort of representation of the proposition expressed by an utterance, and this will often look like some attempt at finding an eternal natural-language sentence. This will not be the intention. In fact, in most cases the given representation would fail miserably as a candidate for an eternal sentence for the sorts of reasons surveyed in section 1.3; the point will usually be to alert the reader to those aspects of the proposition expressed that have been *pragmatically* supplied, so points of difference from the semantic representation encoded by the linguistic expression will be highlighted. Precisely because of the non-existence of eternal sentences, the propositional representations may contain numbers and other symbols (like the hopelessly inadequate but suggestive 't's for temporal reference) in a bid to represent elements of meaning that natural-language sentences cannot represent.

NOTES

- 1 However, Kent Bach is an exception to this. He ties 'what is said' as closely to linguistic meaning as one can without equating the two, so that for him 'what is said' need

not be propositional (Bach 1994a, 1994b, 1997/99a, 2001). His views are discussed in chapter 2, section 2.5.

- 2 Talk of linguistic error and misuse raises some quite fundamental questions about the nature of language. It seems to presuppose the existence of a common public language which is used by a particular population of individuals, some of whom have a less complete grasp of it than others. However, on an internalist and ‘idiolectal’ view (such as Chomsky’s – see any of his writings) there is really no issue of ‘correctness’ or of what is ‘actually’, as opposed to ‘mistakenly’, encoded by a linguistic form (hence my scare quotes in the text). An individual’s linguistic knowledge (including form/meaning mappings) simply is her language and its utility for communication with another individual depends on its degree of similarity to the language (idiolect) of the other. Nevertheless, whether we individuate languages as social entities or, much more finely, as individualistic entities, the issue for a cognitive pragmatic account of utterance interpretation remains the same: both communication failures due to disparate encodings, and communication which is successful only because either the speaker or the hearer recognizes such a disparity and adjusts for it, are real phenomena and have to be accounted for.

For discussion of so-called language misuses, in the context of different conceptions of ‘language’, see Davidson (1986), Chomsky (1987; 1992a; 1995) and Chng (1999).

- 3 All page references to articles by Grice are to the reprintings in Grice (1989b).
- 4 According to this approach to ambiguity, there are, strictly speaking, no ambiguous lexical items, phrases or sentences, but rather distinct lexical items (e.g. ‘bank₁’ and ‘bank₂’) that happen to have the same phonological form, and distinct phrases/sentences (e.g. ‘I went to the bank₁’ and ‘I went to the bank₂’, ‘[visiting relatives]_{NP}’ and ‘[(e)_{NP} [visiting relatives]_{VP}]_S’), each making a distinct contribution to truth-conditional content. In accordance with this, context is said to play a ‘pre-semantic’ role in distinguishing ambiguous forms (see Perry 1997, 1998). That is, disambiguation is a process of figuring out from the context which of several lexical items or phrases with the same surface perceptual properties has been uttered, that is, which semantic entity we are dealing with.

However, earlier work in the application of truth theory to the semantics of natural language did not allow itself this sorting of perceptually identical forms into distinct linguistic entities and, as a result, ambiguity presented a problem. Davidson (1967, 1970) tried out the following two possibilities:

- (i) For any *a*, ‘is a bank’ is true of *a* in English if and only if *a* is a bank.
- (ii) ‘John went to the bank’ is true for an English speaker *x* at time *t* if and only if either John went to the financial institution and the circumstances surrounding *x* at *t* meet condition C, or John went to the wall of the river channel before *t* and the circumstances surrounding *x* at *t* meet condition D.

In the first of these, the idea is that, provided one can translate the ambiguity in the object language into an ambiguity in the metalanguage, truth will be preserved. This turns out not to work (see Parsons 1973). In the second attempt, the idea is that the theory should specify a single T-sentence with a disjunction of truth conditions, together with reference to different contextual conditions for each disjunct. This too is problematic (for discussion, see Parsons 1973; Lycan 1984, chapter 2; and Cohen 1985).

- 5 This is not to say that the semantics of pronouns is a cut-and-dried matter; it is not. The point here is that we do not need to get the semantics of pronouns straight in order to convince ourselves that understanding utterances of them presents a task for pragmatics.

- 6 Katz was led to abandon the Chomskyan psychological conception of language, of which he had been a staunch supporter, by his philosophical interest in necessary truths and his belief that they are expressible in natural language. He claims that Chomsky's mentalism denies the possibility of genuine necessary truth in natural languages, since the most it can offer is the concept of something's being necessary *relative to human cognitive capacities* (Katz 1981: 4–6).
- 7 Katz (1981) notes conceptualist scepticism about the possibility of doing semantics; here is one of the examples he quotes: 'It seems that other cognitive systems – in particular our system of beliefs concerning things in the world and their behaviour – play an essential part in our judgements of meaning [sense] and reference, in an extremely intricate manner, and it is not at all clear that much will remain if we try to separate the purely linguistic components of what in informal usage or even in technical discussions we call the "meaning of a linguistic expression". I doubt that one can separate semantic representation from beliefs and knowledge about the world' (Chomsky 1979: 142). The same position is reiterated in more recent work, for example Chomsky (1995: 26). Little wonder then that Katz, the prime mover in developing the 'semantic component' of a grammar, prefers a Platonist conception. Recall Chomsky's scepticism regarding a theory of utterance interpretation, pointed out in the introduction; the current quote indicates a deeper, more thoroughgoing, scepticism about the prospects for theories of meaning of even quite a narrow sort.
- 8 Searle (1992: 131, 155) too emphasizes that every intentional (i.e. representational) state has an 'aspectual shape', by which he means to emphasize the first-person perspective we cannot but have on the objects we perceive and think about. I take it that 'aspectual shape' is another term for 'mode of presentation'.
- 9 As Perry (1977) notes, Frege's view that each thought is the sense of some sentence (though there are sentence senses which are not thoughts because they are incomplete) was severely tested by indexicals like 'here', 'now' and 'I'. Having recognized the 'special and primitive way in which we each present ourselves to ourselves' (Frege 1918a/77: 12–13), he could not but accept that there are incommunicable senses. So it seems that even those most drawn to it cannot hold the Effability Principle in its strongest form.
- 10 Clearly, the point depends on the assumption that the conceptual content of the description is a crucial component of the proposition expressed (or statement made), so that differences in descriptive content entail differences in proposition expressed (or statement made), as is most clearly the case on attributive (as opposed to referential) uses of descriptions, in the sense of Donnellan (1966/91). This is indeed Wettstein's assumption, since the main purpose of his 1979 paper is to argue against the Fregean view that indexical expressions (names, pronouns, demonstratives), uttered in a particular context, acquire a particular sense or descriptive content, which they contribute to the proposition expressed by the utterance. He favours the direct reference view of names and other indexicals, associated with Mill, Kripke and Donnellan, according to which the context in which the indexical sentence is uttered 'reveals *which item* is in question' rather than providing 'some unique characterisation' of it (Wettstein 1979: 96). So his argument against the possibility of finding that one among a number of eternal descriptions which is the very one that enters into the proposition expressed by an utterance of an indexical sentence is ultimately directed towards this bigger theoretical end. It seems to me that the directly referential view of indexicals is incompatible with the strong effability positions.
- 11 It might be objected that, in fact, the dependence of the intended reference on the domain of discourse does not affect the proposition expressed but comes in at some other level

altogether, perhaps as an implicature. This is a standard Gricean kind of gambit which we will see more of in chapter 2. The result here would be that ‘what is said’ is that the prime minister of Britain, whoever that actually is, (i.e. Tony Blair) is in the next room and ‘what is implicated’ would be that the prime minister of Britain, according to Lucinda, (i.e. Peter Mandelson) is in the next room. This runs counter to a quite robust intuition that what the speaker of (25) said (that is, the proposition she expressed) was that Lucinda will be delighted to find that the person who she believes to be the current prime minister of Britain is in the next room. If the speaker in the same set of circumstances were to add: ‘Of course, the prime minister is not, in fact, the prime minister’, she would not be understood as contradicting herself, but as expressing the proposition that the person whom Lucinda thinks is the prime minister is not in fact the prime minister.

Furthermore, as Recanati (1987b: 68–72) points out, different elements in a sentence may be interpreted relative to different domains of discourse, so in the current example while ‘the prime minister of Britain’ is interpreted relative to Lucinda’s belief world, the predicate ‘is in the next room’ is interpreted relative to another domain (the actual world). In principle, there could be several different discourse domains relative to which different parts of an utterance are to be interpreted. Recanati suggests that, at the level of sentence meaning, every predicate comes with a variable ranging over domains of discourse, so that, for any utterance, these variables have to be contextually determined in the process of deriving what is said (the proposition expressed, the truth-conditional content of the utterance). If this is right, it follows that every sentence is intrinsically context-dependent and non-eternal.

- 12 Dan Sperber, Deirdre Wilson and I have discussed a range of other cases where processes of pragmatic adjustment of a lexical concept eventuate in an *ad hoc* concept in the proposition expressed, a concept which may be narrower than the original one, as in the example in (27), or wider as in the examples in (28), or a combination of the two. See, for instance, Sperber and Wilson (1997/98a), Wilson (1995), and Carston (1996b/97a). *Ad hoc* concept formation and its role in the proposition expressed is the subject of chapter 5 of this book.
- 13 An obvious (Gricean) sort of response here would be to say that the proposition expressed (what is said) by the utterance does contain within it the precise linguistically encoded concept, even though what is meant is something looser. On this view, the speaker does not endorse (that is, ‘mean’ or communicate) the proposition expressed and in most cases that proposition is clearly false (e.g. France has six equal sides). The broader concept which is communicated is registered at some other level of utterance meaning (the implicature level). This is also the standard relevance-theoretic position (Sperber and Wilson 1986a/95b: 232–5). In chapter 5 I argue for a different relevance-theoretic account.
- 14 I am ignoring here crucial questions about what the interrogative and other moods actually encode. It does look very much as if a proper analysis of moods will reveal another area of intrinsic underdeterminacy in natural language (see Gazdar 1981; Wilson and Sperber 1988a, 1988b; Clark 1991, 1993a).
- 15 Wilson (1999b/2000) provides an illuminating exposition of how psychological research on the ‘theory of mind’ capacity and work within the broadly Gricean inferential pragmatic tradition interrelate, and how both of these bear on the more general mental ability for metarepresentation. She also brings in a third strand of research on metarepresentation, one which focuses on the various ways in which metarepresentational elements may occur within the content of utterances; for instance, cases of quotation, direct or indirect, and of allusions to, and echoes of, other people’s thoughts and utterances. Some cases of such ‘metalinguistic use’ are discussed in chapter 4.

- 16 The issues of how exactly the mind-reading ability and the system for interpreting ostensive stimuli are related, what it means to think of them as mental modules, and how they are situated in our overall cognitive architecture are far from settled. Sperber and Wilson (1986a/95b) and Wilson and Sperber (1986b) adopted Fodor's (1983) distinction between modular systems of perception and motor output, on the one hand, and non-modular central systems, on the other, and they took pragmatics to belong at the non-modular centre (along with other apparently context-sensitive reasoning systems). However, in more recent work, Sperber (1994b, 1996) has made a case, largely based on evolutionary considerations, for a thoroughgoing modularity of mind, within which central inferential systems are also modular, including those systems that manipulate metarepresentations, such as the theory of mind (or metapsychological) system and the pragmatic (or ostension comprehension) system. Smith and Tsimplici (1996) suggest that a concept of 'quasi'-modularity is more appropriate in the case of these and other systems dealing in conceptual representations. Carston (1997b) discusses, in a preliminary way, the idea of a pragmatics module.
- 17 I am not giving a fully comprehensive outline of relevance theory in this book, although many of the main concepts and distinctions of the pragmatic theory developed within the broader framework are introduced and discussed in chapter 2 (section 2.3), and there is a glossary of terms in appendix 1. For introductions to the framework, see Wilson and Sperber (1986b; 1986c), Blakemore (1992), Wilson (1994; 1999), Noh (1998b, 2000, chapter 2). For more advanced accounts, see Sperber and Wilson (1986a/95b), Sperber and Wilson (1987) and Sperber and Wilson (1995a).
- 18 Premack (1990) presents arguments for the coevolution of the human linguistic system and 'social modules', including the theory of mind, and Sperber (1990) gives a succinct demonstration that a linguistic system has adaptive value only for a species already able to engage in ostensive-inferential communication. These ideas are further developed in Sperber (2000) and Origgini and Sperber (2000). Given the assumption that the language faculty and the metarepresentational capacity coevolved in humans, Sperber considers the question 'which of these two, the linguistic or the metarepresentational, might have developed first to a degree sufficient to bootstrap the coevolutionary process.' He considers both possibilities (language first, or metarepresentations first) and concludes that the more plausible scenario is that a metarepresentational capacity first developed in a social environment involving both competitive and cooperative interactions, that that capacity made possible ostensive communication, perhaps as a side effect, and 'the beneficial character of this side effect turned it into a function of metarepresentations, and created a favorable environment for the evolution of a new adaptation, a linguistic ability' (Sperber 2000: 127).
- 19 The issue of the language (or languages) of thought is clearly relevant here. Carruthers (1996) and Horwich (forthcoming) support an account of language and thought in which a person's Mentalese just is their natural public language. Fodor (1975, 1987a) is well known for his view that the medium of thought, though syntactic like natural language, is distinct from (and precedes, and enables, the acquisition of) a person's public language; however, he seems to assume a more or less one-to-one mapping between words and the basic constituents of thought, that is, concepts (see Fodor 1975: 152–6, and Fodor and Lepore 1991: 333). Both views are incompatible with the essential underdeterminacy thesis. While any settled answers to questions about the nature of thought are a long way off, I suggest that the arguments presented in this chapter against strong effability (the linguistic encodability of thoughts and propositions expressed) are just as much arguments against the view that thought consists either of public-language sentences or of representations whose constituents are in a one-to-one relation with the

constituents of natural-language sentences. For discussion of the implications of linguistic underdeterminacy for the view of thought as ‘inner speech’, see Elugardo and Stainton (forthcoming a).

- 20 The property of ‘coherence and appropriateness to situations’, which Chomsky frequently mentions as one of the features that makes natural language unique (along with ‘unboundedness’ and ‘freedom from stimulus control’), strikes me as also explained by these wider cognitive capacities involved in ostensive communication (that is, language use) rather than as a property of the language system (competence) *per se*. (See, for instance, Chomsky 1966: 4–5; 1988: 5.)
- 21 Gross (1998, chapter 1) makes (and elaborates on) three claims for the ‘utility of context-sensitivity’: (i) it is essential for various aspects of the learning of language; for instance, the learning of many predicates depends on exposure to exemplars in the environment; (ii) it makes for a considerable increase in the efficiency of communication, since appropriate exploitation of shared knowledge and circumstances reduces the effort and time needed in both production and comprehension; (iii) it bestows a greater flexibility on language, making it fit for a much wider range of purposes than it would otherwise have. He goes on to endorse the view that context-sensitivity is not eliminable, even in principle, from natural language use, that is, in the terms of section 1.2.2 of this book, underdeterminacy is an essential feature of language.
- 22 Higginbotham’s (1988) treatment of the referential use of definite descriptions parallels that of the other referential cases (pronouns, demonstrative descriptions) and captures the specific contribution of the definite article in terms of conditions enumerated in the antecedent of the conditional T-statement:

- (i) If u is an utterance of ‘The dog is hungry’, and the speaker of u refers with ‘the dog’ to x , and (a) the speaker does not refer to anything else with ‘the dog’; (b) x is a dog; (c) x is obvious or familiar, then [u is true iff x is hungry].

The assumption that the referential use of a definite description is a matter of linguistic semantics is highly debatable. A currently popular view is that the referential use is pragmatically inferred (see Recanati 1993; Bezuidenhout 1997b; Rouchota 1992, 1994b; Powell 1999), on the basis of a radically underdetermining linguistic semantics for ‘The F’. If this proves to be correct, then there is no T-statement of this sort for definite description sentences.

- 23 In addition to applying the modified truth-conditional approach to these cases, Gross (1998, chapter 3) considers two other possible ways of handling them, which can be dubbed the implicature approach (or the Gricean gambit) and the ambiguity approach. According to the first of these, there is no part context-sensitivity of predicates and the sentence ‘The book is green’ simply expresses a false proposition if the book has some other colours on it, though it may implicate (via considerations of relevance) the true proposition that a certain part of the book is green. According to the other line of thought, there is not a single univocal sentence of the form ‘The book is green’ that is uttered across the different contexts expressing different propositions. Rather, there are a large number of sentences of this form, to each of which the truth-theory must assign a distinct T-statement, for instance:

- (i) ‘The book is green₁’ is true iff the cover of the book is green.
 (ii) ‘The book is green₂’ is true iff the spine of the book is green.
 (iii) ‘The book is green₃’ is true iff the whole cover and spine apart from the lettering are green.

Gross provides a range of good arguments against both of these attempts to deny the part context-sensitivity of certain predicates.

- 24 In a response to this argument of Lewis, Harman (1974) observes that just as we can know how to translate between two languages without knowing the meaning of the expressions in either, so we can have knowledge of truth conditions without having knowledge of meaning: ‘there is a sense in which we can know the truth conditions of an English sentence without knowing the first thing about the meaning of the English sentence’. When we first come upon the sentence ‘All mimsy were the borogoves’ in Lewis Carroll’s *Through the Looking Glass*, we know the following:

(i) ‘All mimsy were the borogoves’ is true iff all mimsy were the borogoves.

Harman’s claim is that we can know this without knowing the meaning of the sentence ‘All mimsy were the borogoves’. For interesting discussion of this general issue and a deflation of Harman’s point, see Higginbotham (1989). On the relative merits and shortcomings of translational approaches to semantics, see Lepore and Loewer (1981) and Lepore (1996).

- 25 For further discussion of the relevance-theoretic view of these two kinds of semantics, see Sperber and Wilson (1986a/95b, chapter 4), Blakemore (1987, chapter 1), Carston (1988: 175–8), Wilson and Sperber (1988b: 134), Sperber and Wilson (1995a: 257–8), Carston (1999b).
- 26 Unfortunately, different uses of the term ‘logical form’ abound. There are so-called ‘regimented’ notions of logical form, such as those of Frege or Russell, which are intended to replace the ambiguous, vague and context-sensitive sentences of natural language, in the interest of expressing scientific propositions in a way that makes their logical properties transparent and, perhaps also, makes their relation to the external world transparent. Stanley (2000: 391–2) distinguishes this ‘*revisionary* conception of logical form’ from the ‘*descriptive* conception of logical form, [according to which] the logical form of a sentence is something like the “real structure” of that sentence’. Clearly, linguists’ notions of logical form, such as Chomsky’s LF (‘the level of linguistic representation at which all grammatical structure relevant to semantic interpretation is provided’ – Hornstein 1995: 3), and the relevance-theoretic notion belong to the descriptive conception. These have in common a concern to specify the (underlying) structure of a sentence (or phrase) in such a way as to reflect its semantic properties, its meaning, as opposed to its surface syntactic structure, but they appear to diverge in a range of ways that have yet to be properly explored and assessed. For interesting discussion of some of the ‘descriptively’ conceived notions of logical form within current linguistics and philosophy, see Higginbotham (1993a), Neale (1994) and Larson and Segal (1995: 100–5).
- 27 Levinson (1988, 1989, 2000) seems to be of the view that if you do not assign something of a fully propositional nature (thereby admitting of truth conditions and the description of sense relations) to natural-language sentences, you simply cannot say anything of interest about linguistic semantics. In fact, although a full account of the conceptual and procedural encodings of natural-language words and sentences does not exist (yet), the programme for giving such an account is clear enough and relevance theorists have made many concrete proposals concerning the semantics of particular natural-language expressions. See, for example, Blakemore (1987, 1988, 1989b, 1990, 1997b), Blass (1990), Breheny (1999), Carston (1988/91, 1993, 1994a, 1994b/96a), Clark (1993a, 1993b), Groefsema (1995a, 1995b), Iten (1998, 2000), Papafragou (1998a, 1998b, 1998c), Powell (1999), Rouchota (1994a, 1994b, 1996), Žegarac (1991, 1993). Work in other frameworks also shows that giving an account of natural-language seman-

tics in terms of an intermediate (conceptually oriented) system of representations is very much a live option (see Katz 1972; Jackendoff 1983, 1990).

- 28 For the most part, I have encountered this propositional view in discussion and in some unpublished work, and can give only one reference to it in available written form: Chng (1999). Noel Burton-Roberts has defended the position strongly on several occasions in discussion; I think he would stand by the second version I give in the text.
- 29 A very similar sort of argument would be endorsed by those direct reference theorists who, following Kaplan (1977/89a), make a distinction between the semantic character of indexicals (roughly their encoded linguistic meaning) and their semantic content, that is, their truth-conditional contribution in a particular context. On this view, although (i) and (ii) are universal truths of some sort, they are not necessary truths:
- (i) If she is kind some female is kind.
 - (ii) If that spy is clever then someone is a spy.

Although there is no entailment between the contents of the antecedent and the consequent in either of these cases, the appearance (illusion) of an entailment is explained in terms of the semantic character of the pronoun 'she' and of the complex demonstrative 'that spy'. The character of 'she' is such that on any proper use it will deliver as its content someone who is, in fact, female; the character of 'that spy' is such that on any proper use it will deliver as its content someone who is a spy. See Kaplan (1989b) for a discussion of such logical truths that are not necessary truths, and Braun (1994) for an application of the idea to complex demonstratives, such as 'that spy'.

- 30 There are other manifestations of the view that the meaning of any natural-language sentence (including indexical sentences) is a proposition. These are usually of a highly abstract sort, for instance, the 'diagonal proposition' of Stalnaker (1978), the 'external proposition' discussed by Recanati (1993: 289–91) and, perhaps, the 'created proposition' of Perry (1988). What each of these concepts is designed to capture is the context-invariant conditions which must be met by *any* utterance of a sentence for it to express a true proposition. For example, the diagonal proposition associated with the sentence S 'I am happy' is the proposition that there is an x such that x utters S and x is happy. As Recanati (1989a: 237) puts it, in a very useful discussion, 'the diagonal proposition globally indicates the conditions under which the utterance expresses a true proposition, but it does not tell us which of these conditions are contextual conditions, i.e. conditions which must be contextually satisfied for the sentence to express a definite proposition, and which are truth conditions proper, i.e. conditions which must be satisfied for the proposition expressed to be true'. In other words, no distinction is made between conceptual aspects of sentence meaning, which enter into the proposition expressed, and procedural constraints, that do not. This extraction of a very general proposition can obviously be done and may be useful for certain purposes; for instance, Stalnaker (1978) suggests that if for some reason an addressee cannot figure out what proposition has been expressed by a particular utterance, he can at least add the diagonal proposition to the common ground (or set of contextual assumptions). However, it is clearly not the sort of representation of linguistic meaning which can function as input to the pragmatic process of figuring out what proposition has been expressed, because it erases the distinction between two kinds of linguistic meaning, a distinction which plays an essential role in guiding those processes.
- 31 This is what he calls the **Connection Principle**: an essential property of *mental* features of brain states (as opposed to *non-mental* features, such as axon myelination, for instance) is that they are either conscious or potentially conscious (see Searle 1992:

155–61). From this follows his repudiation of an assumption held in current cognitive science: that we can have unconscious knowledge which is, in principle, inaccessible to consciousness (Searle 1992: 197–248), an assumption behind the Chomskyan account of language, for instance. Both Searle and, following him, Nagel (1993) maintain that neither language nor visual perception are strictly speaking psychological; they each simply involve some physical (hardware) mechanism that functions in certain constrained ways. Chomsky (1994: 197–201) presents compelling considerations against the Connection Principle, the mind–body dualism it entails and the alleged non-mental nature of language. I cannot pursue this debate here, but I believe it is possible to take on board the thesis of the Background, or something closely akin to Searle’s conception of it, without automatically having to accept these further claims.

- 32 It is not perfectly clear to me whether Searle intends his ‘Connection Principle’ to apply to everything he is prepared to call mental or just to intentional states (and, of course, qualitative states such as pain). The mental and the intentional are not coextensive on his conception, since the Background is mental though not intentional (Searle 1991: 290). As I’ve pointed out, it does seem that at least some elements of the Background can become intentional states when, for some reason or other, something hitherto ‘taken for granted’ becomes apprehended or believed (see Searle 1992: 184–5) and so is presumably no longer part of the Background. It seems, then, that at least some aspects of the Background are accessible to consciousness. The concept of a ‘manifest assumption’ does as good a job as any of making sense of all of this.
- 33 In his early work on speech acts, Searle (1969) made a claim that follows from this principle: the study of the full range of speech acts people can perform with linguistic expressions in contexts can be confined to the study of explicitly performative sentences. Gazdar (1981) shows that this is false: even explicitly performative sentences do not always determine the speech acts they can be used to perform and it is not the case that for every speech act achieved pragmatically, there is a corresponding explicitly performative sentence.

Recanati (forthcoming a) provides a more sustained exploration of Searle’s Principle of Expressibility, and its relation to his thesis of the Background, than I have given here; in particular, he argues against Searle’s position that Background-dependence applies not only to the truth-conditional content of utterances but also to thoughts, beliefs and intentional states quite generally.

- 34 Andreas Kemmerling has pointed out to me that, although many people attribute this view to Frege, he did not hold it in quite this form. For him, Thoughts were ontologically prior to whatever may be considered their parts or constituents, so it is misleading to say that the *sense* of a sentence (i.e. a thought) is *determined* by the sense of its constituents. The level of unstructured thoughts is basic, and while, for their own epistemic purposes, humans impose a structure on thoughts (perhaps a function/argument structure), numerous different structurings are possible. According to Kemmerling (personal communication, 1999), Frege would have accepted the Compositionality Principle only in a version which contains relativizations to methods of decomposing thoughts (senses) and sentences: ‘Let M be an acceptable method of decomposition of Thoughts-and-sentences-expressing-them, then the sense/reference of an M-decomposed expression is determined by the sense/reference of its M-constituents and the manner in which they are M-combined.’ (Frege is most explicit on this in his (1892b); see especially p. 49 of the (1980) reprint. For useful discussion of Frege’s view, see Janssen (1997a: 420).)

Like most linguists nowadays who refer to a Principle of Compositionality, my discussion in this section is concerned only with compositionality as a property of aspects of human psychology. The compositionality of linguistic meaning (logical form) is a func-

- tion of the mental language faculty and the compositionality of propositional thought (or ‘content’, or a psychologized version of ‘sense’) is a property of mental representation. This makes for an important difference from Frege’s Platonist conception: the only decomposition of sentences, or of thoughts, that is relevant is the one that human minds actually manifest.
- 35 The Compositionality Principle has a long history and a range of interpretations, varying in the strictness of the syntax–semantics relation required, and in the way in which component terms such as ‘meaning’ or ‘semantic value’ are defined. It can be viewed as an empirical claim about natural language, or as merely a methodological principle or evaluation criterion in the development of semantic theories. For discussion of these and other issues, see Welsh (1986), Janssen (1997a, 1997b) (the latter a short version of the former), Grandy (1990b), Partee (1995) and Woodfield (1999). For linguistic data which seem to weigh against the general truth of the principle, see Lahav (1989, 1993), Pelletier (1994) and, in what appears to be a conversion to the linguistic underdeterminacy view, Fodor (2001). For a position similar to the one I am proposing (that is, sentence-type meaning is compositional, but determination of the proposition expressed depends on non-compositional processes of pragmatic enrichment), see Powell (2000, forthcoming) and Blutner (1998, 2002); the latter also makes a strong case against the possibility of any principle of ‘semantic/pragmatic’ compositionality.
- 36 Some caution is necessary here, since we are able to think about our thoughts, or at least some of them, thanks to our metarepresentational capacity. For instance, I may think about a belief I once had, say, the belief that humans are basically good, that it is not true and reflect on how I have come to this conclusion. The form of this reflective thought is: ‘It is not true that “humans are basically good”,’ that is, the mental representation of the original thought is itself represented; it is embedded in the representation which is my current thought. Furthermore, ‘metarepresentational thoughts’ may be incomplete in a way that is not possible for the simple thoughts I am confining myself to in this section. For instance, on the basis of overhearing someone’s conversation I might form the thought/belief that there are no longer any serins in Britain, without having any idea what a serin is (is it an animal, a bird, an insect, an outdated occupation, a mode of dress, etc?). My thought is not the same as that expressed by the speaker, who knew what she was talking about; mine is incomplete, semi-propositional, and it too has a metarepresentational component, namely ‘serin’. For my current concerns, I am ignoring the many interesting issues that arise in this realm of reflective (or representational) thought/belief. See Sperber (1982/85 and 1997a).
- 37 Levine (1988: 233) argues for the ‘pointer to a percept’ conception of mental demonstratives rather than the ‘percept itself’ conception: ‘I have many percepts in play at any one time, yet I do not demonstratively pick out each object of which I have a percept. The act of focusing my attention on one object within my perceptual field and thinking of it “*this* is red”, seems to involve a separate representation from the perceptual representation of the object itself. That is, there is the percept – present to mind, representing an object – and there is the act of pointing to it as a way of picking out the object it represents for selective attention.’ I find this distinction persuasive and so adopt the ‘pointer’ symbol for the rest of the discussion.
- 38 What I’m adverting to here is sometimes called the dual-component view of thoughts: thoughts have a ‘narrow’, subjective, internal aspect and a ‘wide’, objective, external aspect, so may be individuated in two distinct ways. Just as two thoughts with the same narrow content may have distinct wide content (as with the tomato-directed thoughts described in the text), so two thoughts with different narrow contents (and so different causal roles in mental life) may have the same wide content (that is, the same truth con-

ditions). There are many contentious issues in this area, especially regarding the notion of narrow content; see, for instance, McGinn (1982), Fodor (1990), Recanati (1993, chapters 11 and 12) and Jacob (1997).

- 39 There are phenomena, which I have not explored here, that might undercut my conclusion to some extent. I mention three examples. Perry (1986: 145) says 'there is no reason that thoughts that employ representations in the language of thought should not have unarticulated constituents, just as statements that employ sentences of natural language do.' He makes a case for the existence of beliefs of the 'It is raining' sort, which exclusively concern the local weather conditions of the believer and so need not have any location constituent (even of an indexical sort). Sorenson (1991) argues that there is vagueness in Mentalese, that is, that some of its predicates and quantifiers have borderline cases (e.g. the conceptual equivalents of 'tall', 'lazy', 'messy', 'numerous', 'somewhat'), so that sentences, hence thoughts, in which they appear are not truth-evaluable. However, he claims that the sources of vagueness in the essentially private language of Mentalese are different from those in public natural languages used for communication. Giaquinto (1997) discusses cases of what he describes as 'informative but propositionless' thought; these are thoughts which apparently include a constituent whose content is indeterminate across a range of possibilities. One of his examples is the thought 'The Venus Fly Trap is not an animal', where there is a range of possible 'animal' concepts and the thinker doesn't seem to have any particular one in mind, so that the thought does not determine a unique proposition but rather a proposition schema. It is nonetheless an informative thought in the sense that from it can be inferred certain true propositions, e.g. a Venus Fly Trap is not a cat. This is an extremely interesting and plausible possibility, which may not, however, weaken the general point that, indexicality aside, Mentalese sentences do not underdetermine their truth-conditional content, are not context-sensitive. After all, the point about these schematic thoughts is that they simply do not have determinate truth-conditional content.