LANGUAGE AND INTERESTS

Dinah Murray

Department of Linguistics
University College London
Gower St
London WC1

CONTENTS

Acknowledgements

Amongst the many people on whose support I have relied in writing this thesis, two in particular stand out. Firstly I want to thank my husband, who has supported and encouraged me through the many years it has taken me to write this thesis, despite never having actually read it (we have not wanted him to become my teacher). Without his patience and tolerance it would never have been written at all. And in its financial aspect his support has freed me from having to plough any particular academic furrow - without that freedom it would have been a very different book. I cannot estimate the degree to which his off-the-cuff comments and queries, and his written papers, on a variety of philosophical issues relevant to this thesis, may have influenced me over the years. But I'm certain all such influence has been beneficial. Secondly, I want to thank my supervisor, Deirdre Wilson, who has encouraged and supported me through the long years of my writing this, despite having actually read it. Her advice has always been excellent, her comments have always been constructive, and our disagreements have always been stimulating. Without her confidence in the project, my own would likely have run out long ago.

In one way and another, such a large number of people have contributed to this thesis that many are bound to be left out. I hope they will forgive me, and I thank them all for their time and interest. But it would be positively churlish not to name certain of the friends who've helped along the way. I'm

extremely grateful to all the following: Jeanette Buirski and Jeff Mason, for their persistent interest and lively comments - their philosophical background has helped greatly in putting issues and questions that have arisen into a wider context; Alan Campbell, Annabel Cormack and Ruth Kempson, for their critical acumen and good sense, which have helped keep me on the straight and narrow; Jean Aitchison, for commissioning the occasional paper and appreciating my prose style; Dick Hudson and Terry Tomkow, for introducing me to frames and schemas in the work of Minsky, Neisser, and Anderson; Robyn Carston, for introducing me to A Lover's Discourse (Barthes) - thoughts arising from which profoundly affected some of my conclusions; Mike Lesser, who so often seemed to have got there first, for discourse of unparalleled acuity and liveliness (his transformation into a businessman has severely diminished the intellectual life of this town); Peter Slepokura, for conscientiously reading the first draft from beginning to end; and finally, Linda Mason, for turning the final draft into a thing of pristine beauty. Some contributions of other individuals are mentioned at relevant points below. Needless to say, no-one but myself can be blamed for any of this.

February, 1985. London

Pre-Thesis: Mis-Words[1]

> Not strictly part of the thesis, this is a preliminary look at many of the problems it explores. Reading it will serve to orient the reader to those problems, although my views have changed in certain ways as they have developed.

The problem that vexes me in this paper was elegantly stated by William James in 1890. What kind of a "mental fact", he asks, is the intention of saying a thing before it is said? "As the words that replace the anticipatory intention arrive, it welcomes them successively and calls them right if they agree with it, it rejects them and calls them wrong if they do not. It has therefore a nature of its own of the most positive sort, and yet what can we say about it without using words that belong to the later mental facts that replace it? The intention to-say-so-and-so is the only name it can receive" (p. 253). Despite the popularity of this quotation with workers in the field, and despite much useful and interesting research, no-one seems to have tackled the James problem head on. Mechanisms for getting from intentions to formulations have been proposed, and assumptions made about the nature of those intentions: my aim here is to question some of those assumptions, and by doing so to shed some light on that problem, and on the nature of mind. It is widely accepted that the capacity to account for errors is the best test of theories of speech production, and that is what I look at below.

A very widespread assumption is that, to get from a pre-verbal intention to an utterance, the speaker must have

access to a "mental lexicon" in which the meanings of the words of his language are defined. The supposition is that the orderly display of word meanings and their interrelationships which the linguist attempts has - must have - a mental or even neurological equivalent. Most research in this area has assumed the necessity of such a mental lexicon, and some has even supposed that James's mysterious "intention to say so-and-so" can be replicated by definitions of the words to be said. The assumption of a mental dictionary rests in turn on a crisp distinction between knowledge of the meanings of words and knowledge of the world. As we see from the difference between dictionaries and encyclopaedias it is often practical to make such a distinction, but there are contexts in which it looks arbitrary.

On the face of it, slips of the tongue involving proper names are no different from other slips, and a mental search for the right proper name is just a special case of a search for the right word. But there are important differences, arising from the fact that proper names refer to individuals rather than - like, say, bachelor - sorts of thing. No matter how well organized a mental lexicon might be, it will not contain entries which define the meaning of proper names without recourse to descriptions of the world. To demonstrate that I know what "Cleopatra's Needle" means, I can either point to it, or tell you what I know about that object. "It's a tall, carved, monolith; stands on the Embankment near Temple Station; was brought there a long time ago", and so on, with progressive vagueness. The word "bachelor" is exhaustively defined by "unmarried man", that is its dictionary meaning and anything further about bachelors will

count as information about the world. But there is no justification with a proper name for saying, now it is defined, anything more is extra to its meaning. In a sense, proper names are not part of the language like other words - one looks them up in encyclopaedias, if at all, not in dictionaries. And yet, although proper names may not strictly be part of The Language, they are part of speech - and sometimes we get them wrong, or have to search for them, just as we do with other words.

Let us look at some examples of getting proper names wrong. At one time two publishers were putting a lot of work my way, one called Pat, the other Julia. Both were extremely thoughtful and kind to me personally, going out of their way to be helpful; unfortunately, I didn't like Pat at all. To my irritation, I kept thinking - and sometimes even speaking - of Pat as Julia. At first I was at a loss to explain this mistake, but then I noticed that it only occurred in certain contexts. When I was thinking of Pat as the helpful publisher who would go out of her way to make my life easier, I thought of her as Julia; when I thought of her as that unlikeable, boring person, I thought of her by her right name. I clearly did not wish to acknowledge to myself that I was taking advantage of the kindness of someone with whom I felt no personal sympathy. Or take another case. A and B are discussing the relations between a mutual friend, Hilda, and her flat-mate, Alice. A tells B how Hilda suffers from the constant comparison with Alice: Alice is a success at everything, she sails through exams, walks into jobs, attracts the most attractive men, and so on. B guessed who A was talking about, despite the fact that A referred to Alice as "Annie"

throughout. Annie was a colleague of A's, and a friend - also it now appeared, an envied rival.

A sort of proper name slip which must be familiar to most people is sibling confusion. Or, closely parallel, the black maid servant may be thought of as Viola because that's what th former black maid was called. As well as cases of substitutio of proper names, I have also collected some blends: a 4-year-old's "Little Red Muffet"; a grown man's "Das Kampf"; m own "Anger Under the Elms" and "Desire and After" (transpose t first words for the correct titles). What all these cases hav had in common, is that it is possible - with the relevant knowledge - to construe a uniting category for the target and error words, or for the blended phrases.

They differ from strictly semantic errors in that the relevant knowledge is not language-community wide, the uniting category not encoded in the language. Slips of this kind are r confined to proper names. A friend of mine with 5 children and cats explained the fact that my cats were sniffing the hem of h skirt by saying that it had "brushed against the children's feeding bowls" - when she meant, of course, to say "the cats'".

Cases like these suggest that the difference between dictionary and encyclopaedia does not derive from a dictionary' being restricted to 'linguistic meaning', but rather from its including only what we all know, if we know the meaning of the word at all. Only stable categories which hold good from any point of view get encoded in the language.

Thus, in the contexts in which I thought of her as Julia, the relevant distinguishing characteristics of Pat were the sam

as those of Julia; that is, they were both female editors who gave me work and treated me inordinately well. Or again, in the context of that conversation, Alice and Annie were both successful, envied, rival female friends, and so on.

So, what does all this mean for the assumption that speakers must have access to a 'mental lexicon'? That assumption is predicated on a conviction that, in order to find the words which mean what we want to say, - in order in fact, to know what they mean - we must have definitions of them in our minds. A good example of a semantic slip which fits this view, is shrinkled, a blend of shrivelled and wrinkled. "Shrivel" and "wrinkle" have obviously related meanings, with "shrivelled" entailing "wrinkled", and a lexicon would display that relationship. It seems a natural step to explain the error by the adjacency of the two words in a mental lexicon. Perhaps we can explain the proper name slips in the same way, by being more liberal about what can get into a dictionary.

Let us look more closely at the Pat-Julia case. Say they were both entered in my mental lexicon, marked as Work-giving, Female Kind Editor, plus some distinguishing characteristics each. When I thought of Pat as "Julia", I wasn't thinking that she was Julia. The wrong one was being mis-named by me; I did not have the wrong one in mind. Clearly, one has a hold of the thought-about object independent of the linguistic representation provided by its name. Is that hold on the object conferred by its definition? The problem is that for a proper name we must have a flexible open-ended definition, which cannot exhaust the meaning of the name, and can never, therefore, be an adequate

substitute for it.

Also, being flexible and people being what they are, everything about the "definition" can change, even the sex marking, and even its heading, the name itself. What gives the "definition" its identity through time is its being about the same object, which in turn depends on the thinker/speaker knowing what object it is that has these shifting attributes, independently of any particular description of it. It is therefore possible to know what some words mean - to speak and understand and think with them - without access to a definition of them. And we must explain some blends of similar meanings, and some cases of meaning-likeness between target and error word, without postulating a mental lexicon.

All this goes to show not only that proper names cannot be usefully accommodated in a mental lexicon, but also that a mental lexicon cannot be necessary for speech production. For, as I have remarked, proper names are parts of speech. A theorist who wishes to maintain a mental lexicon to explain meaning-related blends and target-error words, and the "tip-of-the-tongue phenomenon" (in which the searcher often comes up with semantically similar terms), must find a different explanation for the parallel phenomena with proper names. The existence of semantically related blends, etc., looks like very shaky grounds for believing in a mental lexicon. Indeed, the widely known fact that we blend sentences with similar meanings - which we don't, of course find ready-made - also requires an explanation which doesn't rely on prior adjacency and should already have cast doubt on the idea that semantic word blends, etc., constitute

evidence that we possess the mental equivalent of a dictionary. All the same, none of this goes to show that we don't have a mental lexicon, only that there are grounds for supposing we don't always need one.

Although nothing I have said has proved that we do not possess a mental lexicon in which most of the words of our language are explicitly defined, there are still further reasons for doubting its existence. As we shall see in more detail later, the idiosyncratic category relations we have found are not confined to proper names - the cats/children case is not a freak. Also, at least in the case of speech production, it is not clear what the explanatory value of a mental lexicon is supposed to be. If in the ordinary life situation of looking for the right word there is a floating definition waiting to home in on the mot juste, then the mystery is simply displaced. For it must be just the right definition for the word we want to say, and the question now is, how do we find that, and recognize it as right? Definitions of definitions would lead at once to an infinite regress. If a speaker's intention is sufficiently focused to pick out the right definition, why isn't it good enough to pick out the right word? As far as speech production goes a layer of semantic representations of word-meanings seems otiose.

* * *

Cases like those I have discussed, in which there is a discernible meaning-relatedness which is not strictly semantic, are usually referred in the literature as "Freudian".[2] The fact

that there are no satisfactory criteria for distinguishing a category of Freudian errors has been seen by some researchers as a difficulty in the way of studying them (see, e.g., Ellis). But, (a) it is just as much a difficulty in the way of studying semantic errors that there is no neat cut-off point between them and Freudian ones, and (b) uncertainty about how to categorize data does not prevent it from being data. As far as I know, Baars et al are the first and only researchers in the field to attempt to demonstrate the existence of Freudian errors experimentally.

With Motley, Baars devised a method of inducing speech errors in the laboratory. Their first experiments were with phonological errors: "The subject saw these word-pairs, one at a time, for 1-2 seconds each, GOOD BOY GO BACK GIVE BOOK BAD GOOF." Subjects produced the target error - GAD BOOF - significantly frequently with this phonological bias. They went on to use a semantic bias, such as preceding GAD BOOF with TERRIBLE ERROR, and again found a significant increase in spoonerisms - subjects were inclined to say "bad goof" in these circumstances. Next they attacked Freudian slips. They prepared lists of word pairs in which half the target spoonerisms were electric shock related, such as SHAD BOCK/BAD SHOCK, while the other half were related to the sexually attractive properties of females, such as LICE NEGS/NICE LEGS. All the subjects were male, and they were divided into three groups: "one group was told that it might receive electric shock during the experiment while another group had a very attractive female experimenter, provocatively dressed, and a control group received neither treatment" (p. 24). The

results were very satisfactory, with the appropriate groups doubling the relevant error rate.

The explanation Baars offers for the brilliant success of his experiments he calls "the competing plans hypothesis". Thus, in the semantic case, reading the words terrible error primes an intention to express that idea, which competes with the plan to speak the target words. The plan to express the idea of a frightful mistake is not a complete or closed plan, or BAD GOOF wouldn't do. A plan, Baars says, is "a representation of a reasonably complex action, existing prior to that action, and feeding into some set of programs and sub-goals which can carry out the action in detail" (p. 7). The function of a plan, then, is to set some programs going with the purpose of carrying out an action; the details are specified in the programs, not in the plan.

The idea that slips of the tongue are caused by competing plans, is a stronger version of Freud's own hypothesis that they are caused by "simultaneous excitations". Although many slips can be explained by the competing plans hypothesis, it seems to me that there are reasons for preferring a weaker formulation. When Mrs. Thatcher (allegedly) said, "Yes, yes, we'll discard the matter in committee", her plan to do just that inappropriately triumphed over her plan to claim "the matter" as a topic to be pursued. And "What a good colour these trousers were" is a result of competition between "what a good buy they were" and "what a good colour they are". But there are a number of cases in which, though the notion of competition is explanatory, it does not seem to be competing plans which cause the slip.

Reading a children's tale out loud, I spoke of the pirates - returning to their ship after an all-night party - as "rolling to the ship", when it should have been "rowing". A thought about the drunken pirates had intruded, but I certainly had no plan either to say or do anything about them. Or when someone wrote "regular week-end gatherings of foster parents and their children will be helled", it was surely because a thought about what those weekends would be like had intervened, rather than any kind of plan.

Even Baars's own experimentally induced Freudian slips don't seem to me well explained by his hypothesis. Baars likens the way a plan sets programs going to the way a general commands his subordinates, leaving the details of execution up to them. But in this case, what is the general telling them to do? Presumably he is commanding them to <u>do something about</u> the threatened shock or the attractive lady. A plan is usually a serious intention to carry out a specific action, achieve a particular end. We have no reason to suppose that the subjects had such serious intentions, though it is evident that they were thinking about the electric shocks or the luscious experimenter. Under the circumstances, they had no chance to do anything constructive - such as flee the lab or make a pass - but the word-pairs offer a chance to do what they can to express their thought. The moral seems to be, those that can, do; those that can't, talk.

So, the picture we have so far is of competing thoughts about this and that: thoughts which are struggling to be expressed (either in action or in words). Baars says that "the competing plans hypothesis leads to a conception of normal

error-free speech as inherently involving a consideration of alternative plans for production which are reviewed and edited prior to speaking" (abstract). The competing thoughts hypothesis leads to no such conception. On my view it is only when it is what to say that we are thinking about, that we need to "consider alternative plans for production". Otherwise, which thoughts achieve expression when, depends on how pressing are the interests which give rise to them. It is the will rather than the faculty of reflection which determines which thoughts triumph by being expressed.

Although the idea of competition is helpful in explaining blends, spoonerisms, anticipations, perseverations, etc., there is a class of slips to which it seems irrelevant: those in which there is no serious sign of a competing alternative. Such cases include pounds for inches, number for letter, misprint for slip of the tongue. Here the speaker has produced the wrong member of the right category, and there is no evidence that the right member was considered and rejected. These cases are more or less strictly "semantic". Then we also have cassette for transistor (radio), zoo for fair, motivation for compensation, in which the connection is less obviously semantic, has more, perhaps, to do with the lives the speakers lead. And at the idiosyncratic extreme there are the proper name cases discussed above, and slips like children's feeding bowls.

The Freudian idea of a repressed element struggling to make itself known may apply in one or two of these cases, but is obviously not generally explanatory. To look for a good general explanation, one should first find a good general description of

what is to be explained. I have spoken of "connecting categories": what is a category? As I have been using it, it could be anything from a natural kind - which holds universally - to the-objects-on-this-table - which form a coherent lot only here and now in this room for these people.

In fact, none of my cases is at either extreme, but they lie along the continuum between. I have described the middle set of errors above as "having to do with the lives people lead". The speaker looking for his transistor radio while announcing a search for his "cassette", used both these objects to provide portable, electrical, entertainment - which was what he wanted the radio for. The speaker who spoke of the fair as "the zoo" would use either as a place to take the children and have a good time. In the last set of cases, the relation between the error and the lives of individuals is even more obvious. Perhaps rather than "uniting categories", we should be talking about "roles played in life". But what of the language-encoded categories at the semantic extreme? Should they, also, be described this way? Whether they should be or not, they certainly could be. The difference between them and the "pragmatic" cases being that in the case of, say, pounds and inches they both play the role of unit of measurement in anyone's life (when they play a role at all).

So, I am suggesting that each of these cases of misnaming can be partly explained by the fact that the objects represented by the confused words at the time of speaking play the same (though not in all respects the identical) role in the speaker's life. What remains to be explained is why there should be errors

at all. The speaker is not actually thinking of the wrong object, or it would not be the wrong object; so how does the name of a wrong object - one not thought of - come to mind? Let us look again at the Pat/Julia case. Recall that it was only in certain contexts that I thought of Pat as "Julia", namely those contexts when the common role both played in my life were what mattered. Now, concerns, matters, interests, are the active ingredients of the mind: they are what account for particular intentions or thoughts. It is, I propose, this active nature of a concern which must explain the fact that the wrong name gets "activated".

We now have a basis for integrating the explanations of slips arising from "competing thoughts", and those I have just been discussing. In the first place, meaning-related blends like shrinkled, clittering up, slogging himself, husculine blend alternative expressions of the same concern (the wrinkles, the mess, the effort, the butch male); and the same is true of sentence blends like what a good colour these trousers were, and it's handy living on a North/South access. And then there are cases in which a different concern intrudes, reveals itself, such as Mrs. Thatcher's (alleged) discard the matter in committee, and Baars's Freudian result.

The picture of the mind which emerges is one of seething activity, arising from the many interests or concerns of the thinker/person. These concerns may be immediate and fleeting or persistent and inexhaustible, trivial or weighty, weak or strong. What they are like depends on the relations between their objects and the person concerned: what they have in common is that they

give rise to particular thoughts or actions about those objects. We count as thinking or talking about something if we make a predication of it; whether we are content with one meagre predication depends on how interested we are. Even in the near mechanical case of reading a familiar tale out loud, the reader was for a moment interested enough to think beyond the given words to what the pirates and their party were like. Anything relevant to a concern may be activated by any other relevant thing (word, sentence, object, concern).[3] In the light of a concern, i.e. from the point of view conferred by it, differences between things may be, or seem to be, irrelevant.

Given this picture of the mind, what does James's "intention to say something before it is said" look like? I see an intention as a particular product of some interest: interests are fulfilled, intentions expressed, in deeds as well as words. An intention to say something before it is said is exactly as mysterious as an intention to do something before it is done (very), and we can no more say anything without words about the one than we can about the other. The peculiarity of the intention-to-say cases arises from their end, not from their initiation.[4] There is a further peculiarity in cases where doubt arises about what to say or do, for when the doubt is about what is to be said then what to say itself becomes an object of concern. It may seem from these remarks that I am identifying mind with will, and making language a slave to thought; these inferences are contradicted by my forthcoming thesis.

Chapter One: Introduction

Writing this book has not been a matter of displaying a pre-ordained structure, of following a route clearly mapped out in advance. Only its aim has been clear from the start: I have wished to understand as well as possible the relation between language and thought. Pursuing this aim has involved taking as little as possible for granted, and having to count on clues, hunches, luck. A recurring theme of the thesis is that thinking is not a matter of finding truths and trying them out. An interplay between certainty and uncertainty, between what may and what may not be taken for granted, is a central characteristic of serious thought which necessitates a reliance on guesswork. The thesis itself illustrates this theme.[1]

My central assumption, obviously enough, has been that there is a relation between language and thought. That I have not questioned. Like most people who've thought about it at all, I've felt sure that some such relation must obtain. Yet it cannot be more than a hunch while we cannot say what the relation is. The reason it's so hard to say what it is, is that it's so hard to separate language from thought, so hard even to discern where the boundaries lie. No direct attack seems possible, so alternative approaches must be sought by anyone trying to tackle the problem. My own approach may appear perversely circuitous all the same, and requires some explanation.

Both psychologists and philosophers may feel I've plunged heedlessly through their territory, with no regard for the local

customs, without bothering to learn the language, and even failing to pay tribute where it was due, noticing only what suited me. Unfortunately there is a good deal of truth in the accusation. My general defence is that, because of the aim I wa pursuing, I have had to pass through so many areas with differen customs and languages, that learning them all would have been an impossible task for anyone - certainly it would have been for me And where I've failed to pay tribute it's because I haven't realised it was due. Had I paused to become expert in any of those areas (e.g. psycholinguistics, memory theory, learning theory, "cognitive science", imagery, "inner speech", theory of mind, theory of meaning...), I could have stayed there for the rest of my life. That said, I have in fact tried to learn as much as I could by dipping as judiciously as possible into the relevant literature, counting on luck, friends, and good guesswork to guide me.

In general, the reader should remember that my quite eclectic and wide-ranging bibliography indicates extensive dipping and not breadth of learning. It should be treated as a guide to further relevant material rather than as evidence that I've digested that material myself. I stress this point because anyone who forgets it is liable to think I'm attacking positions I've never come across and addressing issues I've never discerned.

On the whole, I've found that the method of psychology condemns it to dealing with very small issues, from the perspective of the laboratory. But once I'd developed a general view of my own based on observation of human life at large, I

looked for, and found, support in experimental psychology. Philosophy cannot be used in the same way, since its substance is debate. Some of the issues I address below have been live since Plato or before, and a full understanding of current comment would entail a grasp of two thousand years of argument. Despite my scant knowledge of that debate, I have presumed to argue with two widely read and influential philosophical papers: Grice's on conversation (1975) and Putnam's on meaning (1975). In each case, doing so involved formulating an initially inarticulate dissatisfaction with what they had to say, and was extremely helpful in clarifying my own thoughts.

Although I take issue with much of what Grice says in 'Logic and Conversation', I wholly concur with two of the points he makes. One is that the maxims which guide conversation are a special case of those which guide all cooperation between people. The other is that accounting for relevance in a conversation must involve the interests and desires of the people concerned. That first point is a particular instance of what in retrospect I realise to have been a crucial strategy in my own endeavour: namely, that of treating language as a special case of some more general phenomenon wherever possible. In this way, what is peculiar to language, what distinguishes it, may become clear. I was careful, for example, when defining relevance, to give it its most general possible definition and not restrict it to a linguistic application. Actions, pictures, tools, people, etc. can all be relevant or irrelevant as well as utterances. Forging a definition for all those cases involved bringing in interest, which is a crucial factor across the whole range. In the long

run, what may have been for Grice a quite subsidiary point, has been a dominating theme of the thesis: namely, that understanding why something is relevant to a person will entail attributing an interest or interests to that person.

As the last paragraph suggests, the question of relevance has proved a central concern in this thesis. But why? Why relevance, what has it got to do with the question of the relation between language and thinking? Historically, the fact is that I didn't know where to begin in tackling the major question. But I found the notion of relevance (and interest) essential in my explanation of slips of the tongue, and all Grice's conversational maxims turned out to follow from the obligation to be relevant. Furthermore, my own supervisor, Deirdre Wilson, was treating relevance as at the core of the pragmatic theory she was developing with Dan Sperber. But there is presumably a reason for this history, an explanation for why relevance appears centrally in all these contexts. From the perspective I have now, I can see that it must be because coherence and intelligibility are prerequisites of both effective talking and effective thinking. And the underlying principle of both coherence and intelligibility is relevance.

The coherence of any collection or sequence depends on whether its parts are relevant to the same interest(s), and if so, how. And only what is coherent can be understood. Since 'interest' must cover the range from a merest flicker to a raging obsession, it follows that reason depends essentially on unreason. It also follows from the interest-dependency of intelligibility, that there must be an intimate relation between

relevance, interest and meaning. In exploring that relationship I once again adopted the strategy of considering the topic - meaning - in its broadest context before narrowing the analysis to linguistic meaning.

Having considered some of the ramifications of non-linguistic meaning, I turned to Putnam's views on meaning in language, as expounded in 'The Meaning of "Meaning"'. There I found a great deal to agree with, in particular that linguistic meaning is essentially social, and that identity is interest-relative. In my view, the defects of his position derive principally from his tendency to treat society as basically homogeneous, its constituents divided only by degree of scientific expertise. Hence he fails to allow for there being distinct communities of interest within a society, each with its own expertise and its own capacity to establish linguistic meaning. Furthermore, he seems to forget that societies are ultimately composed of individuals, and claims that because linguistic meanings are social, they are therefore "not in the head". In fact, the opposite conclusion surely follows, even without bringing in considerations concerning meaning in general. And my main finding about meaning in general was that it is interest-dependent. Interest being a primarily psychological characteristic, meaning is always at least partly "in the head". What distinguishes linguistic meanings is not that they are "not in the head", but that they are in many heads, and therefore beyond the control of the individual. It follows from this, inter alia, that the relation between language and thinking is in part a relation between a society and the individuals which

compose it.

Parallel to that distinction between language as primarily social, and thinking as primarily individual, is the fact that while speech is a public affair, thinking at least begins as a private one. This makes thinking harder to investigate. How can I find out about another person's private world, unless they choose to disclose it to me? And if they do, then the tool they use to reveal their thoughts to me will of course be language. It has seemed to me that the only way out of this bind was to have recourse to introspection. A possible alternative strategy would have been to scour the literature for references to people's inner lives - examples abound. But that might have been deemed just as dubious by anyone who disapproves of non-experimental psychological inquiry, and it would have spoiled my pleasure in a good book. As a psychological tool, introspection has been out of favour for most of this century, for the good reason that its data cannot be checked. But the methods of the laboratory, although they yield results which can be checked and double-checked, and turned into the most rigorous statistics, seem incapable of yielding any but the most trivial insights into the inner life of the mind. Anyway, experimental results can only provide a basis from which <u>inferences</u> about inner processes can be made (just as everyday behaviour provides such a basis). However careful the experiments, there is no empirical guarantee of psychological conclusions drawn from them.

Even given the drawbacks of experimental psychology, doesn't the uncheckability of introspection completely disqualify it as an alternative? Even supposing that I am a reliable witness to

the contents of my mind, nothing follows about other people's minds from my findings about my own. Firstly, I would reply that any theory of mind should be capable of accounting for mine as well as for any other. Secondly, the findings of introspection aren't as inaccessible to other people as all that - anyone can do it, i.e. introspect, with a little effort. Insofar as other people's introspection yields results which parallel my own, then my findings are supported. At the beginning of Chapter 6 I describe the method I've used, and there and elsewhere I urge certain easy mind-games on the reader in the interests of providing empirical support for my claims.

Having sketched some of the general peculiarities of my approach, I shall outline the form of the thesis, chapter by chapter. The first chapter proper is the debate with Grice I discussed above. In it I conclude that in the right circumstances, the obligation to be informative follows from the obligation to be relevant, while the obligations to be truthful, orderly and succinct all follow from the obligation to be informative. So, for understanding living speech, understanding relevance would be essential. In the next chapter, I attempt to analyse the notion of relevance. I conclude that, to be sufficiently general to cover all cases, relevance need not entail informativeness, although it often does. The fundamental prerequisite of a relevant remark in a conversation is that it should contribute to a common interest or concern (clearly a special case of a very general social rule). A remark may do that by arousing, expressing or making a difference to the outcome of a common interest or concern. Providing information

is a special case of making a difference to the outcome of an interest. But relevance also has a non-social application, for example, one train of thought may be relevant or irrelevant to another. To generalise the account to these cases, private as well as common interests must enter the picture.

In the next chapter, I apply my findings about relevance to some samples of actual conversation. I find that the obligation to be (relevantly) informative applies when there are perceived gaps in the fabric of mutual certainty within a community of interest. Such gaps arise from doubt and ignorance relative to common interests, and appear as issues and questions in discourse. At any given moment, the background of relevant mutual knowledge for understanding an utterance will be determined by what has made a difference to the current form of the interest(s) in play. Contributing new information entails changing that background by affecting the interests of the people concerned. Mutatis mutandis, these conclusions apply to personal as well as common knowledge.

My next move is from the external realm of public space to the internal world of the mind. My first aim was to find out what sort of psychological phenomenon an interest is. I concluded that as well as being a repository of information, an interest is a productive force from which thoughts and images as well as speech and actions, flow. We all have a great many interests at varying degress of arousal. The more highly aroused, the more likely an interest is to produce some activity either in the external world or in what I've called 'cognitive space'. Unless and until an interest is satisfied (a desire

fulfilled, an issue resolved, etc.) then activities produced by it serve to maintain its level of arousal, by feeding back into it.

The imagery we produce only feeds back reliably into interests which give rise to it, though it may sometimes arouse other interests, and thereby bring other information into play. But our internal verbalisations inescapably activate interests beyond those which gave rise to them, because they connect with the language system, which is independent of any individual's interests. Certain further effects of verbalisation follow, amongst the most salient are: firstly that it is possible to exercise some control over the level of arousal of one's own interests; secondly that minimum effort gives access to maximum information; thirdly that one can make judgements in a relatively 'objective' fashion; fourthly one can fall prey to self-delusion; fifthly one can engage in lengthy internal debate.

Since meaning is clearly a central feature of language, which intertwines with questions about relevance, that is my next topic. Of the next two chapters, the first is devoted to non-linguistic meaning, and the second to the debate with Putnam's views mentioned above. Meaning turns out to be not just connected with relevance but actually to be the capacity to be relevant. The three aspects of meaning, what N means by X, what X means to M, and what X means, can all be stated in terms of what interest X expresses or arouses, or what difference X make to it.

Where X has a conventional meaning, then there is a fixed relation between what it expresses and what it arouses; these

will at least overlap. And linguistic meaning has the facility not only of hitting a precise spot in a set of interests, but through the sentential form, of specifying or requesting a precise modification, i.e. of conveying or requesting exact information. The chapter on Putnam narrows the focus to word meanings. In it I conclude that the extension of a term 'X' is whatever we people call 'X'. Its intension is the current form of the interest that bears that name, that is reliably aroused by that name. Its definition is a description of the features that distinguish that form, such that all those within some community of interest who use that word will recognise those distinctions. Psychologically, word meanings do not have an existence independent of the whole set of a person's interests and assmptions, but are a relatively fixed and rigid subset of them.

In the next chapter I conclude that neither interests nor cognitive space are uniquely human. A set of interests, in the sense I have been using the term, is attributable to all higher animals. The outmoded psychological concept of 'drive' was a crude attempt to capture the phenomenon of interest: it is interest that gets the hungry rat through the maze. Learning in all creatures capable of it consists in the modification of interests. But, thanks to the stable and well-motivated connections and distinctions given in any language, and the potent and reliable call on our interests exercised by words, human beings acquire and have accessible a vastly greater range of information. The capacity to project material into cognitive space, to recall the past and anticipate the future, is almost certainly not exclusive to the human race. But we alone can

imagine ourselves speaking, and listen to what we say - "Language is a tool for amplifying the imagination".[2] Through language we can detach ourselves from our own private interests and assess our possible actions and utterances from the point of view of others' interests. Cognitive space is the playground of the imagination, home of the absent mind, rehearsal room for the drama of life, and courthouse of the 'superego'.

In my conclusions, I attempt to translate my findings into the vocabulary of some contemporary cognitive theories. 'Propositional attitude' and 'interest' as I have used it, cover the same range, with the important exception of 'beliefs' in the sense of assumptions, which form the background of all our attitudes. Far from assumptions being couched in a 'language of thought' which is more objective and context-free than natural language and subsisting independently of our interests, an assumption is just a difference which has been made to the current form of an interest. The form of an interest, its 'schema', is constituted by a set of assumptions; its content is the energy which activates it. So an assumption is an effect on the future expenditure of energy. The system of richly interconnected interests I envisage is roughly equivalent to the network of schemas assumed in some current theories - though interest is crucially more general than purpose. But, spreading activation inescapably depends on what matters to the individual as well as on the prior existence of connections, and those connections must include the ad hoc and ad hominem as well as those laid down in the language. Information acquired relative to one interest may affect other related interests directly, or

it may do so via formulation and the quest for truth - which is not automatic.

This view of 'mind' as consisting of two radically distinct components, one - the network of interests - postulated to account for observable behaviour, but not itself observable, the other - cognitive space and material therein - itself susceptible to observation, has evident parallels with Freud's The Unconscious and The Conscious. But I hold that he delimits that distinction wrongly. Despite the apparent centrality of desires or drives in psychoanalytic theory, as much as current cognitive theory does it rests on the false premise that 'rational beliefs' and desires inhabit different realms. In my view, the unconscious network, through its capacity to be modified, is the instrument of all learning and the repository of all information. And it has as an extensive subpart the self-consistent and well-motivated semantic system of an individual's language.

Insofar as an interest system constitutes a world view then in being partly comprised of the semantic system of a language, what language that is will affect the individual's world view. However, pace Sapir, language can exercise no 'tyrannical hold' because it structures only part of an interest system. Meaning always goes beyond the strictly language-given. Furthermore, the structure of the semantic system itself depends on common interests, and each of us can affect those, and thereby affect the available meanings within a community of interest, a 'society'. Speech is a tool for manipulating interest systems. Language is one of society's most effective ways of owning its members; but my interests don't cease to be mine when society

owns them.

Language serves thought by providing a vast reliable range of information which is generally effortlessly on call. Speech serves thought by freeing it from the interests which give rise to it.

Chapter Two: Conversational Rules

The most influential work that has appeared in the last couple of decades on conversation is that of the philosopher, H.P. Grice, in his paper called "Logic and Conversation"[1] (1975). In it, he makes the following general comments on conversation:

> Our talk-exchanges do not normally consist of a succession of disconnected remarks, and would not be rational if they did. They are, characteristically, to some degree at least cooperative efforts; and each participant recognizes in them, to some extent, a common purpose or set of purposes, or at least a mutually accepted direction ... at each stage, some possible moves would be excluded as conversationally unsuitable. We might then formulate a rough general principle which participants will be expected (ceteris paribus) to observe, viz: "Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk-exchange in which you are engaged". One might label this the Cooperative Principle.

Although conversations are rational, the sorts of inference and implication participants constantly make go beyond what would be strictly justified by the logic of what gets said. Grice's maxims, with which he spells out the Cooperative Principle, are meant to explain how we derive "implicatures", those implications with no basis in a truth-conditional logico-semantics of what is said "in the strict sense". In other words, those aspects of interpretation which do not follow from a context-free grammatical combination of verbal meanings.[2]

The maxims can be summed up as:

1 a) "Make your contribution as informative as is required (for the current purposes of the exchange).

b) Do not make your contribution more informative than is required."

2) "Try to make your contribution one that is true".

3) "Be relevant".

4) "Be perspicuous".

These may sound more like advice on how to write an examination than maxims which are normally observed in discourse. Grice himself says that he has stated his maxims "as if the purpose (of conversation) were a maximally effective exchange of information: this specification is, of course, too narrow, and the scheme needs to be generalized to allow for such general purposes as influencing or directing the actions of others". He also expresses some particular reservations: he thinks that (3) "Be relevant" may render (1b) redundant: he wonders if (2) shouldn't be given a status elevated above the level of the other maxims: and he finds the question of relevance highly problematic - as, indeed, it is. I take the maxims one by one below, and see how they work in the analysis of actual conversation.

* * *

What is it to be as informative as is required for the current purposes of the discourse?

Example 1

1 a. Wimbledon's got a fair share of amenities. Sandy

Denny. Virginia Wade went to my secondary school.

2 B. This is tremendous.

3 A. Celebrities.

4 B. I can offer you Eleanor Bron, Esther Rantzen.

5 A. Eleanor Bron?

6 B. Yeah.

7 C. Where did you grow up?

To inform means to pass on knowledge, normally with the presumption that others are ignorant or doubtful of that knowledge. How much of this snatch of real discourse counts as informative by those criteria? Only 1, 4, and 6 are relatively straightforwardly informative by this definition. At 1, A proposes that Wimbledon has a "fair share of celebrities", and goes on to offer instances of her proposition. At 4, B instances two celebrities from her milieu, and at 6, she removes Eleanor Bron from the realm of uncertainty in which A's question has (ostensibly) placed her: Yes, Eleanor Bron _is_ one of the celebrities I offered (in 4).

On the face of it, questions are requests for information, rather than themselves informative utterances. However, in most contexts, the standard question forms convey "X wants to know..." to the listener, and indeed, "X wants the addressee to provide the answer if possible". The addressee is thereby informed, comes to know, that X wants to know..., and that X is hoping the addressee will supply the answer. So, with 5, A conveys to B that she would like to know whether Eleanor Bron was one of B's examples. And, with 7, C lets B know that she wants to know where B grew up (and where, therefore, Eleanor Bron and

Esther Rantzen grew up).

That leaves only 2 and 3 still lacking any obvious informative content. Are A and C informed of anything by B's "This is tremendous"? It seems from A's response that she gathers from it that B takes an ironic view of celebrities/of boasting about celebrities/ of Sandy Denny and Virginia Wade as celebrities. A's "Celebrities" indicates (informs B) that she acknowledges celebrities a topic for sardonic comment. In the context of this analysis, it becomes clear that 4 will yield more information than I attributed to it above.[3] In producing her own couple of celebrities, B is letting A know that she isn't so contemptuous of boasting about celebrities as not to do it herself.

In Grice's view, if an utterance violates one of the maxims of conversation at the level of 'what is strictly said', the hearer will seek to construe it as really fulfilling the maxim at some other level. There is some doubt as to whether the information (N wants to know, etc.) imparted by a question counts as subsisting in what is strictly said. My own inclination would be to say, Yes, if there is such a level; but Grice, I think, would say No.[4] Assuming Grice is right, questions join utterances 2 and 3, in prompting a search for information implicated but not uttered. Thus, C wants to know where B and Eleanor Bron, and Esther Rantzen grew up, is implicated by 7; B implicates that she takes an ironic view of celebrities, etc.; and so on. Only 4 of example 1 presents a problem: it does not violate the maxim of informativeness, yet it carries information over and above what it strictly says. In Grice's terms, it

implicates that B isn't that contemptuous of boasting about celebrities. And that is clearly information appropriate to this stage of the discourse, just as is the information imparted in what was strictly said. It seems that we may look for information "over and above" what is strictly said, whether or not the maxim has been violated, flouted, or disregarded.

So far, I have been discussing information in itself, rather than information for the current purposes of the discourse, as required by the maxim. In order to see what those "current purposes" may be in example 1, some background must be filled in. The parties to the discourse are three fellow students who know each other only moderately well. They have met in order to "pass the time of day", to chat, to get to know each other better. They have talked about travelling abroad, about moving away from London, B and A have both said they've lived in London all their lives; B has asked A whereabouts in London, and been told Wimbledon; B (knowing that A and a senior colleague N had gone to the same school) remarks that N must also have grown up in Wimbledon, and that is where example 1 begins.

So, what are "the current purposes" of this discourse? Its general point is idle chat, getting to know each other. Of such talk-exchanges, Grice says they have "the second-order purpose that each party should, for the time being, identify himself with the transitory interests of the other". That characterizes this conversation nicely, except that here the people talking also have a common interest in finding out about each other, about each other's origins, history, opinions, hopes, etc.

In line 1, A proffers some information about Wimbledon and

its celebrities. The parties to the discourse have just learned that Wimbledon was where A grew up, so they can construe 1 as information about A's background (why else would they want to know about Wimbledon?). In addition, since C has just reminded them that their senior colleague, N, comes from there, too, it is informative about her background also - in particular, it is about the school N as well as A went to. And, it may also be taken to indicate that A regards N as a celebrity. And B's ironic comment (2) can be seen as reflecting her opinion of N's status as celebrity.

Line 1, then, being about A's background, is legitimized by the general aim of finding out about each other, which we can take to be current throughout this discourse. And it has the added bonus of being about a mutual colleague - we may presume that finding out about mutual acquaintances - i.e. gossip - is a possible aim throughout such conversations, too. Celebrities are mutual "acquaintances-by-fame", and if the conversation had gone on to focus on them, to become gossip about them, then I think that would also have fulfilled the maxim.

Although 1 was transparently about Wimbledon, if A had gone on to tell the others that Wimbledon could be reached by British Rail or underground, was served by the 74, 93, 85, and 39 buses, was abutted by Richmond, Kingston, etc., then they would have been hearing a good deal more about Wimbledon than they had any wish to. At that stage in the discourse, this information was not "required for the current purposes of the discourse". It would therefore have violated that maxim ('a'. And it would also have violated 'b' - do not give more information than required.

If one is as informative as required, one is thereby not giving more information than required - maxims 1a and b do not seem to be distinct in application.

Wimbledon is not a subject-matter that is likely to be fruitful in this conversation, the parties are not interested in being informed about it. Subject-matters which do concern the parties in this discourse happen to be, above all, people.

A great deal of information about mutual acquaintances would be accepted as "required" for current purposes. Being as informative as required is clearly a matter of degree relative to a subject-matter. Which surely means that it's a matter of relevance: the information about public transport in Wimbledon is unwanted because it is irrelevant - nobody wants to go there.

It seems to me that both parts of maxim 1 are "taken care of" by the maxim "Be relevant". As an example of 1a being violated, Grice tells the following story:

> A is writing a testimonial about a pupil who is a candidate for a philosophy job, and his letter reads as follows: 'Dear Sir, Mr X's command of English is excellent, and his attendance at tutorials has been regular. Yours, etc.' (Gloss: ... A cannot be unable through ignorance, to say more, since the man is his pupil; moreover he knows that more information is wanted ...).

Grice labelled this maxim (Be informative) "Quantity" - but, as this example shows, it is not a simple matter of quantity. The

letter gives the wrong information for the purposes it is meant to fulfill, it is qualitatively wrong. Had he gone on to add, "furthermore, he is the most impressive philosopher I have ever encountered," the beginning information would have looked still odder. On the other hand, if the whole letter had consisted of, "Dear Sir, Mr X is the most impressive philosopher I have ever encountered, Yours, etc.", then it would have been as informative as required - what more could a referee say in recommendation? Yet it would have been less informative than Grice's imagined letter. The point is, surely, that A has not put down what he knows is wanted for the purposes of the letter-exchange. The central question being how good a philosopher X is, what A has said is irrelevant. If my conclusions so far are right, we have one major maxim less, and twice the weight placed on the admittedly problematic maxim, Be relevant.

Grice's next maxim (as it transpires, the first) is "Try to make your contribution one that is true". He spells this out further as, "Do not say what you believe to be false" and "Do not say that for which you lack adequate evidence". The information conveyed by an utterance is to be genuine information, not mis-information. The maxim therefore only has immediate application to those utterances I described as "straightforwardly informative": 1, 4, and 6. There is no reason to doubt that those three utterances are true. At the level of "what is strictly said" none of the other utterances in example 1 is assessible for truth. On the other hand, given that every utterance there turned out to be informative at some level, they can presumably be assessed as true in terms of whether the

information they impart is true.

Grice anticipates that conspicuous violation of this maxim will yield an ironic interpretation. Although 2, "This is tremendous" is formally a statement, the gloss is: B cannot really hold that Sandy Denny, Virginia Wade, and N, make up a "tremendous" list of celebrities; she therefore intends the remark to be taken ironically, and not as a genuine opinion. Her statement, though not a true one, thus implicates her true opinion - and A's response to it (which in itself is neither true nor false) implicates that she is prepared to endorse that opinion, and that implicature can be assessed for truth, or at least sincerity.

Questions, as we have seen, can be regarded as standardly imparting the information X wants to know ... C's question at line 7 expresses a genuine want of knowledge, but A's "Eleanor Bron?" does not. A does not, as she seems to, want to check whether the said lady was one of the celebrities mentioned by B, she is in no doubt. Eleanor Bron's name had been articulated with perfect clarity and audibility, other parties to the discourse therefore had grounds for deciding that A's question was spurious, and thus for seeing some explanation other than knowledge-wanting. Immediately after B has answered C's question "Where did you grow up?", A chimes in with, "Someone told me I looked like Eleanor Bron the other day". The point of the question was to signal that she'd like to expand on the theme of Miss Bron (as a way of saying something about herself). That is what she'd truly like to do, that wish is genuine.

Grice finds parallels to his conversational maxims in the

way all social exchanges are organized. The general equivalent of the maxim of truthfulness he puts as follows: "I expect your contribution to be genuine and not spurious". The truth of a question cannot be assessed as such. What can be assessed is whether it is true or false that the asker wants to know... If that 'information' is false then the question is not a genuine, but a spurious, question. Similarly, B's "This is tremendous" was a spurious assertion, but genuinely ironical. In applying this maxim, then, talkers must judge whether the information that seems to be imparted is true or false information, whatever level of analysis yields the information in question. All this raises the question of whether "Be truthful" is itself distinct from "Be informative". One is inclined to say that if a statement is false it cannot be informative, but can only purport to be informative. For one would deny that someone who had been misled into believing p, where p is false, had thereby come to know p. And to inform is to make known. Rather than having a status elevated above the others, this maxim seems to be subsidiary to the requirement of informativeness. And, I have argued, in its full form that maxim itself derives from the requirement of relevance.

Be Relevant is the next maxim; succinct as it is, Grice gives little guidance as to what he means by it.

Some clues are available, though. In his list of non-talk analogies to the conversational maxim, he gives the following parallel to relevance: "I expect a partner's contribution to be appropriate to immediate needs at each stage of the transaction; if I am mixing ingredients for a cake, I do not expect to be

handed a good book, or even an oven cloth (though this might be an appropriate contribution at a later stage)". So, in non-talk transactions, the equivalent of relevance is appropriateness to a need, fitness for the immediate purpose. Now look at his non-conversational analogy to informativeness: "If you are assisting me to mend a car ... if, for example, at a particular stage I need four screws, I expect you to hand me four, rather than two or six". In each case, what goes wrong is that the supposed helpmate fails to provide just what is wanted at a given moment. The difference between the analogies, namely that one brings in quantity, has no application in talk-exchange (though it does in exams - you don't answer two or six questions if you're meant to answer four). The only way informativeness for the current purposes of the discourse can be measured is via relevance. And degree of relevance is not a simple matter which can be quantified like nuts and bolts.

Grice gives us three examples of (imaginary) discourse to illustrate this maxim.

1) A: I am out of petrol.

B: There is a garage round the corner. (Gloss: B would be infringing the maxim 'Be relevant' unless he thinks, or thinks it possible, that the garage is open, and has petrol to sell; so he implicates that the garage is, or at least may be open, etc.)

2) A: Smith doesn't seem to have a girlfriend these days.

B: He has been paying a lot of visits to New York lately. B implicates that Smith has, or may have, a girlfriend in New York. (A gloss is unnecessary in view of that given for the previous example).

In both examples, the speaker implicates that which he must be assumed to believe in order to preserve the assumption that he is observing the maxim of relevance.

He says of these examples that they violate no maxim, "or at least ... it is not clear that any maxim is violated", and they are the only examples he gives of (non-conventional) implicatures arising when no maxim has been violated. He remarks that cases in which "an implicature is achieved by real, as distinct from apparent violation of the maxim of Relation [relevance] are perhaps rare". Here is the example he gives of such a violation: 3) At a genteel tea party, A says Mrs X is an old bag. There is a moment of appalled silence, and then B says The weather has been quite delightful this summer, hasn't it? B has blatantly refused to make what HE says relevant to A's preceding remark. He thereby implicates that A's remark should not be discussed and, perhaps more specifically, that A has committed a social gaffe.

The implicatures arising from violation of the maxim of relevance are of quite a different order from those which arise in (1) and (2) immediately above. There, the implicatures were a spelling out of the relation between one utterance and its successor. Here the implicatures derive from the absence of any such relation. B rejects Mrs X is an old bag as something to be discussed, as a possible topic. The element of blatant refusal present in this story is not essential to cases of violation of the maxim. Cases in which a person makes a remark which is not relevant to its predecessor abound in normal conversation. In each such case, the person whose remark is not immediately

relevant indicates a disinclination to pursue the topic of the previous discourse. Thus, in my own example 1, C's question, "Where did you grow up?" lets her audience know that she does not want to natter on about celebrities, but wants to change the subject. The peculiarity of Grice's case is that "Mrs X is an old bag" is a remark which invites comment - unlike, say, "Yes, well, I'm with you there, I think we all are" or "mmmm". As Grice puts it, the problem is, "how to allow for the fact that subjects of conversation are legitimately changed". A definition of relevance which makes this allowance is required. At the moment I wish only to remark that calling the connections we make to preserve relevance by the same name "implicatures" as the inferences we draw to account for its absence, can only obscure the issues.

Example 2

1. A. What do you think of M, as people who don't know her very well?

2. B. Do you want to turn that off?

Here, though B has blatantly refused to answer the question (yet), she has not really changed the subject. When it is understood that the "that" which B asks about is the tape recorder with which I was taping the conversation, then it can be seen as a relevant response - but an <u>indirectly</u> relevant one. It is relevant because, as we can work out, whether the machine is off or on may make a difference to what gets said in answer to the initial question. B's response is <u>directly</u> about the currently recording tape, and is relevant to A's question via the machine's relevance to its answer.

Relevance is rarely a simple relation between consecutive utterances: an utterance can be, and usually is, relevant to more than its immediate predecessor. To return to example 1, analysis revealed that it was much more richly informative than it first seemed to be. By analysing it against a background of relevant information, including the interests of the parties concerned - their "current purposes" - it was found to yield much more information than was given by the words in it. All but line 6 of it was relevant to more than could be learned from what is said. For example, "Wimbledon's got a fair share of celebrities" was relevant to A's background, N's background, A's opinion of N, as well as (less interestingly in this context) to Wimbledon and celebrities. Further information was derived from knowledge of what the utterance was relevant to: in this case, a loosely related collection of subjects.

I have been illustrating the complexity of relevance. Now consider -

Example 3

1. A. There wasn't anything with his secretary?
2. B. No, no. Driving too fast, and getting up to look after the sheep at 5 in the morning.
3. But I think if you can't repair even that basic level of activity ...
4. And F was using it as a weapon.
5. Well, T wouldn't listen; I mean he wouldn't sit down and talk.

The discourse of which this is part had an extra-ml. purpose. the parties to it have met together with the aim of helping B reach a

decision on that question.

As well as being highly complex, the relations of relevance in this snatch of discourse are also structured. The whole discourse, in fact, is structured by the exigent question of whether B should or should not go to Caster. That question has two major sub-topics or -issues: what would it be like if she went?; would it be worth it? - would it do any good? It is relevant to those questions via How T and F's marriage is, and whether T will listen to reason. For the answers to those questions will make a difference to what the visit would be like, and whether it would do any good (could she help repair their marriage?). - Just as the answers to those questions will make a difference to the over-all issue. (For a much more detailed analysis, see Chapter Four , below.)

In this specifically purposeful conversation, we find a hierarchy of issues/topics/questions, and sub-issues/topics/questions, and possible sub-sub issues, etc., with no principled limit on depth apparent. In this context, only remarks which have a bearing all the way up the hierarchy are valued; any others were passed over, ignored, or dismissed. In casual chat much shallower levels of relevance are permitted. But even then, there must be some structure of relevance: whenever two consecutive utterances are relevant to each other it is via some subject-matter of greater generality than that of the individual utterance. At the level of "what is said", utterance i may be about X, while utterance j is about Y; but, if they are relevant to each other, it must be because i and j are both about Z. This is why "Be relevant" must apply at once to some level other than

what is strictly said.

The discussion of how to apply Grice's maxim of relevance has yielded some clues as to what relevance is. I pursue those clues through the next two chapters. Meanwhile, there is one further maxim to consider, namely, Be Perspicuous. Grice's non-talk analogy to this is, "I expect a partner to make it clear what contribution he is making, and to execute his performance with reasonable dispatch". In order to be perspicuous in conversation, one should "avoid ambiguity", "avoid obscurity of expression", "be brief (avoid unnecessary prolixity)", and "be orderly".

Grice uses two well-known examples to illustrate deliberate ambiguity (the flouting of this maxim): "Never seek to tell thy love, Love that never told can be", and "Peccavi" (I have sinned/Sind). I think the point he is making with them is that, if someone is being deliberately ambiguous, he intends the hearer to adopt either both possible interpretations, or the less "straightforward" one alone. Possible ambiguities are rarely noticed in conversation, because only one interpretation is normally relevant; deliberate ambiguities will only take hold where both interpretations are relevant. Only then will the need to decide arise, and that applies both to accidental and deliberate ambiguity.[5]

To illustrate obscurity, Grice projects a conversation between two adults in front of a child. One is conspicuously obscure, hoping the other adult, but not the child, will understand. Talking French, Pig Latin, criminal cant, etc., are established parallels to this strategy, in which outsiders are

prevented from finding out the "contents of the communication".[6] Deliberate obscurity is a peculiar case, and infrequently encountered in conversations between adults; but accidental obscurity is relatively common.

Example 4

1 A. Remembering how much you crammed into four days in Holland ...

2 B. DON'T! That bastard turning up at 4 o'clock in the morning. You can always have one as a baby-sitter and let him sleep on the carpet you know. There's two of them are quite nice ...

3 A. Who?

In this example, B has been so obscure as to violate the maxim irreparably.

Example 5

1 C. It was that quick. It's got to be serious. Still, presumably she's OK if she's home already. I just don't know.

2 D. Presumably. The Royal Free isn't ... well, never mind.

3 C. Well, to be in and out so quickly - she only went in on Saturday.

4 D. Oh, it can't be bad.

Though this example may be obscure to us, to C and D it was all quite as clear as was needed for mutual understanding. In example 4, B's mistake was to assume that her audience had sufficient background knowledge to construe the relevance of her remarks, and therefore to know who and what she was talking about. Whereas in example 5, both parties do have the relevant

common background. Having that enables D to discern that the "it" in C's first sentence refers to N being taken to hospital, while the second "it" refers to her illness; that the "quick" in the first line speaks of a time-span of a few hours, while the "quickly" in line 3 is of a time-span of a few days; and so on. Each knows what the other is talking about (what their utterances are relevant to), and once we know that, the passage ceases to be obscure. Relevance plays a role in our perception of both ambiguity and obscurity.

To illustrate deliberate prolixity, Grice contrived the following ingenious example:

Compare the remarks:

a) Miss X sang "Home sweet home".

b) Miss X produced a series of sounds that corresponded closely with the score of "Home sweet home".

The author of (b), Grice suggests, uses it rather than (a), "to indicate some striking difference between Miss X's performance and those to which the word 'singing' is usually applied." The ingenuity of (b) lies in the fact that, although it is prolix, it does not stray from the point but rather makes a special point by means of its prolixity. Prolixity in actual discourse - deliberate or otherwise - is rarely as clever as this, long-windedness is usually boring. It is boring because it introduces material the audience does not want to hear, has no interest in. If there is a point being made, it could be made without the extra talk - which is <u>extra</u> because, in short, it is irrelevant. That is why it is deemed long-winded, prolix, redundant. One judges prolixity by the standards of relevance.

Grice does not expand on the theme of being orderly, but he presumably means something more than observing the grammatical rules which make individual utterances orderly. It is an orderly succession of utterances that is required, or perhaps merely an orderly succession of turns, in that each lets the other take the floor.[7] In the former case, within a succession of utterances by one speaker, the requirements of ordering I suppose come from the need to put first anything which is essential if the audience is to understand what comes after. If your hearer does not have the relevant background for understanding S, then establish that background at S-1, rather than at S+n. In this light, Be orderly is equivalent to, Avoid obscurity. But this species of good order can only hold through successive utterances of the same speaker. Since one party to a discourse is not normally in a position to know what the other parties may be going to say, successive remarks by different people cannot be orderly in this sense. Good order between successive speakers consists, surely, in their making their remarks relevant to the current topics. And of course, that requirement also holds within passages from one speaker.

* * *

So, taking Grice's maxims of conversation one by one, I have found that in one way or another each is bound up with the notion of relevance. Grasping the information which any given utterance is meant to impart depends on grasping what it is relevant to. And making one's own remarks as informative as required for

current purposes entails making them relevant.

Being informative also entailed being truthful, since misinformation is not information.[8] Relevance, itself, being a relation between an utterance and something else, requires the parties to a discourse to go beyond what is strictly said. And so judging informativeness and truthfulness also must involve recourse to levels other than what is strictly said. Even the last maxim, Be Perspicuous, which, as Grice remarks, differs from the others in applying to how what gets said is said, turned out to depend on judgements of relevance.

If any maxim should be elevated above the others, it should not be Tell the truth, but, Be relevant. It remains to be seen what kind of a relation relevance is, and what it is a relation between. I have barely touched on its structured complexity in this chapter.

Chapter Three: Relevance

We saw in the last chapter that considerations of relevance dominate conversational practice. In this chapter we shall pursue the question of what relevance is. As a starting point, we shall take it for granted that it is a relation, and ask what it relates. So far, we have been discussing the relevance of utterances: questions, statements, exclamations, etc. But, as Grice points out, conversational relevance is a special case of a more general phenomenon.[1] As well as utterances, thoughts, actions, illustrations, and people etc. can all be spoken of as 'relevant' or 'irrelevant'. Ideally, I would want any definition of 'relevance' I was working with to embrace that whole range. However, since it is firstly for the analysis of conversation that I wish to define it, I shall focus first on the verbal cases.

Assuming an utterance is the first term in the relation, X is relevant to Y, what sort of thing is the second term? An obvious candidate would be purpose, which would include explicit question-asking and question-answering, and decision-making. Let us take a particularly purposeful stretch of discourse, such as example 3 in Chapter Two, in which the parties had got together specifically to decide whether or not B should 'go home'. Even if we disregard that issue, all the other issues which are of interest remain (see next chapter for details). But answering the question of, for example, B's mental health, would hardly count as a purpose; it's just that everybody concerned would be

interested in the answer. 'Purpose' is clearly too strong a term, 'issue' or 'question' seem nearer the mark.

When there is something to be decided, or discovered, or resolved, - explicit or not - an utterance is relevant insofar as it abets that decision, discovery, or resolution. This suggestion is very close to that proposed by Keenan and Schieffelin in their paper 'Topic as a discourse notion'. They there treat utterances as responses to "questions of immediate concern", with the "discourse topic" being "the proposition or set of propositions that the question of immediate concern presupposes." (1976: 344)

A characteristic example of what they mean is, Question: What's in here? presupposition, Something's in here. It's hard to see in this case - and all their cases are like this in their essentials - what the presupposition adds. 'Something' simply functions as a blank place into which an object name is to be slotted, and that is precisely the function of the question word 'what'. This restricted notion of 'topic' may be a result of the simplicity of their data, which are almost entirely child-adult dialogues. In practice, 'something's in here' will be just one among many assumptions relevant for interpreting the stretch of discourse in question (see below and next chapter). Are they right, though, in thinking that this background of assumptions is determined by a 'question of immediate concern'? If so, then we have two candidates for the role of <u>object of relevance</u>: questions on the one hand, and sets of assumptions on the other.

The suggestion that understanding relevance in discourse entails understanding to what questions or issues remarks are

addressed, is borne out by the examples I have given so far. If that is correct in general, then any assertoric utterance will be relevant only if it is addressed to answering some explicit or implicit question that concerns the parties to the discourse. And it would follow from this that any conversation in which the speakers produced consecutive relevant remarks, i.e. any coherent and intelligible conversation, would be one in which information got exchanged. It is certainly true that when there is a question of immediate concern, then remarks not addressed to it will be irrelevant. And if one wants to understand relevance from the point of view of reasoning, of inference-drawing, decision-making, problem-solving, then it is the conformity of utterances to this standard of strict relevance that will be of interest.

However, it is certainly not the case that there is invariably some prior question or issue in a discourse. [illegible] may relevantly inform another person of a rainbow, or of a charging bull, or even, in adult-child dialogue, of a [illegible] cow. In none of these cases has there been any question that is being addressed; what makes it relevant information [illegible] a presumed prior interest. The case of the [illegible] cow in particular bears this out. [illegible] another adult about a harmless cow [illegible] had been found, or a green cow, or a cow which was interesting in some other way. But to a small child [illegible] interest, quite as wonderful as a rainbow [illegible]. It seems that 'interest', [illegible] than 'question' to designate what [illegible]

are or are not relevant.

Although not all the cases I have been discussing have involved utterances being addressed to questions, they have all had the characteristic of immediacy. Rainbow, bull, cow, they all to varying degrees demand attention now. But what about carefree idle chat? It is no less coherent and intelligible because it lacks immediacy. We talk about this and that, discuss friends, recall the past, and so on. Consecutive utterances relevant to one interest will be succeeded by, or sometimes interrupted by, utterances relevant to another. One interest after another can become the current interest, so long as there is no more pressing concern. The more exigent an interest, the more likely it is to have remarks addressed to it, and the more likely it is that remarks not addressed to it will be treated as irrelevant. But it is also true that those remarks which were regarded as irrelevant because they did not bear on the matter in hand, will still have been relevant to the interest which motivated them. So immediacy or exigency is not an invariable condition of relevance, though it plays an important role in the workings of the conversational maxim, Be relevant.

I remarked above that, as well as questions, sets of assumptions were candidates for the role of objects of relevance. For any utterance, judging its relevance would involve interpreting it against a background, and that background would then function as the second term in the relation X is relevant to Y. In effect, this is the approach adopted by Sperber and Wilson in their account of relevance.[2] They take for granted a pool of assumptions, many of them available to every speaker in any given

discourse, and treat relevance as a property that propositions have in varying degree when combined with some subset of that background pool. We shall return to their account of relevance below. Meanwhile, the question of what sort of thing an object of relevance is, remains to be answered: should we say it is interest, or should we say it is some limited set of assumptions?

If an object of relevance is a set of assumptions, then the question immediately arises of what picks out that set. An obvious possibility is that it is an interest which picks it out, hence the dual candidature of interest and background. In fact, it is clear that that must be the case. For if we have been pursuing our interest in X by talking about X, and we then stop talking about X and start talking about Y, then what we know or assume about X will stop being relevant, whilst our assumptions about Y will become relevant. At any given moment in a discourse, the relevant background will be determined by the interest which is in play. So, interest must be regarded as the primary object of relevance.

Let us assume for the moment, then, that in the relation X is relevant to Y, X typically obtains between an utterance and an interest; and let us now ask what sort of relation it is. A rough first stab would be that an utterance is relevant to an interest if it contributes to or furthers that interest. In effect, Sperber and Wilson's theory of relevance is a precise spelling out of one version of that suggestion. For they argue that in order to be relevant, an utterance must be informative relative to some background set. If we assume, as I've suggested, that the background set is determined by the interest

in play, then an utterance will further that interest - and thereby be relevant - by contributing new information to its background. They define relevance in terms of the non-trivial implications which can be newly drawn by combining some proposition with the propositions which constitute the background at a given moment. Degree of relevance is then treated as resulting from a trade off between informativeness and processing effort. Relative accessibility will crucially influence processing time, so that its combination with the most accessible background assumptions will always partially determine an utterance's relevance. And which assumptions are most accessible will of course depend on what interest is in play.

Is it the case, then, that in order to be relevant a remark must be informative? Shouldn't we reject this suggestion for just the same reasons that we replaced question with interest as object of relevance? If questions were indeed the invariable interests to which utterances are addressed, then, as I remarked above, all conversation would consist in the exchange of information. On the face of it, counterexamples abound, the great bulk of them being non-assertoric utterances such as questions themselves, exclamations, and commands. In addition, repetitions, summings up, and statements of what everybody knows everybody knows, would all seem to fail the informativeness test, and although any of those may be irrelevant, none of them is necessarily so. However, as Sperber and Wilson point out, as well as supplying information about the world, speakers may also relevantly inform each other about themselves. So, for example, when X asks Y a question, she is informing Y that she wants him

to produce the answer, from which fact certain new conclusions follow. Or when N and O stand before the rainbow going "Oh!" and "Aaah!" and "Look at that!", etc., they are telling each other how much they appreciate the natural wonder.

From Sperber and Wilson's point of view, analysing utterances of all sorts in terms of their informativeness gives them a uniform way of handling a very wide range of material. But from my point of view, it obscures the fact that - for example - N and O's exclamations as they stand before the rainbow are about the rainbow, not about their own states of mind. Unlike, say, "Hello, nice to see you!", these remarks are only secondarily relevant to the mutual relationship of the exclaiming parties; they are primarily relevant to the natural wonder. What the speakers are doing is expressing their common interest in the rainbow. Similarly, when N asks O a question, she is expressing her interest, and thereby letting O know it and with luck arousing in O an interest in answering it. As well as being relevant to an interest by changing it by supplying new information to its background, an utterance can be relevant to an interest by expressing it, or in the case of the hearer, by arousing it.

A relevant utterance, then, must either express [illegible] add information to an interest. But in order to conform to the conversational maxim, Be Relevant, a more demanding [illegible] must be met, namely that it be a <u>common</u> interest which the utterance affects. However it need not have been a common interest <u>previous</u> to this conversation: [illegible] O's interest in X - say, by expressing [illegible]

then X will ipso facto have become a common interest. Or if X has been common interest between them previous to this conversation, and still is, then either party can relevantly make a remark addressed to that interest, whenever there is no more pressing interest to be dealt with.

Having sketched an account of relevance for utterances I'd like to see how it generalises to non-conversational contexts. As Grice points out, the rules of conversation are a special case of the rules that govern all cooperative practice. The obligation to contribute to the common interest obviously applies in some shape or form to all exchanges with other people. How well do the forms of relevance we have found in conversation apply to exchanges of other sorts? For a start, if the same interest is not aroused in every party to a cooperative exchange, then the cooperation will fail, either because someone is not participating or because someone is pursuing a different interest. Arousal of a common interest is a prerequisite of the pursuit of a common interest, accordingly, anything which aroused it would presumably be relevant to it.

We each have a repertoire of gestures like smiling, weeping, shrugging and pointing which may be used to express a common interest. When they are, then they are meant to be understood, and understanding them will entail grasping what they're about, i.e. what interest it is that they express. As well as these directly communicative expressions, any action which furthers an interest could be viewed as an expression of that interest if it emanates from it and is a result of its arousal. Be that as it may, those actions which further an interest are clearly relevant

to that interest in a way which parallels the third, strictest, form of relevance; which consisted, in a conversational context, of the retailing of information.

Shall we say, then, that in order to be relevant something must arouse, express, or further a common interest? Unfortunately, no. Consider Grice's example of two people tinkering with a car engine, and imagine that a spanner of a particular size is missing, so that the much less convenient monkey wrench has to be used instead. The absence of the spanner in no way furthers their common interest, but it is certainly relevant to it. It is relevant because it makes a difference to the way in which the interest is pursued. Equally, negative information which prevents an interest being pursued in a particular direction, will be just as relevant to it as information which furthers its pursuit. In fact, Sperber and Wilson's account of relevance in a linguistic context allows in all relevant information: what we need is a generalized definition which parallels this by capturing all effects, positive and negative, on the form of an interest or the manner of its pursuit, or its outcome.

At any given moment in the pursuit of an interest, it will be possible to speak of the outcome of the last action, of the next action, etc., right up until the final outcome, if any there be. At each stage, its form will be partly determined by the form which preceded it, and will partly determine the form which succeeds it (for elaboration of this point, see Minsky, 1975 and below). Let us say, then, that something may be relevant to an interest if it makes a difference to its outcome, where its

outcome may be its own changed form. That seems to cover the whole range of the strictly relevant, including the acquisition of relevant information.

So, in order to comply with the maxim, Be Relevant, in a cooperative exchange, an action, utterance or thing must arouse, express, or make a difference to the outcome of a common interest. 'Interest' must be taken to cover a range from idle wonder to gritty purpose, taking in issues, questions, and desires; and a possible 'outcome' of an interest may be its own changed form. However, this is still not a completely general account of relevance, for it is not a relation which is confined to cooperative exchanges or particular occasions. Sequences of actions or utterances produced by an individual may be just as relevant to one another as they would be if produced by different people. The maxim of relevance, as spelt out above, is clearly a special case of the very general social (and even animal) rule: Contribute to the common interest. A common interest is only that because individuals subscribe to it; interest is a primarily individual and psychological phenomenon. So something may be relevant to an individual because that person has some particular interest, whether or not it is held in common with other people.

Let us define relevance, then, thus: to be relevant is to arouse, express, or make a difference to the outcome of an interest. That will cover both public and private issues, though in practice social relevance entails the interest in question being a common interest. And if we allow that anything which has made a difference to some interest, will make that difference, then this definition also takes in rephrasings, summings up, and

any material from the background. This is surely correct, since information I have acquired before this moment informs my thoughts and actions now, and hence is relevant to them now. The background to an interest is relevant to that interest because it is made up of what has affected that interest in the past, and thus makes it what it is now.

* * *

Now let us take these three ways of being relevant - i.e. arousing, expressing, or changing an interest - and give a provisional sketch of what they amount to in practice. So, what is it for an interest to be aroused? What distinguishes an interest which is aroused from one that is not? The first thing to say is that it's not as straightforward as that - arousal is a matter of degree. As I show in subsequent chapters, it is the normal human condition to have several interests concurrently active, and by presumption therefore to some degree aroused. The single-minded concentration we can bring to a demanding task or engrossing situation is the exception rather than the rule. When, as is normal, many interests are active at the same time, not all of them are active to the same degree: whilst some are leading to either actions or utterances or to cognitive activity, others are merely alert to the relevant.

In a conversation in which all parties are addressing a common interest, then the same interest will be aroused in each of them: utterances relevant to it will be provoked, and considerations relevant to it will spring to mind. When a common

interest doesn't impose too heavy a demand, then an utterance produced by A as relevant to it, may arouse a different interest in B. In which case, B may well use it as a link and aim to arouse this different interest in A, thus creating a new common interest, and changing the subject. Indeed B may have been waiting for the chance throughout the conversation; a highly aroused interest will seize on the relevant. So, arousal has to do with both attention, or receptivity, and production, or creativity.

In what way does arousal affect the background of assumptions? It makes the relevant background more accessible for the interpretation of input: for guiding one's expectations, and spotting what's out of place; for supporting or arguing against a case, and coming to conclusions. Arousal also makes those assumptions available as a basis for informed action, including speech; they guide the production of output as well as the interpretation of input. Even an interest so faint as to be almost extinct may be rearoused by some object or some phrase, bringing with it at least some of the assumptions which once informed it.

So, now let us move on to <u>expression</u>; what is it for an interest to be 'expressed'? The cases of actions or utterances one is inclined to label 'expressive', are only that, hence we use that word to pick them out. That is to say, an expressive action or utterance may arouse a like interest in a witness, but will not make any difference to the form of that interest. But it does not follow that an action or utterance which does affect the outcome of an interest, and results from the arousal of that

interest, cannot by that token be an expression of that interest. In other words, it does not follow that something which is relevant in the strictest sense, i.e. by changing the form of an interest, cannot also be relevant in a weaker sense. Doesn't the baby express its appetite at least as clearly by failing on the tit as by bawling for it? And don't I express mine just as clearly by getting up, getting myself something to eat, and eating it, or by announcing that the smell of cooking is making my mouth water, as I do by saying "I'm hungry"? So in general I think we should say that anything which springs from an interest as a result of its arousal will be an expression of that interest.

Now let us take the strictest sort of relevance, and see what it is for the form of an interest to be changed. Firstly, I should explain that the notion of the form of an interest has been introduced in order to distinguish this species of relevance from that of arousal. For arousal clearly also brings about a change in an interest, but it is a change in its state rather than its form. Making a difference to the outcome of an interest is equivalent to changing its form, so if something affects its outcome, eo ipso its form is changed. But what does this amount to? Let us take a very simple case. I want to drink some of the coffee that is in a cup beside me; I reach my hand towards the cup, and thereby take a step towards the final outcome, i.e. towards satisfying my want. Now I have moved my hand towards the cup, there is no longer that step to take; it is no longer the first priority. The priority now is to take hold of the cup, next will be picking it up, etc. At each stage, the immediate

future of the interest depends on its immediate past, and so on for the whole sequence rather than a random assemblage.[3] In what way is the form affected when the final outcome is reached, and I have drunk some coffee? I suggest that it is affected by closure: no more energy will be expended in this direction, through this interest (unless and until it is rearoused and I reach out for another gulp).

This may be all very well for a simple physical task like having a cup of coffee, but how does it apply to a complex coordinated activity like conversation? For a start, I want to suggest that it's reasonable to see the form of an interest as the set of assumptions which make it what it is at any given moment. The form of an interest, then, will be determined by its relevant background; therefore, changing its form will entail changing its background. An utterance which is relevant in the strict sense, then, will be one that has consequences within that background. And, one may presume, the more consequences it has, the more highly relevant it will be. A precise account of this consequentiality, and of the notion of degrees of (strict) relevance is to be found in Sperber and Wilson.[4]

It is no concern of mine in this thesis to give any detailed account of strict relevance in a linguistic context. Rather, I am concerned to examine relevance at its most general, and see what follows for one's picture of the mind, and of the role of language. If the account of relevance that I have just sketched is right, then doing so is going to involve having interest as an absolutely central concept. So I shall bring this chapter towards a close by surveying what we can say so far about what

interests are. As it is used in this thesis, 'interest' must be taken to embrace a range rather wider than it does in normal use. Needing a word that would take in (at least) plan, purpose, issue, question, hope, fear, desire, wonder, love, and concern, as well as 'mere interest', I chose 'interest' as at least entailed by that whole range; so, what do the members of that range have in common? They have in common at least that they may be objects of relevance, and as such are capable of arousal, expression, and change.

Interests are typically brought forward as explanations of human activity of all sorts, and treated as motive forces or springs of power: N did/said/thought such and such, because N hoped, desired, wondered, and so on. This aspect is brought out in the notion of the expression of an interest. Any activity which in some sense 'springs from' an interest, and can therefore be attributed to that interest, will count as expressing it. In fact, people typically invoke N's beliefs as well as N's interests when explaining N's actions, utterances, or thoughts. On the view I am proposing, interests and beliefs are not separate entities which have to be combined afresh for each occasion, rather, they are invariably complementary. For the form of an interest is determined by the information that has shaped it. This aspect of interest is brought out in the notion that an interest is susceptible to change.

As a receptor, a locus of change, an interest is a repository of information which becomes accessible if and when that interest is rearoused. Memories are made of this. All interests have duration; though some pass in a flash, leaving no

detectable trace, others last a lifetime and affect one's conduct from birth to death.

Most of our interests persist through interruptions and diversions, becoming rearoused from time to time through internal pressure or external stimulus, or both. It is the same interest through <u>time t</u> and <u>time t+1</u> just to the extent that the same differences have shaped it, that it is informed by the same assumptions, at <u>time t</u> and <u>time t+1</u>. This allows for relations of complete or partial identity, and of inclusion and exclusion, to hold between interests. <u>Mutatis mutandis</u>, the same relations may obtain between the interests of different individuals as well as of the same individual across time.

As well as being a locus of activity and of information, and having form and duration, an interest is something that we <u>feel</u> in varying degree. Interest can vary from weak, feeble, slight, at one extreme - 'a merest flicker' - to intense, fervent, fervent, ardent, at the other - 'a passionate concern', 'a burning passion'. This is presumably an aspect of the phenomenon of <u>arousal</u>, which as I have pointed out varies in degree both over time and between different interests at any given time. The more highly aroused an interest is, the more it belongs at the fervent extreme, the more likely it is to lead to expression of one sort or another.

* * *

Much of the rest of this thesis is devoted to examining and testing the propositions set out in this chapter, in particular,

the proposition that to be relevant is to arouse, express, or make a difference to the outcome of an interest.

Chapter Four: Issues and Questions

"Contribute to the common interest or concern", is clearly a precept with very wide application. I will show some of the ways in which it applies in conversation, below. The precept automatically allows for non-informative discourse: parties may tell jokes, tall stories, whisper sweet nothings, share recollections, etc., and still conform to that rule.[1] Here, I am particularly concerned with those situations in which parties contribute to the common interest by aiming to increase relevant mutual knowledge. In my terms, an utterance is relevant if it arouses, expresses, or makes a difference to the outcome of a common interest or concern. Among these three, the third, strictest, species of relevance will be the focus of this chapter. Any change in relevant mutual knowledge is ipso facto a difference in the outcome of a common interest.

Contributions to the common interest are necessarily relative to a community; before any conversational rules can apply, there must be at least two people prepared to converse with each other. As Schegloff puts it, "a person who seeks to engage in an activity that requires the collaborative work of two parties must first establish, via some interactional procedure, that another party is available to collaborate" (1968: 1089). People thus embarking on a discourse enter into an active community of interest; in doing so they undertake to contribute to the common interest. Such 'communities' may vary from a pair of transient strangers, at one extreme, to members of the same

family, at the other. Two strangers are normally bound only by a transient interest; once the interest ceases, their community also ceases. These short-lived exchanges most often consist of requests for information, "where the obligation to respond is constant and general" (Labov and Fanshel 1977: 89). Requests to be told the time or the way are of this sort; Labov and Fanshel (1977) quote Erving Goffman's term "free goods". Everyone has the obligation to pass on available free goods to anyone who has expressed a want of them, once willingness to cooperate has been signalled. One can dodge the obligation only by dodging the inquiring stranger's gaze. It is, however, a strictly limited range of interests which strangers are thus obliged to treat as common. Faced with questions on any number of other interests, a stranger will be entitled to deny community of interest: "It's none of your business", "That's my concern", will be legitimate retorts.

One has no right to ask a stranger questions one has every right to ask, say, a spouse. (Absolute intimacy would consist in wholesale participation in each other's interests, and entail the right to ask anything and the obligation to tell all.) Most communities of interest in which we are active fall somewhere between the two extremes. Each of us belongs to a multitude of distinct, overlapping or concentric communities of interest; some last a lifetime, some an instant. It is not my business here to catalogue their variety, nor to discuss how members can manipulate membership allegiances (see, e.g., Garfinkel 1967; Goffman 1967, 1971). What is relevant is that a community of interest lasts just as long as there is at least one common

interest to bind it. Fellow members are obliged to contribute information relevant to any 'binding' interest, on request. If B has information which A has expressed a want of, B can only deny it to A on the grounds that it is information privy to some community of interest in which A is not included.

The obligation to pass on relevant information may obtain even when no question has been expressed. For example, except in war, perceived danger imposes the duty to issue a warning to any threatened other (cf. "free goods"). In closer, long-term, communities of interest, non-physical threats to one's fellows' interests also impose that obligation. If A and B are friends, partners, kin, etc., and B knows of a threat to A's interests, failure to inform A of it may justly be construed as a betrayal of their community of interest. Failure to warn may be inexcusable; omitting to pass on other relevant information may also merit reproach. Faced with such a reproach, one may either offer an apology or, once again, deny community of interest.

In all these cases, a perceived inequity in mutual relevant knowledge creates an obligation to repair it, within a given community of interest. And, in all cases, collaboration between members of a community of interest depends on their having a background of relevant mutually reliable knowledge. What picks out the background in any given discourse is its relevance to a current common concern. In my terms, the relevant background is what has (at any previous moment) made a difference to whatever the current common concern or concerns may be. Contributing to the common interest in the ways I have been describing involves increasing the stock of mutually reliable knowledge.

Supplying information is not the only way of contributing to the common interest in such a way as to increase the mutually reliable relevant background. Obviously, the mutually reliable will always include the certain, what is known for sure. In practice, it will also include a good deal that, though taken for granted in this community, might be treated as arguable by outsiders. Something is arguable (an issue) just as long as there are alternative outcomes; and as long as there are possible alternatives, no one possibility can be relied upon. Hence it is always in the common interest to exclude alternatives, resolve issues, when they arise. Detailed analysis reveals that a surprising amount of discourse is devoted to this end.[2]

Assertions modified by "I think", "I would say", "it seems to me", etc., are expressly put forward as opinion, suggestion, rather than as 'hard fact'. As such, the proposition within the intensional scope is explicitly open to argument. But, of course, speakers are not obliged to signal the disputability of what they say, and may not even acknowledge it. For the analyst, the best evidence of disputability is subsequent argument, either pro or con a given assertion. If an assertion is evidently true to all concerned, then it requires no supplementary support. If it receives such support, that implies that the truth of the assertion was not self-evident; just as the provision of counter-examples implies that.[3]

Some issues, such as whether or not a door is locked, can be unambiguously resolved at once. But some can never be resolved absolutely, placed beyond doubt for all time: all 'value judgements' are of this sort. And some issues, such as any

suggestions about the future, can only be resolved after some time has passed, not in the here-and-now of a current discourse. The point of argument is (or ought to be) to bring round, make of one mind, agree. Then what is agreed, the common judgement, can be treated as given in any further exchanges between members of that community of interest, even when it cannot be counted on in other communities.

An issue may be voiced as a disputable assertion or question, or it may never be explicitly expressed at all. Evidence for the existence of such an unspoken issue will be that remarks are addressed to resolving it. For example, take the following imagined case. A husband comes home, immediately brushes crumbs off the table, empties all the ash-trays, then goes to the fridge and says, "Oh, you didn't get me any beer". The wife responds with, "I'm awfully sorry, I just couldn't manage the extra load. I did wash all your undershirts, though". The relevance of her response and its internal coherence depend on her having heard his remark as a criticism, not only of her failure to get beer, but of her adequacy as a wife. Her remarks are addressed to that issue, as she has taken his remark (and other actions) to be. Both parties are putting into public space information relevant to the outcome of that issue, to what is (in this sexist little community of interest) the proper judgement of her adequacy as wife: his contribution argues for one conclusion, hers for another.

Issues like that of the wife's adequacy are, I think, clear cases of what Labov and Fanshel call "propositions". These are often, but not always, explicitly stated, they are "persistent

potential reference points in the interaction between speakers" (1977: 122fn). Remarks will tend to be taken as supporting or not supporting "general propositions known to both (speakers) though not necesarily believed by both of them" (1977: 122). This is one of Labov and Fanshel's many insights into the nature of actual discourse that have assisted my own conversational analyses, below.

Example 1(a)[4]

1. A: Hello
2. B: Hello
   That'll be £9.31 pence
4. A: Sorry, I've no small change
5. B: Not to worry, no problem
6. A: I saw the article in the News & Echo
7. (B: silent acknowledgement)

A is a regular customer at a supermarket in which B works on a till, at which the conversation takes place. Beyond their occasional encounters in this setting, they do not know each other at all. In the normal course of events, the converation would have closed with an exchange of farewells immediately after 5, during which their transaction is coming to a satisfactory conclusion. But this time, A introduces a 'new' concern - (6) "the article in the News & Echo" - and secures uptake from her interlocutor. She thus changes the basis of their community of interest: the dreary, predictable task of paying for the shopping is succeeded by a concern of quite a different order. Line 6 (on

its surface, anyway) is a statement about an event in A's life. As such, according to Labov and Fanshel, it "requires only an acknowledgement of a minimal kind" (1977: 101). The sequence from 6-7 bears out this informal rule. Here is what follows it:

Example 1(b)

8. A: cried all the way through
9. B: Yesterday a woman was angry with me.
10. A: ANGRY with you?
11. B: Because it made her cry
12. A: Angry with YOU? She seems to have it wrong; it's not YOU she should be angry with ...

A has learned from the article in the paper that B came from a large Hungarian Jewish family, of which she was the only survivor; that she had been abused and humiliated not only by the invading Nazis but by eager Hungarian anti-semites; that she had spent years in a concentration camp, and had stayed there for 2-1/2 years after the war, because she had nowhere else to go. So, in telling B that she has read the article she is telling her that she knows those tragic facts. By telling her that, by introducing a concern to which those facts are background, A obliges herself to say more. Some comment is called for, and B awaits it.

The comment A produces, "cried all the way through" (8) is meant to convey to B her opinion of the content of the article. If one person weeps for another, we generally take the tears to express sympathy, shared concern, for the other's distress: A is

meaning to communicate sympathy to B by telling her she cried. But, as the ensuing dialogue reveals, the inference from such tears to sympathy does not hold universally: it cannot be relied on in every context, and in particular cannot be relied on in this one. By disclosing that "a woman was angry ... because it made her cry" (9 and 11), B makes the proper attitude towards the article (and her life) an issue. Sympathy and anger should be mutually exclusive attitudes to the same object. So, 9 and 11 provide a counterexample to the generalization on which A has relied in 8.

On its own, line 9 is puzzling: How could anger be a possible attitude toward B on the basis of the article? A stresses "angry" in line 10, siting it in a context of contrasting possibilities, as the most economical way of putting the question I have just given in full. After pausing to let A express her interest, B goes on as though she knew the question was in the air already: 9 and 11 make a perfectly coherent whole. With 11, she 'explains' 9, answers 10, and puts A's attitude in question. On the face of it, B's 9 and 11 are just information about an event in her life. As such, they should invite only a minimal response; but, as we have seen, they also raise an issue.

A's "Angry with YOU?" (12) shifts the focus of the syntactically identical question at 10. This time the appropriateness of anger is not in question, rather, the object at which it should be directed is at issue. Anger is taken for granted (given, presupposed); the question now is not, Why anger, but, Why B? with the implication that there are relevant others at whom the anger might rightly be directed. The underscores

that implication with her next remark - "she seems to have it wrong, it's not YOU she should be angry with" - at the same time producing an evaluation of the woman's anger as wrong. A is taking pains to make it quite clear that, if there is a question as to the right attitude towards that article, she is on the same side as B. She is firmly nipping the possible alternative in the bud; in terminating that possibility, A also terminates the issue of her opinion, first raised by line 6. Ceasing to be an issue, its outcome passes into the background of relevant knowledge that A and B can mutually rely upon. B can count on A's sympathy, and A can count on her doing so.

Example 1(c)

13. B: And that was just the beginning ...
14. (A: silent query)
15. B: There was more, far more. Worse; I couldn't tell it, I can't even think of it
16. A: I'm sure, I'm sure. Some things one can't, they're too terrible
17. B: (sighs) Well, thank you for your kind words
18. (A: silent acknowledgement)
19. B: Well, bye-bye.
20. A: Bye-bye.

With line 13, B is not indicating to A that there is more to be learned about the incident of the angry woman. Rather, she means to tell A that there are more tragic facts to know than those revealed in the article. Though she has aroused A's

interest, put her into a state of felt ignorance, she has no intention of supplying the want. A is to learn only that there was "worse". Although she is excluding A from any community of interest which might be privy to those terrible facts, she at the same time excuses herself for doing so. She would draw nobody into it with her, she herself is its sole, reluctant, member. Many authors have remarked that a negative seems to imply that someone might have expected the opposite (see, e.g., Labov and Fanshel 1977: 104;; Bosley 1975: 6; and Givon 1978). In this case, B's "I couldn't tell it, I can't even think of it" might be expanded as follows. "Given your sympathy, given our community of interest (which I have just acknowledged) you might reasonably expect to be told these relevant facts by me, but I can tell them to no one, even myself".

Like (9) and (11), lines 13 and 15, though ostensibly about speaker events, raise an issue, and A is invited to take sides. Once again she sides with B. (16) "I'm sure" makes it certain between them that B is absolved of the obligation to tell, "some things one can't" universalizes the absolution, and "they're too terrible" specifies its basis. With (17), B rewards A's sympathy with thanks for her "kind words", and initiates a "pre-closing sequence" (see Schegloff and Sacks 1973; and Sacks et al 1974). So, with no issues left unresolved, their discourse comes to a satisfactory conclusion.

* * *

Most accounts of conversation, like this one, bring in some

such notion as mutual knowledge. "Mutual knowledge" is a debatable designation, for some of the argument see Sperber and Wilson (1980), and Prince (1981). For more general discussion of the topic - under various designations - see Schiffer (1972), Lewis (1969), Clark and Marshall (1980), Tyler (1978), and especially Rommetveit (1974). Whatever one chooses to call it, it is the essential background to all communication. In any given discourse, the question arises of how parties 'have access' to the relevant background, of how they pick it out from among their sum total of mutual knowledge. In practice, talkers rarely have difficulty in taking into account what others do or do not know. How they succeed in this is a problem for the analyst, not for the participants. I shall examine that question in the light of example 1.

The relevant background knowledge to lines 1-5 is given by the nature of the transaction which is taking place. It can be nicely handled by a 'frame' type analysis (see, e.g., Minsky 1975). The analyst's problems begin with line 6, "I saw the article in the News & Echo". There is no single News & Echo, it is a weekly paper, and more than a week has elapsed since the relevant issue. And, of course, in any given issue there are many articles. Since they had had no previous discourse about anything except shopping, there is no prior referent to account for either the in (6). Yet A's reference to the article in the News & Echo secures uptake. It is justified, and successful, simply because its referent concerns B.

In disclosing that she has "seen" the article, A is letting B know that she has read it, that she knows what's in it. But

can they now speak on the basis that all the information in the article is now mutual knowledge between them? Admitting at last that I was 'A', if B had relied on my having effective access to all the facts in that article, she would have been wrong. We assume selective recollection in each other, and we are right. So B cannot rely on A being privy to all that information, only some of it: what can she rely on as mutual knowledge on the basis of (6)? Certain facts in the article will have struck anyone as more salient than others, namely those I retailed above. They are such facts as to arouse the concern of any fellow human being. Had there been any doubt in B's mind about which facts were responsible for A's interest, they would have been resolved by (6) "cried all the way through". The concern which A is putting into public space is at least partly made up of distress. It is the distressing facts and events of B's life that provide the background of mutual knowledge.

As well as unambiguously circumscribing the relevant background, A's line 8 is also meant to let B know of her sympathy. But instead of immediately being taken for granted by B, until line 12 it is an issue. Both parties know this, and each knows the other knows; so, for the time being, their mutual relevant knowledge is of this issue against the background of tragic facts already given. It is also true that a fact about some woman was added to their store of mutual knowledge with (9 and 11'; but the woman herself is of no interest, she is a mere counter example. B's next line confirms this interpretation, "And that was just the beginning" ('J. primary but not emphatic stress on that) does not continue from the incident of the angry

woman. It connects instead with the background in which A's sympathy - now presupposed - is grounded; the knowledge imparted in (9 and 11) ceases to be relevant. Although neither party has voiced or directly referred to the distressing facts which constitute the background, B's "And that..." is intelligible. That is so because both parties know what they're talking about, it is neither the woman nor the article that is "just the beginning". Rather, it is what A has learned selectively from the article, namely the horrors revealed therein.

No sooner has she told A that there is more to be known, than B is telling her that she will never know it. The rule, that if you put someone in a position of felt ignorance you will be expected to fulfil the want if possible, is a special case of the obligation to contribute to the common concern, discussed above. That the rule is part of their mutual knowledge (even children know it, breaking it is a favourite childhood tease) is manifest in B's excusing herself for not fulfilling it, and A's accepting the excuse. Although she leaves A none the wiser as to further particulars of her life, A knows that B judges them to be "worse", untellable. In her next turn she lets B know that she believes the judgement without further substantiation, and accepts the excuse. And B's pre-closing, "Well, thank you ..." lets A know that her by now indisputably sympathetic concern is appreciated. As well as events, the objects of mutual knowledge include conversational rules, interests and concerns, issues and attitudes.

Throughout the conversation, 'access' to the relevant background has depended on knowing the current mutual interest.

In the first part that knowledge was given in the situation; the transaction could have taken place in silence. But, from line 6, understanding the common concern crucially depends on understanding a common tongue. Parties can take mutual inteligibility for granted only insofar as they can rely on mutual knowledge of grammar and vocabulary. I wish to argue, however, that mutual knowledge of vocabulary itself depends on a background of common interests. I have already shown that understanding phrases like "the article" (6) and "and that" (13) requires a grasp of the interest that motivates them, in context. It is not such context-determining interests that I have in mind here. Rather, it is those interests (with their attendant knowledge) that we can assume any fellow member of our language community will have been engaged in at one time or another.

To support the claim that common vocabulary depends on common interests, I now briefly discuss lexical comprehension by children of the conversation in example 1. I examined three children, aged 6;5, 4;6, 2;8 on the intelligibility of the vocabulary items in it. The eldest child, who is numerate, literate, and gets pocket money, had difficulty only with the reference of News & Echo, whilst taking it for granted that it was a newspaper. The middle child, though highly articulate, with a vocabulary fit to express his every interest, is neither literate nor fully numerate, nor practised at spending money. His difficulties arose with lines 3, 4, and 6: £9.31 pence, small change, and the article in the News & Echo.

Although he knows what money is, and knows what it is used for (coins and purchases have both engaged his interest), the

difference between one sum and another has never concerned him. Until knowing about money affects his interests (becomes relevant), he will remain ignorant of just what £9.31 might amount to. Small change presents a parallel, though not identical, case. Relative size being a recurrent concern in his life, small is a problem only as a modifier of change - wherein the real difficulty lies. He has change in his vocabulary in the sense of make or become different, as in "We'd better change your trousers", or as a noun in "That makes a nice change". But the question of what small change means in its monetary sense had not so far arisen in his life. Finally, article is a word which is entirely absent from his vocabulary. The explanation for this absence is, once again, the absence of any interest or concern with the thing in question.

The youngest child had all those problems, plus yesterday in line 9. And he had not pursued interests which would yield the knowledge relevant to understanding even that "£9.31 pence" stands for a sum of money. As for yesterday, if he understands it at all, it is as a way of making clear that an event is past and over, rather than as a way of locating an event in past time. His concerns are predominantly with the now and its immediate past and future; what difference does it make to him whether it was yesterday, a week ago, or six months ago that such and such a memorable event occurred?

Someone might want to argue that the differences in vocabulary between the three children are a result of exposure, or experience. It is easy to argue the case against exposure as the determining factor. Take small change again. The middle

child has been 'exposed' to both notes and coins, and is aware that both are money; and he has been 'exposed' to many discussions which have involved talk of "small change" in this sense. That is, he has been present at many exchanges like (and including) that in example 1, 1-5; and he has frequently been there when discussions of "whether there is enough small change for his big brother's lunch or pocket money", have taken place.

The case against experience is nicely put by the philosopher David Hamlyn: "It is certainly not enough ... if a person is to acquire a given concept, that he should be presented with a wide range of relevant experiences unless he is in a position to see them as relevant" (1978: 121). The presence of the four-year-old when small change is talked about has so far been as observer, not participant. Whether or not there is "small change" has never been his concern, even though whether or not there are available small coins (for counting, for example) has sometimes been of interest. Having never participated in a community of interest in which small change is being used as an expression for low value coins, he had never been concerned to understand or use the phrase. Once he is taking part in transactions with money he will be functioning as a member of such a community. Alternatively, he could himself create such a community by asking what the expression means. In fact, my question brought the phrase to his notice, aroused his interest, and he did ask me what it meant. His prior interest in money as a physical object, and in observing shopping transactions, gave him sufficient relevant knowledge to understand the answer.

For adult members of a language community, the meaning of

its basic vocabulary is not in question, is taken for granted. It is mutual knowledge we rely on in conversations with any fellow adult speaker of British-English. Some long-term communities of interest develop specialized vocabularies, such as slangs and jargons, available only to their members, which reflect their special interests and knowledge. Most proper names also are available only within communities of interest narrower than the language community. On hearing "News & Echo" (especially in the context of line 6), most fellow members of our language community will take it to be a newspaper, but only members of a much more restricted community of interest will know which newspaper. Just as, if I say, "Tony is a stage-hand these days", only members of a highly restricted community of interest will know who (which Tony) I mean, though any English speaker will take it that he's a male person. At the most intimate extreme (e.g. loving couples), there is the most idiosyncratic range of mutual knowledge, including always general knowledge available in virtue of membership in broader communities. The richer the background of mutual interests and knowledge, the richer the "meaning potential" of a vocabulary.[5]

Identifying membership of a community of interest is a prerequisite of mutual understanding - a prevailing concern in all conversations, and it is an efficient way of identifying realms of mutual knowledge.

The discussion so far has ignored two complicating factors. One is the possibility of a remark being addressed to several interests at once. The other is the fact that information to which all parties are privy may have been recognized as relevant

only by one party. Both these factors play a role in the following snatch of discourse.

Example 2[6]

1. A: There wasn't anything with his secretary?
2. B: No, no.
3. - Driving too fast
4. - and getting up at 5.30 in the morning to look after the sheep.

Unlike example 1, this is not a complete discourse. It is a fragment of a much longer conversation, in which three parties, A, B and C (who know each other intimately) are taking part. I'm sorry to frustrate the reader by not publishing the whole discourse: it was so private and scurrilous a conversation that the parties to it are understandably reluctant to have it all in print (partly because they agree with my analysis). C, silent here, is B's brother, T is their cousin, and A is married to C. They have met with the express purpose of helping B decide whether or not to make a trip to her and C's home town, where T still lives. This issue dominates the whole conversation, and example 2 is addressed to it too, via a number of other issues. B's response to A's question is a skillful attempt to lay all those issues to rest simultaneously, in such a way as to make the trip look pointless. She is busy looking for reasons why she doesn't have to go.

As guidance in the analysis that follows, I give here a sketch of the issues involved, illustrated by certain assertions

that B has made earlier in the conversation.

Issue i: Should B make the trip?

"I don't particularly want to spend any time to go home"

"Sounds like a fun weekend if I go home ... hmnn, shit"

Issue ii: Could B, by going, forestall the threatened collapse of T's marriage?

If so, all parties agree, it would be worth her going ($\longrightarrow$ i)

"I don't think it'll make any difference if I go"

Issue iii: Is T's marriage already beyond repair? If so, then repairing it (ii) can be no reason for making the trip ($\longrightarrow$ i).

Issue iv: Is it true that T and his wife have had no sex for years? If so, then the marriage is likely to be beyond repair ($\longrightarrow$ iii $\longrightarrow$ ii $\longrightarrow$ i), and either (a) T is likely to be demented from sexual frustration or (b) has sought a sexual outlet elsewhere.

"Why T was going out of his head is because as far as I can see they hadn't slept together for over two years";

Issue v: Is it true that T is demented, irrational? If so, he is not susceptible to argument, therefore there is no point in B trying to argue with him, therefore she is unlikely to be able to help repair the marriage ($\longrightarrow$ ii $\longrightarrow$ i).

"T was going out of his head"

"I can't get through to T, it's impossible"

"You can't talk to him"

Issue vi: Is it true that T has manifest irrational behaviour? If so, it would provide evidence that he is suffering from sexual frustration (iva), in which case it is likely that the marriage has collapsed beyond repair (⟶ iva ⟶ iii ⟶ ii ⟶ i)

A's question, "There wasn't anything with his secretary?" (1) concerns the possibility that T has an alternative sexual outlet (ivb). When B said, "Why T was going out of his head is because as far as I can see they hadn't slept together for over two years", her uses of the past subtly suggested that T was already "out of his head" and that no possibility remained of his having sex with his wife. Although she is attempting to treat T's incipient nuttiness as a given, she knows that it is in fact debatable, and that both A and C doubt it. If it is agreed that sexual frustration is likely to lead to dementia (no one in this discourse chooses to dispute it), then if B could show that T had had no sex life for years, it would lend support to her claim. Whereas, if the contrary were the case, and T had in fact been enjoying some sexual activity, then her claim would be undermined. Hence, A's question is relevant to T's mental health as well as to his sex life (v ⟶ ii ⟶ i).

Even if T has had no affair with his secretary, the possibility is still open that he has in fact had sex with his wife (no one in this discourse is in a position to have a confident opinion on this). B follows her denial that T has had an affair with his secretary (2) with, "Driving too fast, and getting up at 5.30 to look after the sheep". In doing so, she is, in part, addressing the possibility that all is well, after

all, with T's marital sex life. Reckless driving is a notorious form of sexual sublimation; it is also irrational. Given that T is really a lawyer and the sheep just lose money, his pre-dawn sheep-tending - though implausible as sublimation - is certainly irrational. These are two instances of T's unreasonableness, and as such lend support to all B's contentions (vi ⟶ iva ⟶ iii i and v ⟶ ii ⟶ i). This is the background which confers coherence on lines 2-4.

When B puts those cases of T's irrational behaviour into public space, she is not telling A and C anything new. Knowing T, they both know his habits: the sheep and the driving are long-standing concerns. What B is doing is making A and C aware that those facts should make a difference to the outcome of the issues current here. Several researchers have argued that conversationalists display a preference for conveying new information.[7] But, as far as I know, only Sperber and Wilson (MSb: Chs. 1-2) have attempted a fully explicit account of newness, in particular of relevant new information. In their terms, an utterance will provide new relevant information if and only if it combines with a 'background set of assumptions' to yield implications given in neither the utterance nor the background. By being relative to a restricted background, their definition accommodates cases like this, in which information is seen as newly relevant. In my terms, their relevant utterances are those which make a difference to the outcome of a common interest. In order for some fact or event to make a difference to a current interest, it cannot previously have been taken as relevant to <u>that</u> interest, or there would be no difference still

to make.

Issues and questions, because they require to be resolved or answered, impose the particular obligation to make a new difference to the outcome. Addressing a debatable issue involves putting forward facts or events which will support or count against one outcome or another. An issue is only resolved when every alternative but one has been ruled out, but it may be settled by agreeing pro tem on a likeliest outcome. Once enough has been said to settle an issue, there is no more to say. Similarly, once a question has been answered, it ceases to be a question, hence ceases to be a current interest. Grice's "Be no more informative than required" (1975), Sacks and Schegloff's "preference for minimization" (1979; see also Schegloff 1972), and Sperber and Wilson's equivalent specification (1980) all follow from the nature of issues and questions.

A problem for any analysis of actual talk is, as Sperber and Wilson put it, "how does the hearer know which background assumptions, out of the enormous range which could in principle be part of the intended context, are the ones the speaker intended him to use?" (1980: 5). This is the same problem as that of access to relevant mutual knowledge, discussed above. Identifying the relevant background is a prerequisite of comprehension, both for participants and observers. I have been arguing that doing so crucially depends on identifying the interest or interests which are being addressed. A fact or event may be relevantly stated as long as it has made or would make a difference to the outcome of a current common interest. If the only interest to which it was relevant is now dead, then it

cannot be relevant to state that fact or event now. But where there is felt ignorance or doubt, contributing to the common interest will entail putting into public space any information that will repair it.

As we have seen, much more of conversation is devoted to the pursuit of issues than at first appears. When that is the case, both the preferences for the new and for minimization assume the force of rules. Doubt and ignorance are also responsible for changes and increase in mutual knowledge. For the outcome of a now dead issue or question will pass into the realm of the reliable, where it may be relevant to any number of further concerns.

I have been using conversation to illustrate some of the consequences of obeying the general social rule: Contribute to the common concern. Being in a position to do so depends on knowing the common concern, with its relevant background. It is always in the common interest to have the same differences shape the common concerns. The existence of a community of interest depends on its members having at least one common interest. But, because an interest in fact resides in an individual, for perfect community the same differences must be made to each member's interests. Two people's interests will be the same just to the extent that they have made the same judgements. Common knowledge and common interest are interdependent characteristics of all human discourse.[8]

Chapter Five: Interests, Obsessions and Images

My central concern in this chapter will be to explore the question of what an interest is, taking for granted that it is a primarily psychological phenomenon. An active interest is liable to give rise to actions and utterances; it is also liable to 'occupy the mind'. My procedure is to examine an extreme case of this tendency, in the shape of obsession. I use the word 'obsession' in an informal, colloquial fashion which only partly coincides with its technical and psychoanalytical uses. I find there is no criterial distinction between an interest and an obsession, an obsession just being an interest which has got out of hand. Using my findings about obsession, I go on to examine the role of 'imagery' in occupying the mind. And in the next chapter, I move on to those specifically verbal occupants of the mind: 'thoughts'.

What I do in the following pages is use my imagination to illustrate points and, in the end, support conclusions. Without years of unfocused introspection, my imagination would have had nothing to work with and it seems reasonable to assume that other people, like me, from time to time notice what's going on in their minds. If nobody rejects the picture I present of an obsession, then that is as good empirical support as I can hope for in making any claims about the internal character of an obsession. And the same applies to the discussion of 'imagery' that later ensues and to any attempt to report or describe introspective phenomena. What I have tried to do is describe

cases which anyone will recognise. (In the next chapter, I use introspection in a different way, which I there describe.) As well as on introspection, I have also drawn on Roland Barthes' perceptive disquisition on amorous passion, A Lover's Discourse. In order to round up as much empirical support as I could for the account I give below, I exposed the first draft of this chapter to as many people as possible. No one proposed any major change, though some modifications have been made in response to comments. To preserve that support (such as it is), I repeat most of the examples in the first draft, below.

* * *

The pursuit of cognitive achievement is not all that occupies my mind; I don't spend all my time thinking hard. Idle recollections, daydreams, possible conversations, all sometimes fill my mind. If one offers cognitive space to useless memories - ones which make no contribution to other current interests - it is presumably because it pleases one to do so. But some memories are capable of besieging one's cognitive space, of occupying it, being displaced, and returning again, and again. It is not a matter of soliciting them, or of being reminded of them, or even of allowing the recollections to flow, but of being unable to stop them. Recollection in such cases has the character of obsession rather than indulgence. Horror, bereavement, and passion (inter alia) are all capable of giving memories this power: they intrude between and even within the pursuit of other interests. While it may be possible to avoid thinking about

them, thinking of them is inescapable.

'But', someone might protest, 'the sort of case you've been discussing can't be called thinking at all. This can have nothing to do with cognition - horror, passion and grief are feelings, obsession is inherently irrational. Besides, of what do these "recollections" consist? Can they not be mere images? And can even having a succession of related images count as thinking?' These are good questions, and there is clearly a lot of truth in the view proposed: certainly feelings and images are (part of) what we're talking about. But a picture in which obsession was presented as a radically distinct 'non-cognitive' phenomenon would be a misleading one. In the discussion so far, obsession has been distinguished as untriggered, unwelcome, unreasonable, purposeless, and image-laden. But need it be any of those things? Is any of them a necessary condition for an interest to count as an obsession?

It is the nature of obsessions that they assail one's mind even when there is no external trigger. But it's also typical of an obsessed person that almost anything, a word, a song, a place, will arouse the obsessing interest. The likelihood of arousal is a function of the intensity of the interest; whether arousal can be attributed to some cue or not is neither here nor there in distinguishing obsession in particular from interest in general. Equally irrelevant to distinguishing obsession from other cognitive activities is whether what obsesses one is welcome. Take, for example, amorous passion. Imagine someone deep in a happy love affair. Recollections of her beloved's smile, look, touch, pervade her mind, and indeed her body. Images of certain

moments recur; not only does she welcome them, she dwells on them, lets them bring further recollections in their train. Now, suddenly, the affair ends. But the obsession doesn't pass: the same moments return, the same images beckon - only now they just make her miserable.

> "From the start, greedy to play a role, scenes take their position in memory ... This theatre of time is the very contrary of the search for lost time; for I remember pathetically, punctually, and not philosophically, discursively: I remember in order to be unhappy/happy - not in order to understand." (Barthes, p. 217)

Passion is of course a paradigm of unreason. However, even in the case of passionate love, what pervades a person is not chaos, is not random. The recollections, images, thoughts, songs, that assail the lover are related to each other - are coherent - because the passion they manifest relates them. Although they arrive as the material of joy or grief, not in order to abet the understanding, they can also be material to understanding: they are information. As they contribute to the hope or the despair, so they can contribute to the understanding, to 'digesting' the truth. It would be wrong, therefore, to claim that obsessions need be divorced from reason.

Futhermore, not all obsessions are as unreasonable as passion, they may equally well be by their very nature rational. It is perfectly possible, for example, to be obsessed by a problem, the proof of that obsession being an assault not of images but of relevant questions, possibilities, and

considerations. Being a problem, it wants solving, and the cognitive frenzy is addressed to that purpose. So, obsessions need not be unreasonable, nor purposeless, nor image-laden, any more than they need be unwelcome or untriggered. Perhaps, however, some or all of them may turn out to be sufficient conditions for an interest to count as an obsession.

Hope, despair, expectation, desire, any of these may be unreasonable without being obsessive. In a sane person, the less reasonable an interest is, the less likely it is to occupy and reoccupy the mind (that lovers are mad is well known). (What 'reasonable' means will, I hope, appear in the course of the next chapters.) As for being image-laden, even writing a shopping list can be that: as I write mine, I view successive shelves of possible purchases in my mind's eye, to see if I want any of this sort of thing or that. Few activities can be more unlike obsession. Likewise, the need to draw up a shopping list, and indeed do the shopping, often strikes me as distinctly unwelcome, whilst in no way functioning as an obsession. And if idle recollection weren't purposeless it wouldn't be idle. The capacity to occur to one or strike one without a trigger will bear a more protracted discussion.

A 'trigger' is something outside an interest, which arouses it at a particular moment. So an 'untriggered' interest would be one that was aroused without any apparent immediate external prompt. But this is not as clear-cut a distinction as it seems, essentially because arousal is not an either-or phenomenon, but one of degree.[1] An interest that was not aroused at all would be much less likely to be triggered than one which was already

highly aroused. Anyway, it is obviously not only obsessions that enter one's consciousness without any reminder. Let us suppose I am lying in bed waiting for sleep. The thought "Tomorrow's Thursday" occurs to me. My eyes are shut, my mind nearly empty: there is nothing one can point to that aroused my interest in its being Thursday tomorrow. Only if this thought - or other thoughts about it being Thursday - has been recurring all day, striking me again and again, will I count as obsessed by it. But what if I think that, or related thoughts, say, half a dozen times in the course of the day? It's surprisingly often, but surely not obsessively... It becomes increasingly clear that there is no hard and fast distinction between obsessions and other interests. An obsession is just an interest at such a high level of arousal that it's perpetually ready to occupy one's mind, to recur.

Now, say I'm wondering what's on the television tonight. Something distracts me and I fail to find out for a few minutes, then the question recurs. I can't lay my hands on the paper, and the phone rings. The soup boils over, I clear up the mess. Then I remember yet again that I want to know what's on, etc., etc., until my interest is satisfied, the information assimilated. (It's no good just reading what the paper says, if I haven't taken it in, then I'll have to look it up again - because I still won't know what's on the telly.) Once I know, then - obviously enough - I won't want to know; but until I know the answer the question is liable to recur. Similar patterns obtain for wants in general and their satisfaction, for plans and their execution, issues and their resolution, problems and their solution.

The capacity of interests to recur is a function of their strength and the rapidity of their satisfaction. So, if I'm mad keen to have a banana, say, and there isn't one in the house, the thought of eating a banana will occupy and reoccupy my mind until I finally lay my hands on one. If, on the other hand, I only quite want a banana and find none in the house, if I'm sensible I'll forget all about it - I'll lose interest in the idea of a banana. Unlike even the most extreme case of banana-wanting, while it rages, passionate love is insatiable: its hunger can't be quenched, only when the passion fades will the hunger go. Hence an obsession is much more likely to develop out of amorous passion than out of wanting a banana. Equally, a painful bereavement or loss cannot be assuaged, only fade, and is liable to assail its sufferer for years. And a really difficult and interesting problem - such as the relation between language and thought - may occupy and reoccupy a mind with great frequency for years, as this thesis testifies. So, an obsession is an interest which is both strong, and - for one reason or another - not susceptible to swift satisfaction.

* * *

Let us review the picture of 'the mind' so far. What we seem to have is a number of different interests of different strengths and at different levels of arousal. Whenever an interest is 'occupying the mind', let us say it is at peak arousal. The strength of an interest can be regarded as a function of the frequency with which it reaches peak arousal,

itself presumably a function of the level(s) of arousal that it maintains when it is not at peak. At the opposite extreme from obsession, a faint interest requires a potent summons to rouse it from the borders of complete inertia, whilst a defunct interest can by definition not be aroused at all (if it's defunct, it's extinct). The picture given is a dynamic one in at least this respect: the interest at peak arousal at a given moment need not be the strongest interest, but strong interests are by definition bound to achieve peak arousal sometimes. Therefore, what is at 'the forefront of one's mind' is bound to shift.

What is it for an interest to be at 'peak arousal'? to 'occupy' or be 'at the forefront' of the mind? It is for images, thoughts, utterances, or actions (including listening and looking) relevant to that interest to occur. What it is for an action or utterance to occur is perfectly clear - they can be witnessed. But an image that occurs to me is a scene only I can see, a thought an utterance only I can hear. That we 'see' things in the 'mind's eye' is a commonplace - and in the first draft of this chapter I didn't query it beyond in one context suggesting that 'in the mind's body' might be a better phrase. And I was willing to accept the translation of 'picture' for 'image'. Subsequently, comments from people who read that draft, plus further consideration of the cases in it, plus some more determined introspection by myself, plus a little 'research' have all conspired to make me reject that translation.[2] Several people commented that they hardly 'had' images at all, or even that they never did; and I now know that a wide variation between individual reports of imagery has long been recognised in

psychology (see Bartlett, 1954). However, everyone seems to be able to comply with a request to imagine that they're in some spot familiar to them and then report on what direction this or that may be in relation to them. This holds even for those most resistant to the idea that images play any role in their mental lives. (That was the 'research' just mentioned.)

When the imagination provides one with an image, it is more a matter of having a scene before one, than of seeing a picture. Sometimes one is just a witness to the scene - which may be quite static, or not - and sometimes a participant; in either case, much of the time, I believe, the images flit by unnoticed. Certainly in my own case, until I started attending to the contents of my mind, I believed that my mental life was almost completely void of imagery. But now I've noticed it, I know that imagery in the sense described is a background accompaniment to a great deal of the stream of mental chatter I've always known about. Unlike those internal verbalisations that I've been dignifying as 'thoughts', images are hard to 'get hold of': because they're not verbal, we can't trot them out just like that, either for another person or for ourselves. It takes an artist to make these mental representations public, to fix them and make them last. Whereas anyone can put a 'thought' into graspable form.

Before I began thinking about it for the purposes of writing this chapter, I had unreflectingly assumed that my private verbalisations were my only overt mental phenomena worth noticing. I was also convinced that language was the sole cognitive means of getting away from the here and now, and hence

the only vehicle for serious thought. Some of my reasons for shedding those prejudices must already be obvious - what else does an image do, other than take one away from the here and now? (In fact the rest of this chapter is devoted to answering that question.) Despite these changes in my viewpoint, my central working assumption in this chapter and the next vis-a-vis internal verbalisations has not changed. Namely, that they are in effect utterances that only their instigator can hear, that we should think of ourselves as having 'mind's ears' as well as 'mind's eyes'. If the thought of hearing things with the 'mind's ear' strikes you as implausible, let me suggest the following 'mind exercise'. Take a line from some song you know, and first say it silently to yourself, then sing it, equally silently.

Those quasi-utterances I have been calling 'thoughts', in my own case at least include a great deal that doesn't merit the name, such as repetitions, exclamations, curses, and old saws (see next chapter for details). Accordingly, from now on I will generally substitute the expression m-utterance for what I have hitherto referred to as 'thoughts'. In parallel fashion, the expression m-action may stand in for 'image '. Viewed this way, 'images' and 'thoughts' appear much less radically distinct than they seem to be at first. Both are m-activities, occurring as it were 'in the mind's body', which quasi-acts in what I shall henceforth call cognitive space. Amongst these m-activities, m-utterances tend to draw the attention, as speech which is addressed to one does in public space. If this view seems far-fetched, then read on for the further considerations which changed my mind, and may change yours.

* * *

We shall now review the role of imagery in one's mental life, in order eventually to determine the distinctive role of m-utterances. The most obvious thing an image does is make it possible to perceive something which is not in fact present, it can be something in the past or in the future, or be purely imaginary: an image has the power to present these, indifferently, to the mind's eye. Consider the following image: in her mind's body, the thinker is passing through a door. Now let us imagine some contexts in which this image might occur.[3]

First let us suppose that our heroine has an important interview the next day, which is a nerve-racking prospect. She is trying to take her mind off it with a good book, but every now and then, the vision of herself stepping into the interview room assails her. Whether or not she has ever seen or been through the particular door she will pass through tomorrow, and whether or not she remembers it if she has, it is still an image of passing through <u>that door</u> tomorrow. Its faithfulness to the real door she will go through is completely irrelevant to her concern.

Each time the image hits her and she's going through that door again, anxiety floods over her once more. The image <u>reminds</u> her of her worry, of the worry that presents her with that image. It's as though the image yanks the whole thing into view, by making its presence known (again) to the poor lady. Luckily it's a good book, and after a while it grips her to the point at which the image ceases to recur. This, I think, is closely parallel to what Barthes refers to as 'remembering, punctually, pathetically

... in order to be happy/unhappy, not in order to understand' (ref. above). The anxiety is at a very high level of arousal, pressing to enter consciousness, flinging the image to the forefront of the mind, ready to occupy it entirely. The image functions here as vanguard.

Now imagine our heroine is trying to remember the lay-out of a house she has visited. 'In her mind's eye' she begins in the most familiar room, visualising two doors leading from it. She then goes through one of them, shutting it behind her and leaning against it to view the scene, and see what exits if any may reveal themselves from there. From time to time she mutters to herself, "Now let me see, if I turn to my left..." "...right", etc. Here the image (one of a sequence) is required to be faithful, as is each image in the sequence - though only in respect of relative position and not in any other respect. For it is with relative position that she is concerned. Also irrelevant here is the point in time at which this mind-trip occurs: that it could do so at anytime is part of the exercise. Unlike the door image brought on by anxiety, this door image is timeless. A further difference between this and the last case is that whilst we must say that the image occurred to the thinker in each case, in this case it occurs to her because she is looking for it.

In what respects, if any, an image is required to represent the actual depends on the interest that gives rise to it. Also, in what respects, if any, it bears on past or future can be recognised not by any feature of itself, but by the nature of the interest that gives rise to it. Recall that the very same images

signalled both past and future to the empassioned lover whilst her hope still flourished; but once her hope had become despair, they told only of the past. If we want to interpret an image, the interest from which it springs must provide the context in which we do so.

Given the difference in contexts between the two examples above, how differently does the image function in either case? Let us call the first case, and all cases in which an image is 'pushed out' rather than summoned, expressive. For in its simplest sense, to express just does mean to push out - though what is expressed in this sense only gets as far as cognitive space. Once an image has slipped into cognitive space, the interest it's an expression of, by definition moves from a high level of arousal to a higher one, namely to peak arousal. That interest, for some period of time however brief, will now occupy a mind taken up a moment before by some other interest; and each time this happens it makes the thinker once more aware of her concern. In this case, the image both expresses and arouses the interest, but makes no difference to its outcome.

Let us call the other case, in which the thinker steps through the door in order to jog her memory, purposive. Because the interest in this case is a purpose, it exerts a certain demand on the thinker to make a difference to its outcome. Her thinking of the door itself does that, and as she steps through it, she reminds herself of what's on the other side. How does the image help her do this? Thinking of the door makes her aware of having a specific interest (part of the general one) in What's through that door. As she passes through it, looking for the

answer, the image functions to bring that interest to peak arousal, so that what she knows about what's through that door will surface, become visible to the mind's eye. Because images function in this way, the interest is urged on by the image, only in that case the interest needed no urging. An image which serves a purpose also expresses and arouses that purpose, as well as making a difference to its outcome.

Although it was true in those two cases, as it is in many others, that an image serves to arouse its initiating interest still further, that is not necessarily the case. For example, when I lose my temper I frequently perform m-actions of extreme violence on the object of my anger. The glimpse that affords me of the consequences of such violence at once makes me aware that I don't desire them, and thus actually dampens the anger rather than urging it on. It also changes the form of the anger - makes a difference to its outcome - by closing off possibilities (those rejected violent acts) rather than by opening them up. Despite the differences between this case and the other ones, the role of the image in every case may be summed up as that of having a feedback effect on the interest which gave rise to it.

With an image, we can represent a possible state of affairs to ourselves and thereby find out how we react to it, how we feel about it, - although it may make stomachs churn, palms sweat and hearts beat faster[4] - without having to suffer its actual consequences. Roughly speaking, the question an image may usefully evoke a response to is, Does this feel right, is it what I want? That embraces the purposeful cases above like remembering the layout and planning the shopping, as well as the

rejection of violent possibilities in the last case. Even when the image is of some past event, it seems to me that one is intrinsically aware of its desirability or otherwise, in being aware of it at all. As well as enabling reactions of this kind, the feedback effect also makes a difference to the level of arousal of any interest affected.

When the level of arousal of an interest is raised, that has both a long-term and a short-term effect. The 'long-term' effect I hypothesise, lasts just as long as the affected interest lasts. When not in action, an interest will tend to sink gradually away, maintaining a lower and lower level of arousal as time passes.[5] Since the more aroused an interest is, the more likely it is to respond to relevant input, every occasion of its arousal acts to strengthen it by retaining it at a relatively high level. Input thus prolongs the active life of an interest. It also serves to keep the assumptions which inform that interest relatively available, which brings us to the short-term effect of arousal. The more highly aroused an interest is, the more accessible is the information relevant to it, so information relevant to an interest which is at peak is maximally accessible. This follows quite automatically - the interest cannot leave its form behind.

A combination of those two aspects of the feedback effect is what makes it possible for an image to make a difference to the outcome of an interest. In being an expression in cognitive space, an image becomes something one can react to: if this were not so, then how could it effect any change? In order to effect any permanent change in the form of an interest - i.e. a change which is henceforth taken for granted - then the m-action in

question must be assessed. If it is to affect the future course of my actions in real space, then it had better be well-motivated. Take a simple case: m-activity when compiling my shopping list brings me to the cereal shelf; yes, we need cereal, Muesli - which gets added to my list. In order to assess the cereal shelf to see if it contains anything I want, I must obviously bring relevant information to bear. And in order to do that, the relevant information must be available.

So, non-verbal m-activity can do a good deal to assist the business of thinking about something. But it has its limits. In the first place, as I remarked above, it is hard to get hold of or fix - an image is evanescent. In my case at least, no exact, reliable, repetition seems possible either, and I cannot hold it before my mind's eye and gaze at it.[6] It is not therefore possible to produce a cool, measured assessment of it; rather, it's an intuitive, 'gut-reaction' type of response to how the image feels. It does appear that some people are better able to grasp or fix their imagery than I am, and therefore better placed to produce a cool judgement of it. However, even relatively expert image-manipulators find that the second limitation applies to them. Namely, that in forming an assessment of an image, it is only the information which has brought its initiating interest to its current form which is reliably and effortlessly available.

It certainly is not necessarily true that an image will feed back only into the interest which gave rise to it. But it does seem to be abnormal for an image to arouse an unrelated interest, and idiosyncratic and arbitrary when it does so. To take an actual example, the other day cooking something in a tightly

lidded pan, and wanting to know how it was doing without lifting the lid, I imagined the pan with transparent sides. Immediately, after an interval of some twenty-years, I recalled an ad for bouillon cubes that I'd once imagined, in which one watched a whole cow through the sides of a glass pan, being magically transformed into a little cube wrapped in shiny paper. Clearly even an interest so weak as to have lain dormant for years can be revived, rearoused, by an image as distinctive as this one. But having these memories revived after all that time, though quite pleasing, was entirely useless, and the connection briefly made entirely arbitrary.

It seems intuitively likely that the power of the image to evoke the old interest in the case just discussed was a result of its vividness and distinctiveness. Take again that everyday image of passing through a door. Imagine that our anxious interviewee, who has earlier been trying to recall that lay-out - using 'the door image' - is being assailed by her worry again, heralded by 'the door image'. In this context it seems extremely unlikely that the image would rearouse the interest which gave rise to its twin. But if we reverse the order it's a different story. Imagine our heroine has been being bugged by her anxiety all evening, with the door image ushering it in each time, and then tries to divert herself by remembering the layout of that house. It seems much more likely that the door image in this context will spark off the worry again. It seems reasonable to assume that the capacity of an image to arouse an interest independent of the one that gave rise to it, is a function of the distinctiveness of the image on the one hand, and the level of

arousal of the independent interest on the other.

So, the basic limitations on the cognitive scope of imagery are these: it is hard to fix, examine, scrutinise, hence not susceptible to cool assessment; and it exerts no automatic call on information which has not already been perceived as relevant to the interest in play, hence it is not susceptible to assessment in the light of more general information. However 'cerebral' the interest which gives rise to an image, judgement of what that image presents is radically interest-dependent and subjective. Often that's good enough, but it isn't if one wants to decide if something's objectively plausible or correct. It means, for example, that when recollecting that layout, if she wants to check the plausibility of her image of what's through that door, she's going to have to make a considerable effort. On the whole, if it 'feels right' she'll just assume it is right - in much the same way as the image of a banana may strike a hungry person as 'feeling right' ('Yes,' she m-utters, 'a banana would be just the thing.')

However, if she wants to make certain that she's right, or if she's in doubt anyway, she's going to have to check it for consistency with information about other bits of the house. A totality in which exits are also entrances, in which stairs that go up come down, etc., i.e. one that is self-consistent, would be the best possible confirmation. So, in order to check the plausibility of her image, she's going to have to pull out information more general than that which has already guided her answer and created the image of what's through that door. That other information does not automatically present itself although

it's readily available when specifically called for. If I put myself in N's place, it seems to me that marshalling this relatively wide range of information is impossible without issuing a verbal summons. Be that as it may, it is clear that if one is concerned with likelihood rather than desirability, then imagery is a much less useful cognitive tool. However, if, as Sperber and Wilson suggest,[7] making sure that we're aimed at an intersection of the desirable and the likely is a central cognitive endeavour, then imagery obviously makes an extremely useful contribution. As we shall see in the next chapter, m-utterances are either free of or less susceptible to, the limitations that I have been outlining which apply to imagery.

* * *

Let us now review the picture of 'the mind' which is suggested by the discussion so far. It is a dual picture, with, on the one hand, a collection of busy, partly interrelated, interests at different levels of arousal, and on the other, cognitive space in which m-activities arising from those interests occur. The interest system is both productive and reactive, creative and receptive. In its receptive aspect, an interest is what is changed by incoming information, and hence is a locus of information. Therefore, new input is always received and assessed against a background of the taken for granted, and even the most visceral reaction is an informed one. Equally, in its creative aspect, all the output of an interest will be guided by the assumptions which inform it. Amongst the output of an

interest may be material in cognitive space. By performing m-actions and m-utterances we create objects in cognitive space which can then function as input into the interest system, i.e. feedback.

However, the picture is less simple than I've been pretending so far. For one thing, no interest is isolated: all interests contribute to the complexity of the interest system, within which they are related in diverse ways. For example, interests may be related by partial identity (overlap), from one shared feature to many, with a complete overlap being the limiting case of perfect identity: two interests are the same just to the extent that the same differences have shaped them. Or interests may be related because the outcome of one makes a difference to the outcome of another, which may include the effect of one ruling out the other. Or interests may be unrelated except by the most indirect and lengthy of routes. The network which is established by these interconnecting interests must be extraordinarily complex, and activation can pass through it quite independently of m-activities in cognitive space. The different levels of arousal of its various parts in a given individual at a given moment, will depend on the history of that individual. And that will be determined by an interaction between what matters to that individual, what s/he cares about, and input from cognitive, public, and physical space.

The notions of 'peak arousal', and of 'accessibility' are also more complicated than I have so far allowed. For example, I have been speaking as though there would naturally be just one interest at 'peak arousal' at any given moment. Yet, when I put

forward the idea of 'peak arousal', I suggested that the occurrence of actions and utterances, as well as m-actions and m-utterances relevant to an interest, might count as evidence that it was at 'peak': the idea being that an interest at peak was so highly aroused that it required to be expressed, and its expression could take any of these forms. If there are four modes of possible expression, could four distinct interests be at peak at the same time? To the extent that one can do things and say things whilst one's 'mind is elsewhere', then it may be possible for four distinct interests to peak at the same time. However, for myself, I find it hard to create m-actions and m-utterances from different sources at the same time. But driving a car, thinking about dinner, and chatting with a passenger are possible contemporaneous pursuits in a pinch. M-utterances and utterances provide more competition for each other than m-actions do with actions, perhaps because utterances exert a greater demand on the attention. At this level the phenomenon we are dealing with is attention - and its sibling, effort (see Kahneman, 1973) - and the capacity to divide it. Giving something one's undivided attention is a matter of having only one interest at peak.

A further complication to the notion of 'peak arousal' is that when a given interest is at peak, then its most closely related interests will (presumably) also be highly aroused. Where do you draw the line? Since the form of an interest may be seen as a function of its place in the total network, there is not generally any all round cut-off point. But that helps solve another problem with the notion of 'peak arousal'. Namely, if an

interest is at peak arousal when it produces, say, an image, then what is involved in its being further aroused? If information relevant to an interest at peak arousal is immediately available, it seems reasonable to propose that when it is further aroused then more information becomes immediately available. So when an image feeds back into its initiating interest, it's like the effect of a stone sending out ripples: the bigger the splash the wider the spread.

Another consequence of an interest being aroused beyond peak arousal is that a sequence of further relevant m-actions and m-utterances may follow the first one, i.e. the level of arousal remains beyond the peak 'threshold' for something longer than a brief instant. This would follow from the tentative assumption I made above, that level of arousal will gradually decline,[8] combined with the assumption that an interest may be aroused way beyond the productive threshold at which it 'peaks'. But we must also, it seems, postulate an upper limit on level of arousal. Otherwise the feedback effect could result in such self-stimulation that no other interest would ever get a chance to take over. Of course, something like this does occur in the case of obsessions, but it is not the usual case. A contributory factor in this may be that the feedback effect need not always obtain, it may depend on a slightly higher level of arousal than the peak threshold for a measure of attention to be paid to what is produced. As my experience of the effortfulness of introspection makes abundantly clear (see also the beginning of next chapter), we most of the time pay very little attention indeed to the contents of cognitive space.

The findings of this chapter are very speculative, but they are the working assumptions on which the rest of this thesis is based. If they lead to false predictions or contradictions then I will have to discard them. But so long as they go on working, I shall go on using them. The central assumptions are these: that interests are fundamental psychological phenomena; that we all have a great many; that they have a duration which may vary from a brief instant to a lifetime; that they are both creative and receptive; that they are loci of information; that they are richly interconnected so as to form a network; that they are normally at widely different levels of arousal as a result of an interaction between activation from within and input from without; that a high level of arousal leads to productivity, i.e. expression; that there are four possible expressive outlets - actions and utterances, m-actions and m-utterances; that input to an interest may be from physical, public, or cognitive space; that input may affect the form of an interest as well as its level of arousal.

In the next chapter I focus on the question of what special contribution m-utterances may make to this general picture.

Chapter Six: M-utterances

In the last chapter I used illustrations which seemed plausible in order to explore some questions about the mind. In this chapter I draw on actual material I have caught passing through my own cognitive space, in order to pursue some of those questions further. That is, I use introspective reports below, rather than telling introspectively plausible stories. So, whether my reports strike you as plausible or not, they are as true as I can make them. In fact, there are a number of difficulties in reporting introspective data, and I shall begin by sketching those.

One difficulty is that a good deal of what introspection reveals is imagery. As I stressed in the last chapter, in my case at least, imagery is inordinately evanescent, fugitive; so even though it seems to accompany a very great deal of my thinking, attending closely to it is not possible. Furthermore, in order to report it, I have to turn it into words, there is no other way to convey it. M-utterances of course do not suffer this distortion: I can in principle tell you exactly what m-utterance was running through my mind at a particular moment simply by repeating it in public space. And presumably because they have this precisely recoverable form, m-utterances can be grasped, pinned down, examined. It thus becomes possible for actual m-utterances to be the subject matter of this chapter.

Despite those advantages of m-utterances, it is not as easy as you may think to repeat private sayings faithfully in public

space. The main practical difficulty is the sheer speed at which material whizzes through (my) cognitive space. M-utterance succeeds m-utterance very much faster than utterance succeeds utterance, which is presumably because the unvoiced is not constrained by considerations of clear articulation and audibility. Even a hectic gabble doesn't equal the speed of my 'stream of consciousness'.[1] Also, perhaps because of the volume of cognitive material I create per minute, I find that I can remember my precise thoughts only over a period of up to about thirty seconds. M-utterances which, for one reason or another, are worth remembering, tend to keep coming back until dealt with - the rest just slip away. I seem to be able to recall what I've been thinking <u>about</u>, i.e. what interests have been in play, over a period of hours rather than seconds, without being able to remember more than a very few of the huge number of unvoiced formulations that have occurred to me during that time.

A further difficulty is that I don't usually pay much attention to the events in cognitive space, they flow by unregarded. Plans, putative conversations and writings, and self-indulgent daydreams may all be exceptions to this, and can all be focal cognitive objects, i.e. be the centre of attention, in the normal course of things. The more busily I am pursuing interests in physical and public space, the less likely it is that I will be noticing the contents of cognitive space. And the more demanding those outside interests, the less attention or effort will presumably be left over for internal business. It follows from this that I am most likely to introspect when there are fewest outside demands, and the mind is presumably at its

busiest. This is liable to have a distorting effect, in that the picture of cognitive space which emerges is liable to be of it at its most active. What is more, there is no guarantee that anything is occupying my mind the rest of the time when I'm not noticing it.

The introspective technique I adopted was designed to minimise the problems I've just outlined. What I did was the mental equivalent of (technically rather incompetent) stop-frame photography: every now and then at arbitrary moments, over a period of several weeks, I would ask myself what I'd just been thinking, and note down a few seconds' worth of the cognitive stream. As well as carrying out these self-instructions, I found that my generally aroused interest in the contents of my mind made me much more generally aware of them. In my own case at least, the cognitive stream is almost incessant; there are some physical activities so absorbing or demanding that cognitive space is void whilst they're going on, but they are few. The frequency and volume of cognitive material undoubtedly varies greatly between individuals, but I would contend that everyone has some material in cognitive space most of the time. For anyone who doubts this, I suggest the two following mind-experiments. One is to employ the technique I just outlined, and catch a few seconds of your own cognitive stream as often as it occurs to you to do so. The other is to attempt to empty your mind: for any but the most practised meditators the result of trying to do that tends to be an awareness of just how much there is to empty out, and just how fast one goes on producing more.

What justifies my using introspection despite its many drawbacks? - How can I hope to get any closer to answering my central question of what role language plays in thinking, if I don't dare examine actual cases?

An essential presupposition of this chapter is that imagery and m-utterances are on a par with each other in several respects. They have a common origin in the interest system; they are both primarily expressions of specific interests; they can co-occur in cognitive space, just as actions and utterances can co-occur in social space; and they can both have a feedback effect on the interest system. Let us begin by testing those assumptions against an actual case of a sequence of m-activities.

Example 1

"What do we need?" I said (silently) to myself, "coffee, birthday present for Roy, money (NB children's lunch money). A toy for Roy, Habitat?" Here a fleeting image of all the toys from Habitat being swallowed by the set of all Roy's toys spread across his sitting-room floor passed swiftly before my mind's eye. Coinciding with it was, "too near home, he'll have everything. We've got enough coffee for breakfast." At this point I got up and checked cash flow: enough not to worry. Then an image took me down the hill from the school, towards possible non-Habitat sources of toys. "Post Office? No. Boots? surely we need something from Boots... Cat lit? No. Cat food? Well... Damn. Have to wrap the present and bring it to school at pick-up time. Wrapping paper! Smiths. I'll go down the hill with Jenny and lend a hand with the kids."

At the end of this sequence of m-activities, I had settled - though I don't think I told myself so - that I'd squeeze in present-buying first thing, and coffee-buying at the end of the school day.

The interests which gave rise to this sequence were the need to make certain essential purchases, combined with the desire to expend as little time and energy as possible in doing so. Both m-utterances and images help me sort the matter out by expressing aspects of it, and feeding back into it. Thus, amongst what we need to get is a toy for Roy, and amongst the possible places to buy a toy in my neighbourhood (i.e., without going far and spending time and energy unnecessarily) is Habitat. But it has only a small toy department, is just round the corner from Roy's house, and Roy has a great many toys already. The vision of Roy's toys swallowing the whole supply from Habitat helped make it clear that another source for his present would have to be found, in short it had a feedback effect on the interest. And note that that image was immediately preceded by "a toy for Roy. - Habitat?" which presumably had the specific feedback effect of bringing to mind the characteristics of Habitat as a source of toys for Roy - which the image then sums up.

So, images and m-utterances are on a par with each other in each of the respects I mentioned above: they are potentially co-occurring creations of specific interests which may contribute to those interests by feeding back into them. In what ways, then, is a m-utterance functionally distinct from an image? The most obvious gross difference between them is that one is quasi-aural while the other is quasi-spatio-visual. Furthermore,

the aural dimension of a m-utterance is strictly rule-governed and precise: the form of a m-utterance is always partly determined by the language in which it is couched. This has many repercussions, which will emerge during the course of this chapter (and beyond).[2]

In the last chapter, I presented the mind as having two components, a network of interconnected interests, and cognitive space. The network as portrayed there was a highly flexible assemblage of interests connected with each other in arbitrary and accidental ways as well as in well-motivated ones. Furthermore, activation of those connections will always partly depend on the prior levels of arousal of the activated interests. In short, the network can make a fairly chaotic and unreliable contribution to cognitive endeavours. Intuitively, (m-)utterances, being potential bearers of truth values, and potential conclusions or premises of arguments, strike one as having a more rational and less unpredictable basis. Despite their common origin in the seething network, their common destination in cognitive space, perhaps images and m-utterances take different routes. Might the mind have three and not just two essential components? - a computational logic machine as well as cognitive space and the network?

Among the m-utterances I have noted myself producing, are moans, exclamations, curses, cries of joy, and people's names. These are all unlikely products of a logic machine, being more like the barking of a dog than like reasoned proposals or conclusions. But these explosive little bursts of excitement form a separate category from that of most (m-)utterances,

namely, expletives. Perhaps what distinguishes them from other (m-)utterances just is that they don't achieve expression via the logic machine. So, let's look again at the example of rational decision-making I gave above.

Having established two priorities - getting the coffee and a toy for Roy - I put my mind to doing them on the same trip. The only place for toys on a natural circuit with the coffee shop being Habitat, I then consider getting the present there, but rule it out. Since there's enough coffee for breakfast, I stop worrying about fitting coffee in, and aim myself towards the other toyshops which aren't convenient to the coffee shop. I then put my mind to finding other reasons for making a separate trip, and remind myself in succession of shops at which I might plausibly want to get something. I focus first on the Post Office - which is first on the route - then ignore a shoe shop and a furniture shop, and focus on Boots. Given certain background propositions, the sequence of thought up until now could quite helpfully be spelled out in inferential terms. But at this point it becomes much harder to do that. If I am proceeding in a deductive fashion, then the reasoning here must involve taking into account every kind of thing that Boots sells, and checking it against what's missing from our shelves. But in that case how come I think of cat food and cat litter, which we don't need, and don't think of kitchen paper, which we do? Well, maybe one could explain cases like this by concluding that I need rewiring...

If m-utterances do occur as a result of sub-rosa logicising, and their relationship with the assumptions which inform them is

inferential, then what is going on in the case of the images which occur, which may be equally informed by assumptions? For example, when our heroine of the last chapter imagined her lover, that image was certainly informed by the knowledge that there was no one in the world she'd rather kiss. But an image is not a proper object to insert in a logical sequence, quite apart from the intuitive implausibility of an image being the deductive result of a desire. It seems to me that what we think of is circumscribed by what we take for granted, our assumptions set limits on the possibilities that occur. Once possibilities have occurred, then they can be assessed; this assessment is an aspect of every case of thinking about that I noted. It is at least as much in the reaction to rather than in the production of m-utterances that the key to their relative 'reasonableness' is to be found.

Example 2.

Going round the supermarket: "Bread" - I take some and put it in my trolley. "Biscuits" - I choose some. "Sugar" (sugar is not in the immediate vicinity, is not to be got next) "mustn't forget sugar. Sugar; sugar, sugar." - and so on.

Example 3.

On my way to a house visited only once before, "Look out for the phone box at the beginning of the lane."

Example 4.

As I put the weights on the scales, "Seven ounces"; as I shake in the walnuts, "nearly" - I slow down as they tip the scale, "that's it" - I stop.

All these cases are peculiar in that I am apparently telling myself something which I obviously already know, since otherwise, how could I tell myself it? It's as though I were two people, one doing the telling, the other the listening, one who knows and one who doesn't know. But actually, of course, there's only me. What is going on here? Evidently, the answer lies in the feedback effect.

The m-utterances I produce in 2-4 are like signposts I put up at points of choice, i.e. at moments in the performance when there are alternative paths to take (quite literally in 3). The instruction feeds back into the interest in such a way as to focus the energies into specific actions, putting a precise part of the network on the alert, priming it, bringing it to peak arousal. If we assume that m-activities in cognitive space take less effort than activities in physical space - which is surely a reasonable assumption to take on board - then the procedure is fully accounted for. By talking to myself thus I flog up that little extra effort and attention: because I can m-hear myself m-talk. I can respond to my instructions as though to another person telling me what to do. In fact, in example three I

actually am talking to myself in another's words, I have been told to 'Look out for the phone box at the beginning of the lake'. Examples 5 and 6 are explicit cases of the schizoid relationship implicit in 2-4.

Example 5.
"Come on Dinah! Don't just sit there gazing into space, get a move on!"

Example 6.
"Dinah Dinah Dinah... what have you done now?"

In response to 5 I frown and sigh, m-utter Example 3, but after a moment get to my feet and get on with things: it's a vigorous call on a pretty slack interest. In response to 6 I put my head in my hands, frown, sigh and shake my head - a rhetorical response to a rhetorical question. Apart from making the dual embodiment of speaker-hearer clear, what does it add, addressing myself by name? The contexts in which I tend to do it suggest that it sums up my social self for me - after all Dinah is what other people call me (is what I am called).

In all those examples, i.e. 2-6, I address myself with the voice of authority, and in 5-6, society adds its weight. The suggestion that I sometimes address myself in the voice of society is borne out by the irritating frequency with which I m-utter ungainsayable old saws like:

Example 7.

"A stitch in time saves nine."

As well as speaking to myself on behalf of society, I also address myself in the voice of temptation, much as in cartoons:

Example 8.
"Why not give yourself another five minutes? Go on ..."

which followed immediately after 5, in a classic mini-debate between the good and the devil. Since the devil has no authority, she has to rely on her powers of persuasion. As well as playing these rather abstract roles, I sometimes address myself in a particular person's voice - and sometimes I argue with it in my own persona (see also below). All this is only to demonstrate that of the dual roles of m-speaker and m-hearer, the m-speaker can adopt more than one persona and more than one point of view, and what the adopted persona is will make a difference to how the m-hearer responds.

So far we have only discussed variation in the m-speaker's personae, but it is equally possible - and, in my case, more usual - to m-address different people, and therefore to respond in the personae of different hearers. Here is a small but quite typical example:

Example 9.
Looking at a hot-water bottle stopper, "Twisted, it'll break soon." Then in cognitive space I am before the chemist's counter, holding out the stopper, and saying, "I bought this

hot-water bottle from you last Christmas." Then, deflated, no longer facing the chemist, "it was cheap, it was dirt cheap."

From my point of view, the bottle, being a fairly recent purchase which is shortly to become a useless one unless I can replace the stopper, is a threatened but still precious object. From the chemist's point of view it's one of a cheap lot of Chinese imports; and as I imagine the confrontation so his point of view comes to me: I hear what I have said through his ears. Sequences of m-utterances like this one are scattered through the day. That is, being struck by something about which it would be appropriate to speak to another, I proceed to try out ways of 'putting it'. How I put it, i.e. in what words, in cognitive space as in public space, depends on who I'm talking to. Since more often than not I am quite straightforwardly talking to myself, that also has its effect.

Though in the case of Example 9, I don't try very hard because I reckon I'm on a losing wicket, sometimes I come up with, and reject many formulations. For example, I may plan quite carefully the opening remarks of a delicate phone call, or the salient points of a tricky request I'm about to make, or I may write and rewrite an important letter in my head. Often, but not always, I am putting a case. A claim that sounds quite plausible in the context of interest from which it arises, may sound extremely unlikely to other ears. The more debatable my case, the more sensitive to potential objections I must be, and the more subtly I must present it. If I want to be persuasive, I must make use of what my audience and I both know, and not take

too much for granted. The less my hearer and I have in common, the more explicit I will be obliged to be. At the furthest extreme from the sort of private self-addressed m-utterances of Example 1 are the m-utterances of which this thesis is composed. For they, dear reader, are addressed to you - and I don't know who you are - and every one of them has passed through cognitive space on its way to the page (and for every one that made it, a dozen bit the dust).

In this section, we've found that a m-utterance can affect me in many of the same ways that actual utterances do. The dual roles of speaker-hearer which I can play for myself, affect both the form of what I say, and the way I respond to it. An aspect of this which I have not so far brought out is that, just as utterances are, m-utterances are produced with varying 'force', and responded to accordingly. For example, Examples 2-3 had the force of imperatives, of warranted injunctions, and were reacted to as such. While in Example 1, I ask myself a number of questions, which I proceed to answer. I also make comments, (Example 4), suggestions (Example 8), claims, indubitable assertions (Example 7) challenges and jokes.

So, let us take a closer look at the different consequences for me as hearer of these different forces. In varying degree, they all exert some sort of demand on the hearer.[3] Indubitable assertions exert a minimal demand, namely for unquestioning assent. A claim also exerts a demand for assent, but if it received unquestioning assent, then it would have the force of an indubitable assertion. A claim is more demanding than that: it requires the hearer to assess its likelihood in the light of what

she knows. A comment invites assent rather than requiring it; in social space it requires acknowledgement, however minimal, but in cognitive space that would seem redundant. Very often, in fact, I find I greet the information on which it occurs to me to m-comment with pleasure or regret, as in

Example 10.
"Oh no, I forgot to get orange juice."

and many others. A suggestion exerts a demand for assessment in terms both of how well it fits one's wants, and of how likely it is to become true. A question exerts a demand on its hearer to try to answer it. Trying to answer a Wh- question involves looking for something which fills the Wh- slot, i.e., fits that want; trying to answer a yes-no question involves choosing one alternative rather than another, either as being more likely (with certainty as the limiting case), or as being more desirable. An imperative exerts a demand to perform the stated actions. And a challenge exerts a demand to justify a claim as likely, a suggestion as desirable and likely, or an action as warrantable.[4]

As hearer I can listen to these various sorts of demands with a variety of different attitudes. I can listen attentively, or with only half an ear; I can listen credulously, as to an authority, or sceptically, as to one of dubious credentials; and I can treat the demands as serious or frivolous, and as urgent or deferrable. All these factors interact with the inherent difficulty of the task demanded, to affect how much effort I will

put in to meet a demand. In every case except that of actual imperatives, meeting the demand involves matters of likelihood or desirability. Once I have decided what is the most likely outcome, or what is the most desirable, then I don't have to think about it anymore, i.e., I don't have to pass it through cognitive space in order to assess it. Instead, I am just ready for a given outcome, or aimed at a given goal. In general, what I am doing when I meet those demands is attempting to get matters settled, so that they may henceforth be taken for granted.

When I try hard to meet those demands, in what do my efforts consist? A strategy I regularly employ is to repeat the question, suggestion, or whatever. I tend to say it two or three times to myself, then pause for the response, and relevant considerations invariably occur. If - the question (say) not yet being answered - I carry on repeating it, then that, inter alia, will do the trick of keeping other interests out. We found in the last chapter, that, up to some presumed ceiling, the more highly aroused an interest, the longer it would occupy cognitive space (i.e., keep irrelevant material out of it) once it got there. That was because of a presumed gradual decay in level of arousal, so that the higher an interest was above the peak threshold, the longer it would take to sink back below it. We also found that m-activities tend to feed back into the interest(s) from which they spring, and there is no reason to suppose that this is any less true for m-utterances than for images. But the expression of an interest in a m-utterance has a feedback effect which is both much more specific than that of an image, and much more widespread.

M-activities invariably arise out of some broader interest, which must itself be aroused for the m-activity to occur. So when, say, a Wh- question - which is itself an interest - arises out of an interest, it is a narrowly focussed peak of a more broadly aroused area of the interest network. By repeating the question, I keep that narrow peak defined, maintain its distinct identity. Without that restricted focus, it will tend to lose its thrust, and its demand may never be met - though if it's a really good question, then I won't need to make myself repeat it, because it'll recur without any effort on my part - the more general interest will keep bringing me to it. How precisely circumscribed an interest is, and how decisively it excludes irrelevant material, must surely be functions of each other. So, when I put a question, insofar as it is made up of words whose meanings are precise in the context, then it will set a clear boundary on what may be allowed into cognitive space.

At the same time as being an efficient means of keeping the boundaries neat, the channels to cognitive space clear, and cognitive space itself vacant, the types of verbalisation we've been discussing also function as ways of expressing demands. So what one of them feeds back into the interest system is a summons, which is to be answered through quite specific channels and presumably, depending partly on different prior levels of arousal of the various interests at the time, reaches only parts of the network which are tapped by the words used, in that combination. When I formulate a question, I make a precisely defined place for the answer, and only what fits it will come forward. Another point about a m-utterance being in words, is

that it taps meanings which exist independently of my whim, which are not relative only to my interest. In short, it connects my interest with the systematic and well motivated body of assumptions encoded in the language (see below, Chapters 7 and 8).

Useful though it is, there is more to hard thinking than simple repetition of the initial articulated prompt: it also involves assessing the material which has been invited to come forward. In the case of answering a question, it is a matter of assessing whether what offers itself fully satisfies the want expressed. For it to do that, it must not only fill the gap, but also meet the requirement of plausibility, it must be both possible and not unlikely. One is aiming at the truth but not going out into the world to check it; effective certainty is the best attainable result. Of course, one may be sure of an answer at once, but it is hard thinking we're discussing, and this business of assessing plausibility and arriving or not arriving at firm conclusions is characteristic of all hard thinking. Let us call all m-utterances that demand to be assessed in this way, proposals.

If a proposal is truly plausible, then nothing that one currently assumes will rule it out. So, in order to assess its plausibility thoroughly, it should in principle be assessed in the light of one's whole range of assumptions - although in fact, of course, most of them will be completely irrelevant. As we have seen, information relevant to an interest becomes readily accessible when it is aroused. So, when an interest in the truth is aroused, presumably the same thing happens, but unlike the

other cases, it is not just a tight self-consistent subset of assumptions that becomes more accessible. Rather, it is anything that strikes one as possibly having a bearing on, and especially supporting or undermining, the likelihood of some particular proposal. Because in practice some areas of interest are more highly aroused than others at any given moment, one cannot count on bringing information relevant to the less highly aroused interests to bear. One may therefore fail to round up everything which makes a difference to this outcome. But one can still try:
- How?

Recall that the spirit in which one greets a proposal has an effect on how one responds to it. In order to be thinking hard about a proposal at all, one must have taken it seriously, but not credulously. In order to be as 'disinterested' or 'objective' as possible in assessing a proposal, one should hear it with the ears of a sceptical stranger. In that way one's own special interests get set aside, and one is obliged to try to take into account only what one assumes anyone assumes. This is presumably possible because a substantial subset of assumptions is always more or less readily available, encoded in the language, irrespective of what interests are in play. (Although, as we shall see in Chapter 8, they are not all always equally readily available.)

Rising to the challenge of performing this act of the imagination, in order to assess one's own proposals with maximum objectivity, has a number of beneficial consequences. Firstly, it greatly reduces the total set of assumptions which may be allowed as support for the proposal and therefore presumably,

also reduces the effort involved in assessing it thoroughly and therefore increases the likelihood of that occurring. Not only is it a restricted set, but it prominently includes a set of well motivated and reliable assumptions with relatively fixed and warranted interconnections, which forms a fairly tight and self-consistent system, i.e. the language. Also, if, as sceptical hearer, I respond to a proposal with a challenge ("Oh yeah?") I thereby oblige myself as speaker to justify it: Being naturally lazy, I will tend to abandon an assessment once I have found some support; this is often a mistake, since the interest which gave rise to the proposal in the first place will still be relatively highly aroused and ready to provide support. The sceptical listener device is a way of guarding against such laziness and the mistakes it can cause.

When I speak of 'assessing a proposal against a background of assumptions', what do I mean? In the last chapter, we saw that when an image was assessed against a background of assumptions, it was in terms of whether or not it fitted the interest which both gave rise to it, and was the locus of those assumptions. If it didn't suit that interest - if it had undesirable consequences - it would be rejected, whilst if it did suit it - if it had desirable consequences - the interest would accept it and be changed by it. Is the picture radically different when it comes to assessing a proposal where the motivating interest is truth?

It's not a general interest in the truth we're discussing, but an interest in whether some particular proposal is likely to be true. So, let us look at an actual case.

Example 11.

If one is listening to oneself with the ears of a sceptical stranger, "one is obliged to try to take into account only what anyone knows, in other words the assumptions encoded in the language."

became: "one is obliged to try to take into account only what one assumes anyone assumes," plus inserting: "a substantial subset of those assumptions is... encoded in the language."

Those changes happened in several stages, at first I was almost completely satisfied with the first version. But then I challenged myself about it, and it began to sound wrong. The first change to be made was replacing 'know' with the weaker 'assume' - a substitution I've been quite consistent about in this thesis for the good reason that in psychological terms, 'know' makes too strong a claim, and rules out mistaken assumptions, which play exactly the same cognitive role as correct assumptions unless and until they're found out. But thus weakening the claim I was making was not enough to make the proposal sound good. It then came to me that the set of assumptions I wished to single out was not the set of assumptions that anyone makes, but the set of assumptions one <u>assumes</u> anyone makes. Those may or may not be coextensive, so I modified the formulation accordingly. But the proposal still didn't sound quite convincing - <u>is</u> the set of assumptions which one assumes anyone assumes identical with the assumptions encoded in the lanugage? Nicely though it would suit me to believe it - it would have been a handy plank in my argument - it is evidently not true when one thinks about it. For counter-examples, such as

every English speaker being able to assume that every other English speaker knows that President Kennedy was assassinated, abound. After wondering a bit about whether there might be some way of tightening the specification of which assumptions were to be allowed in, feedback from a sceptical supervisor plus my own continuing sense of dissatisfaction made it clear that the formulation is even now not good enough. Though no longer false, it remains incomplete. As it stands, it would seem to suggest that one can only employ already shared assumptions in an argument. But that - evidently false - conclusion is not one I would wish to draw. Rather, I must make any that are not shared explicit and obtain credence for them. To 'make an assumption explicit' is to radically alter its status: it ceases to be taken for granted, indeed ceases to be an assumption, and becomes an assertion and hence an intelligible object which lays claim to being true. So, addressing a sceptical stranger obliges me to find out what has brought my interest to its current point, and put that information into cognitive space where it can in turn undergo the same sceptical appraisal. If that still does not suffice to force my conclusion on the sceptic, then I will seek further support, or perhaps modify my conclusion. In the last few pages, you have witnessed this process in action.

By determinedly playing the sceptical stranger to myself, and making the relevant assumptions explicit, incompleteness and inconsistency become apparent and can be dealt with, and new conclusions can be rigorously reached. It's a device for imposing the discipline of objectivity on myself - or at least trying to. The internal consistency of a set of assumptions, and

their pr acticality in pursuit of the interest they inform, make them a good bet - suggest that they are likely to hold good in the light of other interests too - but constitute no guarantee. Equally, inconsistencies between assumptions relevant to different interests may never appear at all if the different interests don't rub up against each other in practice, and if assumptions from neither set are ever made explicit. We are not always interested in the truth as such, much more often we just want to get some project or other to work, and so long as it does we'll be satisfied.

Mutatis mutandis, just the same pattern is found here as we found in the last chapter for images. An initial reaction of acceptance or rejection takes place. In the case of Example 11, an initial acceptance was replaced by a similarly unanalysed rejection, caused by my treating the first formulation with scepticism. Once a satisfactory formulation has been found, the interest accepts it and is changed by it, until then, one cannot move on. Hence, analysis tends to follow in the wake of rejection rather than acceptance, otherwise it's a waste of time and effort. Since, with a proposal, the interest in question is in the truth, one's grounds for rejecting it will be that it leads to a contradiction or falsehood, even though at the moment of rejection, prior to analysis, one may not be able to state those grounds, and may not even have discerned them. Whether one is after the truth, or whether one is after, say, having a good time, the process of assessing possibilities appears to begin with a 'gut reaction'. If an unanalysed gut reaction is possible at all, then why should this not apply just as much to a desire

for the truth as for any other desire?

* * *

The phenomenom I have been calling 'cognitive space' is obviously a close cousin of what gets called 'consciousness'.[6] But there is an ambiguity in the more usual phrase that I have wished to avoid. At the beginning of this chapter, I discussed some of the difficulties of introspection, and averred that I was often just not aware of, i.e., 'conscious of' the contents of cognitive space. So the use of the term 'consciousness' would lead to the apparent contradiction that one need not be conscious of what is in one's 'consciousness'. The fact is that attention is clearly not a unitary phenomenon, but can be divided and distributed in different directions (within limits - see Kahnemann, 1973). What passes through cognitive space generally demands minimal attention, presumably because no disastrous consequences are likely to ensue in the snug recesses of the mind. Nevertheless, one sometimes pays close attention to the contents of cognitive space, and sometimes one devotes effort to controlling its contents. This expenditure of effort and attention has its rewards.

From sexual fantasies and ego-gratifying daydreams to decision-making and hard thinking, the primary reward of actitivies in cognitive space is self-stimulation, or the feedback effect. Except when its expression constitutes its closure, a m-activity strengthens its initiating interest. This I believe it minimally does even in the cases in which I'm not

aware of being aware of it. Just as I pay peripheral attention, without noticing it, to aspects of my surroundings on which I am not focussing, and adjust to them accordingly, so I pay peripheral attention to the contents of cognitive space, and may likewise be affected by them. Because different interests can take turns in cognitive space at far greater speed than they can in physical space, the feedback effect must provide a sort of keep-fit mechanism for a much wider range of interests than could otherwise be maintained in good shape.

Because the interest network into which a m-activity feeds back is a repository of interconnected information, feedback into a particular interest has the effect of making selective information more available. Because they are in a language, m-utterances tap the relatively tightly organised, immensely reliable, and always accessible sub-system of the interest/information network which any natural language constitutes. Furthermore, because they are in a language, m-utterances have a precise, repeatable, unchanging form, in which they can be detached from their initiating interest, treated as independent objects in their own right, and assessed in a range of different contexts.

Because m-utterances are in quasi-speech, I can perform quasi-speech acts with them. Because of the duality of speaker and hearer, I am the hearer of my own m-utterances. So I may ask myself serious questions which require precise answers, and thus oblige myself to think hard. Or I may tell myself what to do or what not to do, and do so authoratively or not, and respond accordingly by acting on or rejecting the proposal, or by

debating the issue. Thus do I become a moral being, capable of choosing what to do or not to do, in a manner which can be justified. Or I may tell myself what to believe, and do so authoratively or not, and respond accordingly, by accepting or rejecting the proposal, or by debating the issue. Thus do I become an intellectual being, capable of choosing what to believe or not to believe, in a manner which can be justified. But I may fail to listen scrupulously, sceptically, to what are actually dubious proposals. Thus do I become capable also of self-deception. Because assumptions are interest-relative, I may never find out that assumptions relevant to distinct interests are incompatible. Thus I may be a repository of so-called 'irrational beliefs' without any failures of logic being indicated.[7]

The rewards of using language to exploit the resource of cognitive space are mixed: through it one may aim to do only what it is wise to do - and fail; and one may aim to believe only what it is wise to believe - and fail.

Chapter Seven: Meaning I

At this point I want to take up a thought I had when I first began to think about relevance. Namely, that the relevant for any sentence will include the person-neutral linguistic meaning. To illustrate that claim, consider the sentence, "Home-made lemon curd is delicious". Even though in almost all contexts that would be a completely irrelevant thing to say, it will always be relevant to the deliciousness of home-made lemon curd. It will also always be relevant to the quality of home-made things (at least one home-made thing is delicious), to the desirability of home-made lemon curd (if something is delicious then it is desirable, therefore...), to the desirability of making lemon curd at home (if one makes lemon curd at home, it will be delicious, therefore...), and so on for anything which also follows from the meaning of the words as combined in that sentence. In order to explore these connections between meaning and relevance, I try below to pick my way through the minefield of meaning using my definition of relevance as a guide.

The phrase 'person-neutral linguistic meaning' itself implies that meaning need be neither person-neutral nor linguistic. There is also an opposite extreme. At it, for example, we may say that to the distraught lover, a pattern, a song, a place, a particular sensation, a phrase, etc., all mean her lover, her broken heart, the whole damn thing: because that is the concern they touch - i.e., <u>arouse</u> - in her. Similarly, when the anxious interviewee imagined passing through the door,

its meaning was quite different from what ostensibly the same image meant to her when she was trying to remember the layout (see Chapter 5 above). Here both meanings are determined by the interest expressed. So, meaning has to do with both the arousal of interests and their expression.

Now take the use of mean in which 'means that ...' precedes a description of some outcome or upshot. Consider, for example, what this rock is limestone means to a geologist, a potholer, a botanist, a palaeontologist, someone who has been given the task of finding samples of ten different rocks, someone who has been charged with the task of removing tons of it in a fleet of lorries, etc. Here too, meaning is interest-dependent in ways which may vary from individual to individual. The consequences of this rock's being limestone vary for different people according to the different interests and attendant assumptions they bring to it. In short, in this sense, meaning concerns what makes a difference to the outcome of an interest.

So, each of those instances of what X means to N (where X may stand for any discriminable event or thing, but N must stand for a person) is interest-dependent in one of the ways in which X is relevant to N is interest-dependent. But there are also other uses of 'mean', of course. There are those cases in which we attribute meaning to a person, as in N means q by A (where N is a person, A is some action - including speech acts - of N, and q may be either an infinitival or a that-clause (those cases in which A is an act of reference - the uttering of a referring expression or the execution of an ostension - and in which q may be any discriminable event or thing, are expanded on later in

this chapter). Lastly, there are those cases in which we talk about what X means (where X may be any discriminable event or thing) without bringing in any particular person at all. Let us see how these other cases of meaning connect with questions of relevance.[1]

First consider the case in which N means q by A. Let us imagine that we see N press a switch, whereupon a lamp lights up. We might then say that N meant to turn on the light by pressing the switch, and that would most likely be true. But if N, say, had wished to do a drawing of a light-switch in an on position, s/he could legitimately deny having meant to turn on the light, and insist that s/he had meant merely to alter the position of the switch. The light's going on or not going on is a matter of complete indifference to N, it is irrelevant. In other words, what N meant by pressing the switch depends on what N's interest was. What N means, (what a person means) explains, or makes sense of what N does. We understand what N does - i.e. grasp its meaning, in the light of the interest we presume N to be pursuing.[2]

Now imagine M lurking in the shrubbery outside, seeing the light go on. To M, N's pressing the light switch and therefore illuminating the window means, say, that N is home, that the moment for murder has come. In this case, what N's action means to M is of course not part of what N meant: N does not know that M is in the shrubbery and knows nothing of M's current interest. Perhaps, on the other hand, N's being home means to M not that the moment is ripe for murder, but that their assignation can at last take place. "When I turn on the study light," N has said to

M, "come and tap gently on the window". In this case, what N's turning on the light means to M and what N means (to do) by it coincide: N knows, or at least hopes, that M is there, waiting and watching; N's current interest and M's are the same; and the consequences of N's action for their common interest have been agreed between them.

Lastly, let us look at those cases of X means... in which the participation of a person seems not to be required, and which are not interest-related in any obvious way. For example, whether N means to turn the light on by pressing the switch or not, the pressing of the switch means the light goes on. Here one seems to be talking about consequences in the world rather than in any interest system: though if there were a divine - universally interested, omniscient - witness, these consequences in the world would result in changes in the form of her interest network. The notion that in these cases X means... can be treated as equivalent to X has certain consequences in the world applies in many instances, for example, to 'If the leak goes on at that rate it means the tank'll be empty in less than an hour'; 'the extreme winter cold means that all warm-blooded creatures either migrate or hibernate', etc. It won't do, however, for Grice's 'Those spots mean measles', or for 'That chord means the finale is about to begin', and many others. Yet it seems to be the same use of 'mean' in all of these cases, certainly each instance in this paragraph can be translated into an assertion of the form If... then ... - yet it's clearly not logical consequences we're dealing with.

A formulation which fits each of the cases in the last

paragraph is this: If one is certain of X then one may also be certain of ... (the conditional in formal logic being just a special case of this). In short, X means Y is a licence to assume Y if one may safely assume X.[3] So it does not follow that these cases of mean are actually person-free just because no individual need be involved. That is, just because the licence to assume is not taken up by any individual, it does not follow that it is not people to whom the licence applies. Rather than being a person-free, ergo interest-free, sense of 'mean', it is (like the linguistic use) a person-neutral sense. Given that the acquisition of an assumption constitutes an effect on the form of an interest system, even this use of 'mean' - like the others I've discussed - concerns interest and is related to relevance.

* * *

I want now to get closer to questions about linguistic meaning, by discussing two cases of quasi-linguistic meaning: smiling and pointing. So, let us consider the meaning of a smile. Following the precedent set above, we can ask three questions about the meaning of N's smile: What does it mean? What does N mean (to do) by it? and What does N's smile mean to M?

Suppose for a moment that we are in the shrubbery outside the lighted window. Framed in the window we can see N, s/he may or may not be alone; then we see N smile. What can we say about what N's smile means? Very little, unless, that is, N is smiling at us out here in the garden. Assuming that N is actually

oblivious to our presence, we may suppose that N's smile means at least that for a moment s/he is pleased, glad, happy about something. We understand the smile as an expression of pleasure: if it is not such an expression, then we will say either that it was simply a twitch of the muscles, or that it was a false smile. As mere witnesses of the smile, we are in no position to tell; but if it was just a muscle-twitch then we were wrong to attribute meaning to it, while if it was a false smile, to know its meaning we need to know what N meant (to do) by it. If, on the other hand (as is most likely), the smile was real, then to understand it fully we need to know what N was smiling about. We need to know, in other words, what interest it expressed.

If, in fact, N's smile is directed at M (now in our field of vision), then we at once take a step nearer to grasping the meaning of the smile. It is an expression of friendly interest in M, an indication of pleasure at M's presence. Again, if it does not mean at least that, it was not truly a smile at M, but either a false smile at M, or a private smile. If it was a purely private smile, then M's presence is irrelevant to its meaning. To know the meaning of a private smile, the only recourse is to ask the smiler 'What did that smile mean?'

In response to this challenge, N may refuse to disclose what it meant, or may agree to, but is not free to deny that it meant anything without denying that it was a smile. Suppose N is prepared both to acknowledge and to explain the smile, that is to tell us what it meant, what might s/he say? "I was just remembering the wild flowers round Avebury last July" or "I was thinking how warm and happy the cat looked", etc. To know what

that particular smile meant, one must find out what interest was at that moment occupying N's mind. Within its vague meaning as an expression of some degree of delight, a smile may have a very specific meaning. Note, however, that this is not what N meant (to do) by it. N meant nothing by that smile, just so long as it was a purely private smile. N may not have meant to smile at all, may not have known s/he was smiling - yet still be able to vouchsafe the meaning of the smile when asked. But even now we know that N's smile meant that s/he was remembering the wild flowers round Avebury last year, unless we know just what those wild flowers meant to N, then we still know only a part of the meaning of that private smile.

Answering either of our other two questons - What did N mean (to do) by it? and What did it mean to M? - entails considering public meaning. In order for N to mean anything by the smile, M's presence is required. If N is alone, unaware that we observe the smile, and we ask "What did you mean by that smile?", s/he is entitled to respond with, "What do you mean, 'What did I mean by that smile?' I didn't mean anything by it, I didn't know you were there." (Someone who 'smiled' in order, say, to find out which muscles moved, would not have truly smiled.) The consequences of a smile are its effects on other people, therefore N cannot have meant anything by the smile unless s/he knew that someone saw it.

Let us consider the straightforward case were N is smiling at M, in which I have argued the smile must mean at least that the good spirits expressed by it consist of friendly feelings directed at M. (Let's also at this point split the gender: N is

male, M female.) Friendship being what it is, a smile at another person is an assertion of and/or a request for reciprocity, an expression of community of interest itself: "We're on the same side, aren't we?" (The doubt or its absence which intonation would mark in that utterance can also be expressed in a smile.) In general, the meaning of the smile will vary with the nature of the community of interest between the parties. In particular, it will vary according to what their common interest is at the time of the smile. In the right circumstances, the meaning of a smile can be rich, specific, and intelligible.

Suppose N and M are standing beside a vase full of wild flowers, M looks at N, sees that he has been looking at them, and says "Do you remember?". In reply to M's question, N smiles. To M, N's smile means, roughly, "He remembers the wild flowers round Avebury, he remembers the good time we had together, he loves me." And if all goes well, then that is indeed what N's smile means, and also what N means by it: it is a perfect act of communication, in which the answer to each of our three questions is the same. Perhaps, however, unbeknownst to M, N is no longer in love with her, has been gazing unseeingly at the vase of flowers, and has no idea what M's talking about. But he's fond of her, doesn't want to hurt her feelings, and can tell that a smile is required. His smile means, roughly, "Yes, I remember that we had good times together, I like you", and that is also what he means (to convey) by it. (Poor M is deluded by hope, rather than by N.)

In the case just discussed, their communication goes wrong because the interest which N expressed is not the same as that

which moves M, though each of them assumes their interest is the same. It is also possible for N to mean by that smile precisely what M takes him to mean, and yet the smile not mean that at all. N may for reasons of his own, i.e. from some private interest, want M to believe he's still in love, that every memory burns as vividly in his breast as it does in hers, etc. In this case, the smile means not that N remembers, etc., but that N wants M to believe he remembers. Alternatively, perhaps M sees through his smile, recognises what it means, in which case, what it means and what it means to M will coincide rather than what N means by it and what it means to M coinciding. In these two cases, one would be inclined, I think, to say that N's smile was false, not a true smile. In order for his smile to be true, as soon as the smile ceases to be a straightforward expression, but is employed to achieve the effects of such an expression, its sincerity is in doubt. It is because a smile has a public meaning, signifying friendliness, etc., that N can use it as a tool and use it falsely.

Before moving on to pointing, let's summarize what we have found about smiling. The first point to note is that without a meaning, a 'smile' is not a smile. Its meaning depends on the interest expressed, and when it is a smile at another, also on the effects it has on their interests. Although the meaning of any given smile will vary with the current interest, the interest it expresses must be pleasing, either privately or mutually. Understanding the precise meaning of a smile is possible only for the insider, but its general, minimal, meaning is available to any witness. For someone to mean something by a smile, the smile

must be known to be witnessed. And the presence of a witness, combined with the general meaning of a smile, makes a false smile possible.

Though a private smile may be a derivative phenomenon, getting its meaning from that of a public smile, it is a normal and intelligible event. But a private act of pointing would be bizarre. A pointing finger means something like, Over there, in the direction in which my finger, arm, gaze, are pointing, is the object/place in which you are/I wish you to be interested. If I am on my own, and only my interest is in question, I don't need to show myself what it is that I'm gazing at, thinking of. An act of ostension differs from a smile also in that the one who points must always mean (to do) something by it. It follows from the meaning of the act of pointing that N means to single out some particular by extending a finger and arm towards it. If there is to N's knowledge no distinguishable particular to be indicated, then N cannot really be pointing (the direction in which a place is to be found would fall under the description 'distinguishable particular', being particular and distinct). It is possible to imagine the peculiar event of N, alone, pointing and meaning, say, that green cushion. But the point of pointing being to single out some particular - which N must already have singled out - it strikes one as unintelligible or even crazy to point without some M to witness the act and be informed by it.

So, in a normal, sane, case of pointing, there is an M for whom an object of interest is being singled out by N. If all goes well, M knows what pointing means, looks for the object of interest which N means, in the direction in which N points, and

succeeds in singling it out for herself. As with smiling, grasping the precise meaning of this occasion of pointing, entails understanding not only what its general meaning is but also what N means by it, i.e., what N is pointing at. Except in the most abnormal settings, in any direction that one looks there are a number of distinguishable particulars, and a pointing finger is entirely inadequate to single out one of them except by actually touching it, or getting very close. If N can talk, and M can too, then N will help her search for the right object by narrowing down the field of possible objects with a description. But pointing is used far more often by the inarticulate than the articulate, namely by the infant population. (We adults have learned the insufficiency of the gesture, and besides have learned a better way of arousing in our audience an interest like our own, i.e. by expressing ours in words.) If M is sensitive to N's interests, she'll make good guesses even without the help of words. Infants evidently believe that drawing attention to the fact that they have a particular interest (noises of an exclamatory and commanding kind generally assist in this), and indicating the direction in which its object is to be found, will be sufficient to single out for their audience the object they have in mind. This is a typical demonstration of infantile egocentricity. Luckily parents are generally sensitive to their infants' interests, because they get a lot of flak when they get it wrong, and - more seriously - because getting it right is vital for the acquisition of a vocabulary.

So, a successful act of pointing is one in which what it means to M is what N means by it, which in turn must fall within

the limits of potential meaning set by the general meaning of the act. The relations between those three aspects of meaning are broadly similar to those between the aspects of meaning of a smile. In either case, both what N can mean by it and what it can mean to M are crucially restricted by what its general meaning is. And in both cases, its meaning to M precisely what N means by it depends on M's being tuned into N's concerns. Also, in both cases, their having the appropriate general meaning is criterial for their being what they are, i.e. a smile, or a gesture of pointing; and their having such a meaning depends on N having - to however slight a degree - the interest which his action appears to express. There are however, fundamental differences between smiling and pointing, which derive from the difference in meaning. The meaning of an act of pointing is necessarily in the here-and-now of the world out there, while the meaning of a smile can lie in the history of, or in the future of, the interest which gives rise to it. So a smile can be highly charged with meaning - "His smile means everything to me" - in a way that is not available to an act of pointing. A further - related - difference is that an act of pointing must have a specific meaning to both N and M, not given in the general meaning. A smile could just mean "We're on the same side aren't we?" and even a private smile might just signify a general happiness (I think). But an act of pointing signifies that there is, here and now, some particular of interest and what particular it is will vary with the interest of the pointer without limit from occasion to occasion.

I began the discussion of pointing and smiling by claiming

that they were quasi-linguistic phenomena. What justified that description? They are like words in being used repeatedly, recognisably, and intelligibly in our discourse with others - they are part of a shared vocabulary. They have meanings that are not so much fixed as bound to fall within a certain exclusive range - they have a restricted 'meaning potential'. But smiling and pointing also amount either to making a claim or to asking a yes-no question and are in that way more sentence-like than word-like. And pointing may also function as an imperative - witness again the pointing infant.[4]

As claims both smiling and pointing can be denied or affirmed: N smiles at M, M does not smile back - no, we are not on the same side; N points, M shakes her head - no, it is not over there. As questions equally they can be denied or affirmed, indeed the only difference here between a claim and a question is the confidence with which it is put forward. Because they can be used to make claims, they can be used to make false claims, i.e. claims which are not merely mistaken as in the cases just given, but in which the smiler or pointer knows that they are not true. Here a difference between these cases and that of spoken language emerges: a false claim made by a smile or an ostension requires an interested witness, but I can make a false claim in language with only myself to witness it. Indeed, with language I can -if I perversely choose - make an eternal false claim, e.g. water boils at 17°C, while the truth or falsity of an act of smiling or pointing will always depend on the state of the interests of the people concerned at the time at which the act occurs. A further difference, which is presumably related, is that real sentences are composed of words.

Although I've been lumping smiling and pointing together in this summary of the ways in which they are language-like, they are also each like language in distinct ways too. Many of the subtleties which can be expressed vocally via intonation, voice-quality, etc., can also be expressed in a smile. What is more, utterances by which N means, and which mean to M, inter alia, "We're on the same side" abound in natural discourse, from "Hello" and "Goodbye" to "The way Thatcher sucks up to Reagan makes my flesh crawl" or "What a sunset!" or "Cream? Sugar?", or "Yes", etc. In all these cases - just as with smiling - their community of interest is part of what is expressed by the speaker and aroused in the hearer, and is strengthened thereby. On the other hand, a popular view of sentences is that they should be seen as primarily ways of stating facts, or giving true information, about the world. In my terms, when sentences do this, they contribute to the common interest not merely by arousing it or expressing it but by making a difference to its outcome, by changing it. So, if the capacity to inform were a criterion, then pointing would be much more sentence-like than smiling is. But a sentence informs more reliably because the object of immediate interest about which information is being imparted, will be specified under a name or description, not left as a gap to be filled by guess-work. And, of course, the object of interest represented in a sentence need not be present to the speakers.

* * *

I have been showing, among other things, how the meanings of smiling and pointing depend on the interests of the actor and his audience. Without the requisite interest, I suggested, they are either meaningless movements or fakes. So, to some extent even their general meanings within which more specific meanings may be expressed and understood depend on the current interests of the one who performs them. On the other hand, it is because we understand them as expressing interests of the requisite sort that they can be detached from those interests and used falsely.

When M smiles at N in the context of the vase of wild flowers, she may do it with romantic confidence in the perfection of their community of interest, sure that to N her smile will mean (roughly formulated) "Ah those wild flowers round Avebury, oh those heady days, etc., etc." - but she may be wrong. N may neither cotton on to the connection between these wild flowers and those, this moment of affection and those moments, nor even cotton on to the fact that there is any connection to make. However much M smiles, there's nothing in the smile to evoke in N the specific interest which gave rise to her smile. Their community of interest itself she can evoke, because that's what a smile means, but for her to arouse between them a common interest more specific she must give some more spec ific outward sign. Although M's smile may mean "Ah those wild flowers, etc." to both her and N, she would have to be romantically confident to the point of lunacy to employ her smile to convey that meaning. All she can mean by it is everything which follows from their being "on the same side". Nothing beats a smile for intimacy but for specificity language beats it hollow.

When we look at pointing, we see that 'specificity' is not itself a precise term: pointing is both less and more 'specific' than talking. Either by pointing or by talking, one can single out a spatio-temporal particular visible to those present, but pointing - at least at close quarters - is more reliable as well as more succinct. To see that this is so, picture a floor covered with spilt drawing pins, each identical with every other. Now imagine picking out one of them by description. If on the other hand the drawing pins were distinguishable from one another, so that they could be told apart, then each drawing pin could have a name. So that when I said, "Ned", you would know that I meant the drawing pin called Ned, so long as you knew that there was such a drawing pin, and you knew that I was talking about drawing pins rather than, say, boys, or donkeys. And, if you further knew which drawing pin was called Ned, were practised in telling them all apart and could recognise Ned, then you would even know as much as my proximous pointing would have told you. Namely, precisely what is the object of my interest.

Pointing has the advantage over naming in that it requires no special prior knowledge from its witness: successful naming always requires such a background. Naming, of course, has it over pointing in that it does not require co-presence with the thing named in order for uptake to be secured - given all the prior knowledge required. Both pointing and naming have the advantage over description in that - when the relevant conditions obtain - they specify uniquely, and succinctly. A description with no deictic element will have to include references to location in space and time if it is to be specific enough to pick

out just one individual. Otherwise there can be no guarantee that some other individual won't fit the description given. Fortunately unique specification is hardly ever required; knowing what the speaker means involves only knowing what the description means. Even quite a sketchy description will be good enough to establish between speaker and hearer a common object of interest by delimiting the range within which such an object must fall. Two interests are the same just so long as the same differences have shaped them. We point, we name, we describe, in order to arouse in our audience a common interest, or in order to demonstrate our common interest by expressing it. Describing has the advantage over naming in that it can be used to establish or create between the speakers a new specific object of interest. And, unlike pointing, it can do that in the absence of the object in question. Furthermore, both naming and describing have it over pointing in being able to single out not only what is not present but what is not material (ideas, possibilities, the Holy Spirit, etc.) and even what does not exist at all (Vulcan, unicorns, the present King of France, etc.).

* * *

I began this section by remarking that 'specific' was not a precise term, as demonstrated by the fact that pointing was both more and less specific than talking. We have seen that for picking out a spatio-temporal material particular unassisted by background knowledge, pointing is superior to talking. But though pointing, like naming and describing, can specify an

object of interest, none of them can specify what is currently interesting about that object. It might seem that once one has grasped what object it is that's being pointed at, (named/described) then the complete meaning of the act has been grasped, i.e., what it means, what A means by it, and what it means to B. But that is certainly not always the case. Suppose B has asked A "Which of these would you like?", then when A singles one out, it means that is the one A wants - and that's what A means by it, and that's what it means to B. Mutatis mutandis, the same will be true whenever something is singled out in answer to a question.

By far the most frequent occasion for the unadorned production of a name, description, or act of pointing, is in response to a question. But interest of other sorts will account for one's understanding o other occasions. For example, if A and B are waiting for the morning post, and A addresses B with "the post", B will take it that she means that the post has arrived. Or if B has lost his car keys and they've both been looking for them, then A's pointing them out will mean B's car keys have been found, and that's what A will mean by it, and that's what it'll mean to B. In each of those cases, the act is understood because it answers to (satisfies) a previous interest which the actor assumes the audience shares.

However, one of the expressions in question may also stand alone as an exclamation, as in "A rainbow! a brilliant rainbow!" or "the police!". The presumption in such a case being that the object put forward as interesting, remarkable, is unmistakably interesting in itself; nothing further need be said to specify

some currently interesting feature of it. Compare the meaning of a rainbow in "A rainbow is a phenomenon characteristic of mixed weather conditions" and in "A rainbow!". In the exclamation it means that there is a rainbow here and now before our very eyes, and it means that the rainbow is self-evidently interesting, a wonder; and it means that B should look at it! In subject position in the sentence it means none of those things. Pointing, using neither face nor voice, and having none of talking's potential for clamour, is not on its own capable of expressing/arousing wonder. But, as anyone who has been acquainted with a pre-articulate child will know, accompanied by grunts and cries it can easily mean "over there is a wonderful thing! Look! Look!", though we are less inclined as we grow older to see the material of daily life as wonderful.

In the right circumstances, any of these methods of singling out an object of interest can also be used to express a question or an imperative. For example, if some informant is supplying me with the names of a list or series of individuals, then when I say "the man with a black moustache" it will mean What is the name of the man with the black moustache? Similarly, when an infant points and exclaims non-verbally, it may be taken to be requesting the name of the indicated object - certainly giving the name is a normal adult response. Or it may be demanding the actual object, and being fobbed off with its name. With the aid of intonation, names and descriptions can be used as yes/no questions, though the meaning of a question so expressed will in turn depend on the what question this may or may not be the answer. Thus, "Bacon and eggs?" is a possible answer to the

question, What would you like to eat? when what the addressee would like to eat is known to be in question. Equally, given the right social circumstances, "bacon and eggs!" may function as an imperative.

I have been suggesting that when an expression of this kind appears in isolation, it is fully understood only when what is currently supposed to be interesting about the object singled out has been understood. Only when we know that do we know what it means, what M means by it, and what it means to N (if N has understood it). Otherwise, such an expression is a puzzle, raising the question of why one's attention has been drawn to it in the first place. Given a grasp of the relevant interest, these expressions can function as assertions - in answer to questions - or as questions, requests, imperatives and exclamations. Of these, only exclamations can be fully intelligible without some specific prior current interest. In the other cases, understanding what a name or description on its own means will involve understanding its relation with a current interest or interests.

Exclamations, assertions, questions, requests and injunctions are standard types of speech act, and the meanings of sentences which exemplify each type will vary in the same way as in the examples above.[8] Let us, as briefly as possible, see what distinguishes each type. Exclamations have to do with the expression and arousal of interests, i.e. with changes in their state rather than in their form. They signal that the speaker regards the object of interest as in itself worthy of attention, rather than meriting attention just because it's the location of

some modification to a world-view. Assertions, I suggest, specify a difference that is to be made or has been made to the current form of some specific interest. Insofar as the effect of such a difference is on future conduct (including talking and understanding), it is important that any sentence which proposes such a change be true. Yes/no questions are requests to the audience to confirm or deny a truth, requests to put one in a position of certainty. Certainty is the position one is in when some change has sunk out of cognitive or public space and become mere form, become taken for granted. Dubious, debatable, assertions and Yes/no questions both have the effect of deferring the acceptance of change, of keeping alive an interest in their truth.

A WH question is a request that one's hearer specify either some change in the interest system, as in Q: 'Who is that man?', A: 'He's the new foreman.'; or the location of such a change, as in Q: 'Who took the last piece of salami?', A: 'John did.'. Any sentence with a non-specific term like someone, somehow, shome mishtake, somewhere, etc., can function as a WH question, if it prompts a knowing hearer to specify the person, manner, etc.. Other requests are requests to the hearer to affect the outcome of the speaker's interests by affecting the world itself - imperatives being the less polite way of enjoining a person to do that. Evidently, a major advantage of speech over pointing is that it can come in sentences and not just noun phrases, that it incorporates verb phrases, and therefore can create or evoke a specific modification of an interest. That is how talking can be more specific than pointing: by talking one can specify not only

interest, so, when an encounter with a particular phenomenon arouses a specific interest, the relevant name springs to mind with it. As we saw in Chapter 5, above, something which acts as a trigger to some interest may also function as an expression of that interest: words are a special case of this.

* * *

In one way or another, we have found in this chapter that every aspect of meaning was related to some characteristic of interest. 'Interest' and 'intention' are clearly related terms. Have I really, then, just been reworking the familiar connection between meaning and intention (see, e.g., Grice, 1957)? I have generally avoided using the term 'intention' throughout this thesis, largely because it is notoriously problematic (see, e.g., Grice, loc. cit., and Anscombe, 1957) and I've never been at all sure what it meant. However, I think it may be justifiable to define an intention as an interest which is sufficiently aroused to be in competition for an 'output channel'. To make it a well-defined class, it would be necessary to postulate another threshold in the interest system: that above which an interest competes with other interests above that threshold. The existence of such a threshold, and the need to postulate it, are debatable, especially in the context of a discussion of meaning. Intentions are not generally seen as repositories of information, nor viewed as interconnected in such a way as to form a network; nor does one naturally speak of their being aroused, though one does of course speak of their being expressed; lastly, 'interest'

but not 'intention' suggests that what we are talking about when we talk about meaning is what matters. In short, intention-talk obscures connections between different aspects of meaning which interest-talk brings out.

Meaning is, in general, the capacity to arouse, express, or affect the outcome of some interest or interests, i.e. it is an aspect of relevance.

Chapter Eight: Meaning II

A central theme of this thesis is that a wide variety of questions about language and the ways we use it can be answered only by bringing in the notion of interests. It was with delight, therefore, that I discovered Hilary Putnam arguing, in his '78 and '83 collections, that explanation, cause, reference, truth, and translation, are all interest-relative. And in 'The Meaning of "Meaning"' (1975) - with which this chapter is primarily concerned - the relation same L on which much of his argument turns is, as he points out, also interest-relative.[1] However, as I shall seek to persuade you, it is the interest-relativity of this relation, same L, on which his argument founders. I think that if Putnam had taken his own observation seriously he'd have reached very different conclusions. For from it follow the interest-relativity of categorization and of extension itself. I shall first sketch the bones of his argument, then home in on what I see as the weak spots.

A crucial claim he is making is that what he refers to as 'traditional theory of meaning' - by which he means 20th century Anglo-American theory of meaning - rests "on two assumptions that are not jointly satisfied by any notion, let alone a notion of meaning" (1975: 1219). These are "(I) that knowing the meaning of a word is just a matter of being in a certain psychological state, and (II) that the meaning of a term (in the sense of 'intension') determines its extension.". The central thrust of

his paper is that at least one of these assumptions is going to have to go. In the event, he keeps II - that sameness of intension entails sameness of extension, but rejects I - that knowing the meaning of a word is a matter of being in a certain psychological state. An important proviso, however, is that 'psychological state' appears in assumption one, only in the severely restricted sense which conforms with 'methodological solipsism'. That is, with the assumption that "no psychological state, properly so called, presupposes the existence of any individual other than the subject to whom that state is ascribed" (220).[2] So, although Putnam gets carried away occasionally and says things like "Cut the pie any way you like, 'meanings' just ain't in the head" (loc. cit.) his conclusions aren't really that radical.

It follows, Putnam says, from the two assumptions jointly, that "two speakers cannot be in the same psychological state in all respects and understand some term A differently; the psychological state of the speaker determines the intension (and hence, by assumption II the extension) of A." A good deal of 'The meaning of "meaning"' is devoted to attacking that conclusion. His first move is to take a handful of natural kind terms, and argue that their extensions are what they are quite independently of what individual users may know about them. His next point is that there are certain terms which the 'average' speaker uses perfectly competently, but leaves it to 'experts' to know the criteria which fix the extension. "It is only the sociolinguistic state of the collective linguistic body to which the speaker belongs that fixes the extension." (229)

Next, along lines parallel to Kripke's,[3] he argues that natural kind words are 'indexical'. "Water at another time or in another place or even in another possible world has to bear the relation $same_L$ to our 'water' in order to be water" (234). The kind of criterion which an expert will apply, Putnam suggests, will tend to concern 'hidden structures', so that our water being H_2O, any water lookalike, feelalike, smellalike, tastealike which is not H_2O will not actually be water. - Although it does not follow from this that we should regard 'water' as synonymous with 'H_2O'. His final point, before moving to the topic of 'meaning' as such, is that many non-natural-kind words, "most nouns" and "verbs like 'grow', adjectives like 'red', etc., all have indexical features". As with natural kind words, no conjunction or cluster of properties will be synonymous with, for example, grow or red, or with any other 'indexical' terms.

Under the heading 'Meaning', Putnam sums up his conclusions so far, thus: "the extension of a term is not fixed by a concept that the individual speaker has in his head, and this is true both because extension is, in general, determined socially ... and because extension is, in part, determined indexically ... Traditional semantic theory leaves out only two contributions to the determination of extension - the contribution of society and the contribution of the real world!" (241). He suggests that one should "identify 'meaning' with an ordered pair (or possibly an ordered n-tuple) of entities, one of which is the extension" (246) the other being the stereotype. In order, he says, for someone to count as having 'acquired' a word they are required to know a certain amount about stereotypical instances of it. That

such a "conventional idea is associated with 'tiger', with 'gold',, etc." is, Putnam urges, "the sole element of truth in the 'concept' theory" (250). Here, as elsewhere, he identifies concept with intension.

Although I entirely agree with the thrust of Putnam's attack on 'methodological solipsism', I don't think he goes far enough. If he took social practice as seriously as he purports to do, he would have to abandon his scientistic account of 'extension' - as exemplified in the claim that H_2O is <u>the</u> extension of water. The most obvious reason for finding that account objectionable is its extremely narrow application: 'Natural kind' terms do not play a conspicuous role in most people's vocabularies. What is one to say about the extensions of those thousands of other meaningful words? One can subscribe to Putnam's views about the indexicality of a large part of our vocabularies (and I do) without being committed to the view of extension he presents in 'The Meaning of "meaning"'. I also think it would have been more helpful to reconsider the notion of 'psychological state' from a non-solipsistic perspective than to dismiss it as irrelevant - as Putnam sometimes seems to do - or to reduce it to the 'having' of a socially determined stereotype.

It will be instructive to take a closer look at his arguments about natural kind terms. Several years later defending himself against a possible charge of materialist 'essentialism', Putnam sums up his case in 'The Meaning of "meaning"' as follows, "What I have said is that it has long been our <u>intention</u> that a liquid should <u>count</u> as 'water' only if it has the same composition as the paradigm examples of water ... I

claim that this was our intention even before we knew the ultimate composition of water" (1983: 221). Put like this, it seems to me, the weakness in his argument appears plainly: who are "we", who apparently had these intentions for so long? This is not a trivial question.

To support his claims about natural kind terms, Putnam invites us to imagine a variety of scenarios, many of which involve a 'twin world'. That is, a world which is exactly like ours in almost every respect, except - for example - that what English speakers there call 'water' is in fact 'XYZ' rather than H_2O. If that were so, he says, even if the 'water' played exactly the same role in their lives as it does in our lives, then "in the sense in which it is used on Earth ... what the Twin Earthians call 'water' simply isn't water." (224) At first glance this may seem irrefutable, their 'water' and ours just aren't the 'same thing'. But it's what they thirst after in the desert, it's what they freeze and put in drinks, it's what they swim in, it's what their cucumbers are 98% composed of, it's what they boil in their kettles, it's what falls from the clouds as rain... In short, except for its 'hidden structure', it's indistinguishable from what we call water. So what? Putnam the scientistic essentialist would reply, what they mean by 'water' just is not the same thing as what we mean by 'water', that is, it does not bear the relation same$_L$ to our stuff.

He claims that, given that our water is in fact H_2O, if visitors from Earth had gone to Twin Earth in 1750, even though the 'hidden structure' of water had not at that time been discovered on either planet, then although they'd have agreed

with Twin Earthians in calling that wet stuff 'water', they'd have been wrong. "In 1750 we would have mistakenly thought that XYZ bore the relation same L to the liquid in Lake Michigan, while in 1800 or 1850 we would have known that it did not". (225) So, when Twin Earthians talk about 'water', they don't mean what we do, and they never did.

In fact, Putnam himself provides the material for what seems to me a conclusive counterargument: firstly by pointing out that the relation same $_L$ depends on what is <u>important</u>, and hence is interest-relative; secondly by informing us of the case of jade. "Although the Chinese" (-which Chinese? not the scientists I bet), "Although the Chinese do not recognise a difference, the term 'jade' applies to two minerals: jadeite and nephrite. Chemically there is a marked difference." (241) He uses the case to illustrate the primacy of the <u>local</u>, of what we're familiar with, in the indexicality of meaning. But let us retell the story of a visit to Twin Earth, using 'jade' rather than 'water'.

It is 1800. Fifty years ago a spaceship from Earth - the only one ever - visited Twin Earth. The astronauts on it found life there just like it is here except that jade was more plentiful and much cheaper. Knowing a bargain when they saw one, they loaded up with jade on the return journey, which they sold without difficulty to jade craftsmen the length and breadth of China. The present Emperor of China is a renowned jade connoisseur - an expert par excellence. One of the new-fangled 'hidden structure' johnnies, on examining some of the Emperor's jade collection, announces that some pieces are made of calcium magnesium and iron - nephrite - while others are made of sodium

and aluminium - jadeite. The Emperor expresses polite interest. The scientist then shows him which pieces have which structures; he has segregated them as he identified their category, or he could not now tell them apart. Now the Emperor's interest increases, for he can distinguish them: the nephrite figures are all carved from jade brought back from Twin Earth in the middle of the last century, while none of the jadeite was made from the stone in that consignment.

The Emperor thinks of this as a nice tid-bit to pass on to other connoisseurs, whom he is not above upstaging. The scientist immediately conceives an ambition to visit Twin Earth and find out if the hidden structures of all stones, liquids, etc. there are different from those on Earth; also a more practical ambition to check further available samples. Both, in short, conclude that it's highly likely that Twin Earthian jade and Earthian jade have different chemical compositions. Neither is in the least inclined to deny that the pieces in the Emperor's collection with a Twin Earth origin are jade. Nor, once the word had spread to other interested - expert - parties, is any one else inclined to say that Twin Earth jade is not jade. Rather, the jade connoisseurs speak (and think) of it as Twin Earth jade (it has become more highly valued than Earthian jade, being in limited supply). The scientists, qua scientists, don't speak (or think) of it as 'jade' at all, but as 'nephrite', a naturally occuring chemical compound (from Twin Earth) with a contingent similarity to another naturally occurring compound, known as 'jadeite'.

In fact, sans Twin Earth and (presumably) sans Emperor,

something like this story must once have occurred. For once upon a time nobody knew that jade came with two hidden structures, and one day somebody must have found it out. And - as Putnam helpfully points out - the end result was as sketched above, jade just is jade, whatever its chemical composition may be. The first moral I want to draw from this is that Putnam's H_2O/XYZ case won't wash. The reason it won't is that he has not been sufficiently alert to the consequences of his own insight, viz that the relation same $_L$ is interest-relative. From that relation being interest-relative, it follows that indexicality must be, so must the basis of categorization, and so, therefore, the 'determination of extension' must also be interest-relative. Throughout any language community, there are smaller communities of interest whose members relative to the community at large are experts: expertise is also interest-relative. Although Putnam correctly draws our attention to the role of experts, he then proceeds to argue as though only scientific expertise really counts. It will be interesting to consider briefly why Science exerts this fascination.

In the case of jade, distinct chemical structures "produce the same unique textural qualities" (241); it is this which makes the case of jade anomalous. It is normal for differences in chemical structure to be correlated with gross perceptual differences, and that is not accidental. Once the molecular structure of a substance has been discovered, it becomes possible to make predictions about its behaviour in a wide range of circumstances. 'Hidden structures' are important just to the extent that they are consequential. Knowing that, e.g. water is

composed of molecules of two hydrogen atoms to one of oxygen, is a bit like knowing the divine recipe. But 'H_2O' was supposed to have entered the picture as somehow uniquely specifying the extension of water, not as an insight into the 'psychological state' of the divine mind. Why should 'H_2O' be a specification of the extension of water, any more than 'flour, yeast, sugar, salt, water' should count as the extension of bread? Although telling someone that that was what bread is typically made of, might be part of telling them what sort of thing "This is bread" is true of. And equally, its being H_2O does have something to do with the sort of thing "This is water" is true of, and that's what the notion of 'extension' was introduced as - namely as that of which a term is true.

It is presumably thanks to the fact that it is composed of a whole lot of pairs of hydrogen atoms combined with single oxygen atoms, that water is what people the world over pour on their plants to revive them, drink to quench their own thirst, witness flowing in streams and rivers, falling from the clouds, etc., etc. Thanks to its being the sort of thing it is, it plays the same role in many lives; that is, it makes the same differences to the outcomes of (approximately) the same interests, all over the world. When we talk about 'the extension' of a term, we are asking what it is that people thus treat as distinct and call by such-and-such a name. It is a question about social practice and its motivation.[4] Questions about 'indexicality' and social practice seem to be inextricable when it comes to seeing what determines 'the extension' of a term.

Discussing what he sees as the consequences of his H_2O/XYZ,

Earth/Twin Earth argument, Putnam says, "There are two plausible routes that we might take." The first of these, he suggests, is "to retain the identification of meaning with concept and pay the price of giving up the idea that meaning determines extension," in which case "'water' has the same meaning on Earth and on Twin Earth, but a different extension." (245) The other alternative, he claims, is to drop the identification of meaning with concept, but agree that meaning should be identified as "an ordered pair (or possibly an ordered n-tuple) of entities, one of which is the extension". (246) In which case it would be "trivially true that meaning determines extension (i.e. difference in extension is ipso facto difference in meaning) but totally abandons the idea that if there is a difference in the meaning my [Twin Earthian] Doppelgänger and I assign to a word then there must be some difference in our concepts (or in our psychological state)." (246) Let us take a look at the one solid argument he advances for preferring the second of those options.

Here is his whole brief argument: He invites us to "consider 'elm' and 'beech', for example". He has already confessed that he personally cannot tell them apart - "If" he says, "these are 'switched' on Twin Earth, then surely we would not say that 'elm' has the same meaning on Earth and Twin Earth, even if my Doppelgänger's stereotype of a beech (or 'elm' as he calls it) is identical with my stereotype of an elm." (246) Does that legitimise Putnam's conclusion? Is it evidence that we should "totally abandon the idea" that difference in meaning entails difference in 'concept' or 'intension'? It is only evidence for such a conclusion if we subscribe to the solipsistic account of

intension that Putnam is himself attacking. To see that this is so, consider what it means for 'elm' and 'beech' to be switched. It means that English-speaking people on Twin Earth - my Doppelganger's people - call an 'elm' what English-speaking people on Earth - my people - call a 'beech'.

Putnam is perfectly right to insist that what we say about the meaning of any given word cannot be based on what's in any individual's head in isolation. But it does not follow that differences in meaning need not affect the psychological state, the sets of intensions, of individuals. Like Putnam, I am a ltlle vague about just what the difference is between elm and beech, but I know at least this much: they are deciduous trees (Putnam also admits to knowing that much) which there are grounds for treating as distinct, one sort being called "elm", and the other sort "beech". Although, like Putnam and his Doppelganger, I don't happen to know the precise differences on which the distinction is based, I do know that the words mean whatever the people who do know mean by them - this is something I know, about the meanings of words. In other words, it is a part of my 'psychological state'. The idea we should abandon is that of solipsistic intensions or concepts. That way we can save the idea of intension and extension as a reciprocal pair, and not abandon the thought that mind and meaning have some connection.

Let's return now to the question of 'extension'. Putnam wants us to include 'the extension' in any account we give of the meaning of a content word. Conceding that one should, the question becomes How? We have seen that an essentialist account won't do - "H_2O" turned out to be just another way of talking

about water, another description. And the position is worse than that: for in a sense, Putnam is making an impossible demand when he requires us to incorporate 'the extension' in the 'meaning vector' (or whatever). The problem is that 'the extension' of "water" is of course water itself - and we can't actually put that into our description.

So, is there any way of 'incorporating' the extension of a term into an account of its meaning without simply substituting a description, and hence retreating again into intensionality? It seems to me that there are two possible ways of doing so. One would be the Davidson-style 'disquotational' method, in which the word "water" would get into the account once with and once without quotes, as when I wrote above: "the extension of "water" is of course water itself."[5] But, as Putnam points out, this is a procedure which in no way illuminates questions about meaning. The method I prefer, because it tells one a <u>little</u> more, - and of the right kind - is the following. The extension of any term, e.g. "water", just is <u>whatever it is that we people call "water"</u>. The extension is thus given, undescribed, by "whatever it is": <u>qua</u> extension there is no more to be said about it. This account at least has the virtue of bringing in both the indexicality and the other-people dependancy of extension. However, when it comes to giving the <u>meaning</u> of a term, it obviously leaves almost everything still to be said.

What does one need to say, then, in order to 'give the meaning' of a term? The problem is, what to <u>tell</u> someone who does not already know, e.g. what it is that we people call 'water', in order to teach such a person what the word means.

Doing that, Putnam suggests, is at least teaching them about its 'stereotypical' characteristics. For example, the list he offers for water is, "colorless; transparent; tasteless; thirst-quenching; etc." (269) That 'etcetera', that open-endedness, is largely his response to the indexicality of so much of natural language: no matter how long one goes on, one can never succeed in capturing the 'thisness' of an indexical term. So it is only rarely possible to give another person necessary and sufficient conditions for membership in the extension of a term. And that, Putnam rightly asserts, is not what one should be trying to do, anyway. Rather, one should be attempting to impart "Information about the minimum skills required for entry into the linguistic community" (257); that is, one should be teaching learners their 'linguistic obligations'. "What it means to say that being striped is part of the (linguistic) stereotype of 'tiger'," Putnam says, "is that it is obligatory to acquire the information that stereotypical tigers are striped if one acquires 'tiger'". (251) The psychological counterpart of this social obligation, is, it seems to me, that once one has acquired a word one is from then on obliged to hear it as meaningful. That is, one's tiger interest - faint though it may be - automatically responds, is activated or aroused, by its name, by the word "tiger".

Although none of what gets listed as stereotypical is a necessary condition, Putnam suggests - borrowing Quine's notion - that some features are relatively central or unrevisable.[6] If we are talking about tigers, for example, we can imagine all sorts of variations which would still count as tigers, "But", as Putnam

puts it, "tigers ceasing to be animals! Great difficulty..." Here is his sample list of highly central features for a variety of terms: animal, living thing, artifact, day of the week, period of time. These, he remarks, "attach with enormous centrality to the words 'tiger' - animal, 'clam' - living thing, 'chair' - artifact, 'Tuesday' - day of the week, 'hour' - period of time. "And," he adds, "they also form part of a widely used and important system of classification." What makes these particular features part of a system of classification, is, of course, their interest-relativity. They play a highly consequential role in our lives, that is, they make a difference to the outcome of a wide variety of interests. In making the same differences to a number of different interests, a feature renders those interests 'the same' in respect of itself: two interests are the same just to the extent that the same differences have shaped them. These effects on our interests, not only affect the way we deal with 'the things themselves' when encountered, but also affect our conversational practice.

In most contexts in which I meet or talk about things which I'm inclined to call "tigers", and in most contexts in which I hear about things that other people are calling "tigers", then their being animals will be a characteristic of what their being tigers means to me. But now let's suppose there's a football team called "The Tigers", and I'm sitting in the pub with a bunch of football fans... In that context, it's obvious that I will not take my friends to mean one of the animals known as "tigers" when they speak of "a Tiger". Someone might want to say that this "Tiger" and the other "tiger" are really quite different

meanings have nothing apparent in common except the form of their expression, "bank" being a classic case of this. As Schegloff points out (1972), ambiguity is hardly ever a participant's problem. It is, rather, an observer's problem; and the reason that is so is that the observer is, by definition, qua observer, not a member of the community of interest, and not therefore automatically privy to the current common concerns. The fact that children hear puns, notice ambiguities, much more readily than we adults, is I think largely because they haven't learned to tune in to the common interest, and perhaps partly because all their interests are at such a relatively lively level compared to ours. In most cases in which a word which is technically lexically ambiguous appears, its other meaning is relevant to an interest which is currently at a low ebb, quiet, and only the meaning relevant to a current interest gets effectively 'activated'. I don't rule out the possiblity that the word form itself is sufficiently potent to call even an extremely weak interest, but since such a call would normally be overruled in cases of ambiguity, there is no way I can imagine in which this question could be answered. What we are finding here is a special case of a phenomenon I've already discussed extensively, namely that the more highly aroused an interest already is, the more likely it is to be aroused still further, the more avidly it seizes on the relevant.

Let us now go back briefly to the question of the relation between extension, intension, and meaning. The first thing I want to say is that if 'extension' is to play any role in a theory of meaning, then its spatial connotations should surely be

regarded as irrelevant. Take, for example, the meaning of 'meaning'. I have been trying in this and the last chapter to formulate as best I can, what it is that distinguishes whatever it is that we people call meaning, that is trying to formulate what distinguishes the extension of meaning. Yet there is nothing out there I can point at. What I am trying to do is articulate things that I assume we actually all know we all know - or, better, can all assume we all assume. We are surely, as ordinary people, experts on what meaning means. It's only as linguists, lexicographers, philosophers, that doubts arise.

Just as I cannot give the actual extension of any term, I cannot give its intension, either; for its intension is simply the current form of the interest which bears its name. The best I can do is try to convey what it is that makes that interest what it is, that distinguishes it from all those other interests. In naming those features I thus confer them with an extension, make objects of them. But their psychological status is not as objects, but rather as boundaries, which give form to the interests they bound. The more impersonal an interest is, the less freedom there is as to its form, the more rigid its boundaries, the firmer its structure. Word meanings are like islands of near fixity, with reliable and well-motivated interconnections, in an otherwise highly variable, subjective and idiosyncratic network of hopes, desires, etcetera.[8] Putnam occasionally speaks as though being a social practice precluded a thing's being a psychological phenomenon. In fact, if a social practice fails to affect individual members of a society, it ceases to be a social practice. Anyway, as soon as we permitted

interests to enter the picture, we let intensions in at the back door. For interests, e.g. hopes, fears, desires, doubts, concerns, are paradigmatically intensional.

* * *

I began the last chapter with a brief discussion of 'person-neutral linguistic meaning'. We now see that this 'neutrality' applies only within a language community (which itself may have very few members).[9] When we say X means Y, where X is a word, it is equivalent to saying: within such and such a language community, if someone is talking about X (say 'Tigers'), then you may be sure they are talking about Y (say 'footballers'). So that when, in this community, an interest is expressed in X then it is bound to be an interest in Y, or when new information about X is acquired, then it is bound to be new information about Y, and so on. This has the effect that, within any given language community, where X is a word, what N means by X, what X means to M, and what X means will be substantially identical, although X may mean more to M or N than it does in isolation. But the situation is rather different where X is a sentence.

As we saw in the last chapter, sentence meanings characteristically concern modifications that have been or are to be made in the world and/or in a person's view of it, i.e., in their network of interests. Members of the same language community can be counted on to share a substantially identical subsystem of their interest networks i.e. the semantic network

created by word meanings and their interconnections, and to share the rules for their combination. Therefore, as is the case with words, we may say that within a given language community, where X is a sentence, X has a meaning which is not dependent on what any individual knows. This guarantees that what X means to M and what N means by it overlap. However, although their networks of interest coincide in the strictly semantic, M's and N's may diverge widely elsewhere. Since when X is a sentence its meaning crucially consists in its consequences in an interest system, its meaning for M and its meaning for N will diverge just insofar as it has different repercussions in M's and N's interest systems.

Chapter Nine: Arousal

Set in stark isolation, the word 'arousal' suggests brutish appetite rather than the refined apparatus of cognition. Yet "it aroused my interest" is common parlance, hence I chose the word to use in my definition of relevance. Other roughly equivalent expressions are, "it stirred", "excited", "stimulated", "tickled", "seized" my interest. In any of these cases, the consequence of my interest being aroused (etc.) by X, is that my attention is directed towards X, X becomes the object of my current interest. Put like this, it sounds as though the interest is essentially a passive phenomenon, a response to the stimulus of X: because X is what it is, my current interest (I) is what it is. But, except in the event of X being completely unfamiliar, the character of my current interest will depend on two further factors, both of which concern the past history of the interest which X, at such and such a moment, has aroused. One is the history of changes which I has undergone prior to whatever effect X may have on it. That history is responsible for I's current form, and partially determines its future development. The other factor is I's level of arousal immediately before my encounter with X.

In general, the level of arousal of any interest must depend on how wide awake its owner is. For, except during periods of dreaming, what distinguishes sleep is that one's interests are quiet, not busy producing the words and deeds, and 'thoughts' and images, that fill the day. And a sleeping person is the opposite

of laboratory rats have been made to sing for their supper. But even in the most strictly controlled conditions, one rat cannot be counted on to behave like another, nor even to reproduce the same responses itself in the same objective conditions.[2]

The inconvenient unpredictability of rats is far more extreme in human beings. So too is the related problem that, if all varieties of behaviour are to be explained in terms of the variety of drives, then one is liable to wind up with uncountably many 'drives'. As one 1950s psychologist - an observer of chimpanzees - put it, "Every little action has a motive all its own" (Nissen, 1953). 'Drive' was introduced as part of a positivist, reductionist, programme, aimed at demonstrating the possibility of psychology being a True Science. As such, drive was supposed to be quantifiable, and it was supposed to assume only a restricted number of forms, each of which could be derived from a biological need. This is one of many problems on which behaviourist theory eventually foundered.[3] But, despite the demise of the theory, the facts which 'drive' was proposed to account for do not go away.

Drive was thought of as an energising factor, an increase of drive leading to an increase in activity, and conversely. A hungry rat (or person) will tend to look lively when presented with food 'cues' or 'stimuli' which a satiated rat (or person) may ignore completely. Up to a certain point, the hungrier it is, the faster its reaction times to relevant stimuli are liable to be, and the less likely it is to be affected by irrelevant stimuli. Sex, thanks to Freud, and hunger and 'pain avoidance', thanks to laboratory techniques, have been the paradigms of

'drive' for psychologists. With those as models, it was natural to treat the phenomenon in question as essentially biological on the one hand, and 'goal-oriented', on the other. But alertness to the relevant, and the production of relevant activities, were of course what I introduced the notion of interest to account for. As we have repeatedly seen, not all interests are purposeful; and there are indeed for any individual human, uncountably many, only some of which can sensibly be attributed to any biological need.

Plans and purposes, and even hopes and desires can be manipulated to some extent in a laboratory, and they have tended to take over the explanatory role of 'drive' in more recent psychology.[4] Wonder, awe, and love are less susceptible to engineering, but each is as capable as any other kind of interest of occupying the attention wholly, to the exclusion of everything else. In such cases, the relevant activities produced are gazing at, contemplating, 'drinking in', and exclaiming at the object of interest - and, in the case of (requited) love, a reciprocation of these activities is called for. 'Alertness to the relevant' is manifest as 'heightened awareness' in each case. These states of mind and activities which are beyond the reach of experimental psychology should nonetheless be recognised by theoretical psychology. Despite this blind spot, 'goal-oriented' behaviour has proved a fruitful field for research into human, as well as rat, psychology. One of its fruits is the prominence of 'schemas' as constructs for the description of particular, familiar, behaviours.[5]

The other direction taken by the notion of 'drive' was

towards its identificaion with the 'arousal system' of the brain stem. (See, e.g., Hebb, 1953.) Luria quotes Pavlov saying, "Organized, goal-directed activity requires maintenance of an optimal level of cortical tone" (Luria, 1973). Sleeping lowers 'cortical tone' to the point at which inertia pervades the system. Research by Moruzzi and Magoun, published in 1949, demonstrated that cortical tone was regulated by the reticular formation of the brain stem. "This formation has the structure of a nerve net, among which are scattered the bodies of nerve cells connected with each other by short processes. Excitation spreads over the net of this nervous structure... not as single, isolated impulses, and not in accordance with the 'all or nothing' law, but gradually, changing its level little by little, and thus modulating the whole of the nervous system." (Luria, 1973)

The notion of an excitable network of neurons at variable levels of excitation, through which activation from various sources spreads, has fired imaginations and caused waves of excitation to pass through a worldwide network of psychologists. It also has some clear parallels with the picture that has emerged from the discussion in this thesis. The general correlation between this 'arousal system', active in the waking, functioning, person, and my busy interests thrumming away all day long, is so clear as to be undeniable. But to what extent more particular correlations can be established is much more problematic.

If I have the thought, say, "It's time to feed the cat", on different days, do the same neurons fire each time? Or if one

day I have that thought in so many words, and another I think, "It's the cat's supper time", do some of the same neurons fire? Or if one day I have that thought, whilst another I just give the cat his supper, because it's time, will some of the same neuronal activity occur? And what if someone else has that same thought about that same cat...?[6]

These questions merely scratch the surface of the problems involved. Greatly though neuropsychology has advanced our understanding of the brain in the last few decades, it has not reached the point at which it can answer such questions, and it is doubtful that it ever will. Another reason for not embracing any detailed identification between particular mental and neural events, is that activation of the neural circuitry is - as is well known - only part of what goes on in the brain: a multitude of chemical changes is also taking place. We all know that a thought can make our hearts beat faster, our tongues cleave to the roofs of our mouths, our knees go weak, etc. The relation between mental and endocrinal activity is so far much less well understood even than the relation betwen mental and neural activity. So, I shan't assume an identity between the arousal system in the brain and in the mind.

What we are talking about is 'one of the central mysteries of life and mind' - the source of all effort and attention. Let us see what we must attribute to this vital source. Generally, as we have seen, we want to attribute wakefulness to it. We also want to appeal to it to account for different degrees of push at different 'places' in the system, that is, for variation in strength between interests, and hence for variable responsiveness

to input. With no inner push, no interest going out toward the world, the world makes no impression. At one extreme, a sound sleeper can snore through an overhead thunder storm, at the other, a Toscanini can hear the fifteenth violin playing a wrong note. As well as increasing receptivity, the energy we invest in our interests is also creative, producing outer and inner expression.

It's important to remember that this phenomenon of the variable allocation of 'mental energy' must embrace an extraordinary range in human beings. It must include the difference between a quite hungry person, a very hungry person, a satiated person, and someone from whose mind the thought of food could not be further. It must also include, for example, the difference between a woman who has never thought to question her role in life, and that same woman once she has 'had her consciousness raised' and come to see her role as an issue. (Where the hungry person is alert to the smell of food, the feminist is alert to the smell of male chauvinist piggery.) And it must account for the difference in behaviour between someone who has recently heard the phrase 'terrible error' and someone who has not, when asked to say 'gad boof' out loud (see Baars et al, discussed in Mis-words , above). And so on.

As well as talking about different levels of arousal, we are talking about different foci of arousal or concentrated energy: we are not just talking about different degrees of interest, but about different interests - many of which may be active to some degree at the same time.[7] To what extent distinctions between areas of activation depend on distinct stimuli is an interesting

question, and it brings us to the second source of activation in the mind's arousal system, namely input from the 'outside world'. The notion of an interest surely presupposes its being an interest in something. Different interests reach out towards, and are receptive to, different 'objects', what else could distinguish them? There are two kinds of answer to that question, one concerns the possibilities of innateness, the other the nature of the 'objects' in question.

There are, apparently, ants which give off a characteristic smell when dead; as soon as they start to smell, the other ants pick them up, carry them out of the nest, and throw them away. A human observer of this behaviour made an extract of death smell, blobbed it on some living ants and, one at a time returned them to the nest. One by one, kicking and struggling, their fellow ants - who knew a dead ant when they smelled one - marched them off and hurled them to oblivion.[8] This narrow-minded imperviousness to contrary indications is typical not only of ants but of innately programmed behaviour in general. Take a richly preprogrammed creature out of the environment in which its program has developed over the preceding generations, and it will be at a complete loss. Activating an innate action schema of this sort entails energies being spent in a particular way in response to a particular outer stimulus. The individual may never have encountered that stimulus before, but the species has. So even innate distinctions between interests ultimately depend on distinct stimuli from the outer world.

If every interest has an object, what is the object of the ant interest just discussed? The object is that to which the

blueberries, then I probably won't notice the little grey and brown fungi nestling in the grass beside the blueberry bush. But I won't miss the large - edible, good - bright yellow Suillus Grevillei which grows there too. If I'm chatting with a friend at a party, I'll hear our conversation, but not the conversations going on around us. But if someone mentions my name, or that of someone I care about, then I probably will take that in - my attention seized, my ears 'prick up'.[9] Or a sudden loud noise may interrupt our discourse, distracting the attention of everybody in the room. All these are cases of an interest assuming currency as a result of something out there 'summoning my attention'.

Degree of inner push - prior interest - and degree of conspicuousness, distinctiveness, of the stimulus clearly interact in determining whether it becomes an object of interest, in determining whether I focus on it or not. There are quantifiable aspects of some of these stimuli, which clearly affect their salience (their capacity to arouse an interest), such as the brilliance of the mushroom, or the loudness of the noise. But even these cases presuppose distinctness against a background- a bright yellow mushroom won't stand out much against a bright yellow background, one loud (meaningless) noise won't stand out in a cacophony, and so on. What each of the cases has in common is that my attention is attracted by something which doesn't fit my current anticipations (my 'set'). In each case, the refocussing of attention involves either new anticipations coming into play, or new questions arising.

One lesson to be drawn from this is that what 'gets through'

is more than whatever is the current object of interest, whatever is currently being focussed on. Otherwise, nothing from outside could pull the interest away, because it would not be able to reach the attention (the arousal system) in order to affect it. The arousal system is evidently active in a variety of different directions a great deal of the time.[10] Anticipation is a readiness to act and react in a manner fitting to this occasion: the inner push which comes out to meet a heterogeneous world brings heterogeneous information with it. Different interests are activated: different schemas direct the energies, prepare the person as well as possible for what's next on the basis of what has been.

Just as the effects on me of this encounter with X depend on an interaction between X and the current form and level of arousal of my interest in X, so have all my past encounters with it depended on this interaction. What X means to me now depends on this subjective history. The world 'gets through' insofar as it activates the arousal system to some degree. And it means something to me if it activates some 'place' within the system - if it is a reaction distinct from other reactions, an interest distinct from other interests - rather than activating the system at large (as in a startle reaction at, e.g., a loud noise). Interests are distinct from one another insofar as their activation entails a flow of mental energy this way and not that way. If repeated encounters with X bring about no change in the form of my interest in X, if each encounter conforms with my anticipations, then these encounters become less interesting, less demanding of mental energy, and X becomes less likely to be

an object of interest. Once I am thoroughly used to X, then when I encounter it the merest 'activation' occurs; I register it without focussing on it, without thinking about it I'm ready to deal with it appropriately if required.

* * *

People (and the best works of art) defy the process of habituation. It is almost impossible to be as indifferent to the presence of another person in the same room, as one may be to the carpet, the walls. For one thing, the carpet and walls won't (except in the most peculiar circumstances) do anything, for another, they won't notice anything I do. Whatever they may mean to me, I am not constrained to mean anything to them. As a human being, my very survival depends on securing the cooperation of other people, and that entails getting them to see my interests as their own. (After a bath one day, my six-year-old was shivering violently. I suggested that he was hamming it up rather. "But I am cold," he said - I indicated belief - "and if I don't do this, nobody will know.") All my waking life, as I move through a shared world, I must be alert to what I mean to other people, as well as to what they mean to me.

Cooperation between ants may well be achieved entirely by imperatively meaningful smells, but between human beings it takes work, the results of which will never be fully predictable. This is where language comes in, of course. One way in which language facilitates the construction of communities of interest, is by making it possible for its members to express interests to each

other explicitly and directly. Questions, requests, pleas, commands and so on, are expressions of interests which require cooperation for their pursuit. The right to put forward a request, a command, etc., and the obligation to respond, vary widely both within any given society, and between societies. All societies have ways of allocating these rights, because they are an essential foundation of social order, but all individuals know ways of bending and violating their obligations. While language makes possible the passing on of precise and detailed information, at the same time the devil gets in with mis-information, misrepresentation, plain falsehood, and bullshit.

Truthfulness can be an issue for any new assertion. So, as well as providing the means to increase the pool of common knowledge - the essential basis for all common action, for the pursuit of all common interests - language also sows the seeds of a potentially endless succession of issues. It makes it possible to talk about what is in doubt, to debate, to put forward arguments pro and con. Then again, that same power - the power to say what is not so - is also the power to tell stories, and it is the power to project possible futures into public space. Language also makes it possible to reminisce about the past, and to attempt to explain the present. And it gives one a way of talking impersonally, of referring to one rather than to me or you or him or her, to someone and anyone rather than to this one or that one. It also confers the capacity to place events along a fixed time scale (times and dates) as well as in relation to some now (tenses).

All these effects of language could be summed up by saying that it provides the means of acquiring and conferring alternative points of view. It can give one a view that is not just from not here and not now, but from not me. It does this partly via other people telling one (disputable) things, but - more fundamentally - it does so via the (quasi indisputable) meanings of the words we all use and understand. This is presumably what Brentano meant when he remarked that language provided a 'sort of preliminary classification' (1907). In every individual, language first takes hold as a way of expressing interest; as we learn better and better how to do so (how to make ourselves understood), we are acquiring the system of classification that is built into our language. Once we have acquired it, it lies as inescapable background behind not only all linguistic encounters, but all future learning.

Although the relatively fixed system of intersubjective meanings transcends my subjective history, it is nevertheless part of it, and part of what I know. My interest in, say, this cat, will always activate, be informed by, what I know about cats, and that is a matter of knowing what "cat" means. As I pass through the world, being affected by it, the schemas which get activated are in fact organized (held in a fixed relation to each other) by the language I speak. And, as I look around I see nothing I can't name if I want to. An encounter with a familiar and uninteresting object will generally activate, to some slight degree, a certain schema within the whole. But words, those most familiar but potent of objects, appear to activate connections between places in the network too; hence many more than one schema, - or node? - get activated by just one word[11]. Most

words mean more than most things (to most people).

For the individual, acquiring a language means inter alia acquiring a relatively orderly mind, which furthermore is organised along the same lines as those of others sharing the same world. It also means, insofar as we're hearing and using language all the time, that the semantic system is kept perpetually at the ready. So, we all have a rich range of highly reliable and interconnected common knowledge, which we never have to think about, and which we can call on at any time with a minimum expenditure of mental energy. But language does even more for our mental life than that: finally we reach the third source of activation of the arousal system - and that is the material in which I have called 'cognitive space'. The flow of directed energy we have been noting has results; we are creative as well as receptive. Interests near peak arousal are foci of mental energy waiting to discharge in expression of some sort. As we saw in Chapter 5, as well as actions, including speech, among the products of this focussed energy are 'images' and m-utterances

In this chapter, we have been seeing how a differentially activated, informed, arousal system functions as preparation, anticipation, for the immediate future in a non-uniform but largely predictable world. I am suggesting here that in cognitive space we can go one better, and project ourselves into an imagined future, in it we can enact possible futures with merely imagined consequences. It is where the imagination plays. It seems highly likely that all 'higher' animals have imaginations. The capacity to solve certain problems - like the

monkey's using a short stick to get a long one to get a banana - can be readily explained if we assume the animal can rehearse the needed actions within its own cognitive space. Such rehearsal would only be useful when the actions in question presented a problem of some sort, when they were not habitual.

As human beings, the future we're trying to be ready for includes a future of dealing with other people, as well as things. Few things create more problems than other people do; as we've been seeing, solving these problems typically involves talk. It's no wonder that human cognitive space tends to be full of unuttered sayings, with them we put our inner actions into a form with intersubjective meaning. (If all the world's a stage, every mind is a rehearsal room.) Therefore, by creating sayings in my mind, I connect my would-be actions to the knowledge system, to information not so far taken into account. I become audience to myself, and potentially therefore also judge. If I'm seriously trying to project the future, then I'll be alert to the plausibility of what I imagine, and to the truth of my thoughts. If, on the other hand, I am daydreaming, then whether what I say or see 'feels right', sounds right, doesn't matter in the slightest. I can listen to myself telling stories with the same suspension of the demands of truthfulness as I apply to all fiction. Or I can attend to the likelihood or truth of my inner sayings with all the scrupulous rigour of a cynical outsider. Equally, I can apply those strict standards - "<u>can</u> this be right?" - to my (other) considered actions, or I can indulge myself with pure fantasy.

Not all one's interests are directed at the future, although

the great majority for purely practical reasons must be. One can recreate, and thanks to language reflect upon - make judgeable objects of - the past as well as the possible future, and the impossible daydream. The capacity for reflective self-control which language combines with the imagination to provide, is the capacity to reject one's own would be actions, or claims. Hence it is the capacity to aim at the good and the true, and it is therefore the origin of sin.

Chapter Ten: Conclusions

In this chapter I begin by presenting a general sketch of the picture of 'the mind' which has emerged so far in this thesis. It is a dual picture, comprised of the system of interests on which all meaning depends, plus cognitive space, within which we may pretend to act and speak. From this general picture I draw certain further conclusions, particularly about the form and content of interests, and about the nature of interest systems. Next I compare my views with certain other current approaches and find a number of points of contact, despite wide differences in perspective and vocabulary. The most radical differences between my views and those which prevail in this field, concern the relation between beliefs and desires (or as I would prefer to put it, between assumptions and interests). Both are supposed in many quarters to be 'propositional attitudes' of equivalent psychological status, and it is supposed to make sense to treat them as though they were in separate 'boxes', one labelled beliefs, the other, desires.[1] Far from their being independent of each other in this way, in my view every assumption is interest-relative, and every interest is informed. Where interests are given the crucial organising and initiating role I think they deserve, it tends to be in the guise of goals, plans, purposes, etc. (see, e.g., Anderson, 1975; Schank, 1975; and discussion in Chapter 9 above). Though end-oriented interests like those account for a large share of a person's cognitive activity, they do not exhaust it. A recurrent

If an interest were not a repository of information, then what activities it led to would be random, and it would no longer be possible to understand those activities in the light of that interest. If interests were not informed, then they would not be distinct; the uncontrolled flailing of a neonate's limbs bears this out. Input which informs an interest changes its form, and therefore its future manner of pursuit; so interests change and develop and adjust themselves to new information. In that respect they are clearly equivalent to Neisser's 'schemas' or Minsky's 'frames': in my terms, a schema or frame is the form of an interest.[4] It follows from these considerations that a person's system of interests is also their system of assumptions or beliefs.

The form of an interest is its schema, and that is constituted of its set of relevant assumptions, for an assumption is no more than a difference which has been made to an interest. Interests may be connected with each other in varying degrees. At one extreme, they may be embedded in each other so that all assumptions relevant to one are relevant to another, but not vice versa; or they may be connected only in having some shared relevant assumptions; or, at the most tenuous extreme, they may have a purely associative connection - so that, for example, every time N sees a pretty stamp she thinks of M. So, all together, the set of a person's interests constitute a densely, and variably, interconnected network, the parts of which at any given waking moment will be at greatly varying degrees of arousal. For any individual, amongst the connections, and assumptions, which are most reliable, and whose activation is

closest to being perfectly automatic, will be those laid down in that person's language. And those connections are rational connections: it is not a mistake to rely on them, either in discourse or in practice; energy spent activating them is not wasted.

Although the best-maintained, most firmly interconnected, and soundest parts of the system may be linguistic, there is no reason to suppose that the linguistic sub-system is cut off from everything else a person knows.[5] On the contrary, if it were an isolated system, then the information it encodes would have to be duplicated in the rest of the system. For that information is practical, i.e. it affects the way one approaches the world. This is so even at the most 'analytic' - i.e. rigid, reliable, universal - extreme. To take a classic case: if N wants to marry M, and she knows he is a bachelor, then she knows he is not (at the moment) a married man, and that the object of her desire is not ruled out, at least on those grounds. The linguistic subsystem is distinguished primarily by its relative immunity from the whims and preferences of individuals, because it consists of what 'we all' (i.e. we speakers of L^6) assume we can all assume.

Generally, activation of a part of the interest system makes relevant information ready for use. It seems reasonable to assume that activation has a long-term as well as a short-term effect on the accessibility of information. Like the tunnels in a mine, connections which don't get periodically reactivated are liable to fall into disrepair, and unless the assumptions they once joined can be reached from another direction they will become inaccessible and hence unusable. Forms will disintegrate, schemas collapse, unless their interests are sometimes aroused.

We keep the language system so regularly in use that it's no wonder its meanings are generally so effortlessly available. A function of m-utterances is that of keeping the language system busy.

So, the form of an interest - what it takes for granted, what it has been informed by - is its schema. What then is its content? Jointly, the set of all a person's schemas constitutes a network within which activation variably occurs, and from which activities variously emerge. An active interest is replete with energy, to which it imparts a more or less precise direction; an inactive interest does nothing, is mere contentless form. As well as input (impression) causing activation to spread from points of entry into the network, concentrated energy builds up at points in the network, i.e. within particular schemas, until they are bursting for expression (output). One could think of it as horizontal movements of energy on the one hand, and vertical ones on the other. An interest is a potential hole in the fabric of the taken-for-granted; a way in or out of the network, through which change can enter or leave the system. The content of an interest is the energy it directs.

What can we say about what happens, on the one hand when energy spreads through the system, and on the other when it is disgorged in expression? Ex hypothesi, the horizontal movement activates some set of assumptions - but then, what does that do? A regular laboratory finding is that reaction times to consecutive related stimuli speed up, whilst an unrelated intrusion will elicit noticeably slower response.[7] In other words, the whole activated area becomes more highly aroused, and

therefore more alert to new relevant information, more susceptible to change. Also, it must in some way make the information encoded in that part of the system more sensitive to whatever initiated the activation. New information may change the schema by accretion, it may combine with it to yield new conclusions, or it may have the more radical effect of altering previously held assumptions. Input which runs counter to - i.e. would rule out - what has formerly been assumed relevant to the activated interest(s), will have its non-conformity detected, though not necessarily understood. A clash with the taken-for-granted creates a disturbance, a hot spot, an interest in its own right: an issue. And the assumption(s) with which there is a clash may eventually be dislodged. Whereas input which conforms with what has previously been assumed will tend to solidify and entrench that set of assumptions. It will serve to keep that part of the network in shape, maintain its form.

Now let us consider what happens when energy is released in expression. In some way, the manner in which an interest gets expressed depends on its form, i.e. the place in the network from which the energy springs forth, the assumptions which inform it. We have found (Chapter 5) that an interest tends to produce activity which leads towards its satisfaction (which in some but not all cases involves its closure). How does what one assumes affect that activity? Let us take a very simple case, suppose I want to sit down. Swiftly, confidently, with a well-aimed movement, I place my bottom on the chair. In doing so, I assume that the height and width of the chair seat are what they have always been, that it is in the position I have judged it to be,

expressions.[10] Lastly, since talking is generally less effortful than doing, verbal expressions presumably generally use less energy than more active forms of expression.

If talk takes up less energy than action of other sorts, presumably the quasi talking and acting which go on in cognitive space use up less energy still - no actual moving parts whatever being set in motion. Lazy though one may be in cognitive space, thinking can also be hard work - what good can it do to spend even small amounts of energy on pretence? - for that is what m-activities amount to. As in real space, the activities one performs in cognitive space are expressions of some highly aroused interest, the form of which guides the forms of its expression. What m-utterances and m-actions occur to a person over a period of time will depend on what interests come into play during that time. Some of these interests may be purposeful, others not, covering a range from problem-solving to daydreaming, and including such short-lived interests as a flash of surprise, delight, irritation, etc.

Some of the material in cognitive space is there solely in response to encounters with the here-and-now. There are the comments I m-utter as I proceed in the world, on things and events of passing interest. With these I sum up what strikes me, generally what is unanticipated. By putting it in a language, and thereby activating (making accessible) assumptions relevant to it, I make the new familiar. At a moment when expression in public space is ruled out, and there is no action to be contemplated or taken, no conclusion to be considered or reached, no change to be made, but there is nevertheless a head of

focussed energy pressing for release, then m-activity of some kind may be the best one can do. On the one hand, its expression will let off enough pressure that it will soon sink back below the threshold of peak arousal. On the other hand, at the same time, because of the feedback effect, the expression of an interest will serve to maintain it at a sufficiently high level of arousal that it will sooner or later cross that threshold again and achieve expression again - next time, perhaps, in real space.

Unlike exclaiming and passing comment, daydreaming and dwelling on the past are m-activities which are both purposeless and incompatible with any but the most mechanical and undemanding engagement with the here and now. Is not energy devoted to them pure waste? The answer in both cases is No, at least up to a point. One can go on learning from the past long after it's gone by; and even when there is nothing more to learn, one may still get pleasure from it. Only when every shred of meaning has been chewed over and digested, and when all that's left is bitter, does it become entirely futile to dwell on the past. Daydreaming doesn't carry that risk of futility, being as far as I know unfailingly devoted to self-pleasing ends.[11] The danger with daydreams is of forgetting that fantasies are what they are. The only fundamental difference between daydreaming and serious thinking is that in the former but not the latter one has free rein to pursue the entirely unlikely. Much follows from this.

All our serious interests are pointed at, geared towards, the future on the basis of the past - they'd be useless repositories of information if that were not so. The future

indicated by an interest may be immediate or mediate. The immediate future is an unbroken extension of the here and now, it is about to be actual. Mediate futures are not about to be actual, are projected, away from the here and now, into cognitive space as possibilities. As we have seen (Chapters 4-6), a great deal of our cognitive activities in both public and private space is devoted to consideration of the possible. We consider possibilities in order to assess their likelihood and their desirability. As Sperber and Wilson put it, "Assigning a subjective probability value to a proposition involves representing to oneself the likelihood of its being true, and assigning a wantability value to a proposition involves representing to oneself the desirability of the potentially achievable world in which it would be true."[12] I find out about how likely and desirable a state of affairs appears to me (i.e. assign it a subjective probability or wantability value) by imagining - i.e. representing to myself - how things would be if it were actual. Obviously such m-activities are only required when there is more than one available interesting alternative.

The more alternatives one can rule out, the more narrowly one can focus one's energies, the more efficiently one can spend them. Although it's only sometimes fun, serious thinking is more valuable than daydreaming because it leads to the exclusion of alternatives, rejection of courses of action, closure of exits - and all this without lifting a finger. Once there is no doubt left and only one alternative remains, then that will henceforth be taken for granted. Subjectively, 'settled conjecture'[13] is treated as certain; this is the limiting case of likelihood.

Once an issue has thus been settled, it will cease to trouble or disturb, cease to occupy cognitive space. From now on it will function - if at all - merely as an assumption relevant to some interest(s), as an effect on the future flow of energy. It is in the nature of an interest to create activity which tends towards its satisfaction. What satisfies a question or an issue is the annihilation of relevant ignorance or doubt. Even in self-consultation, what brings that about is, typically, debate.

In just the same way as considerations that weigh only on one side appear in a daydream, so considerations that weigh on either side appear in a m-argument. This is realism - in the popular rather than the philosophical sense. Being realistic requires one to resist as far as possible the lure of preference and desire, and keep the line between debate and daydream as sharp as possible. It is never reasonable not to be realistic if the matter has any bearing on potential proceedings in the world. A language encodes a realistic world view, and, being public, cannot be changed to suit an individual's whims. Furthermore, assumptions which appear in cognitive space as pro or con some m-argument, do so in a language, and as statements. As such, they lay claim to objective truth, and in principle call for assessment in the light of everything one knows.

If we proceed in the world - in actual and public space - without adjusting our interests to it, then at the very least we'll try to do what can't be done, and at the worst we may suffer or even die for our mistake. Knowing how to avoid such mistakes is the core of reasonable being; all living beings must have some means of doing so if they are to survive. But the

language-given capacity of inner debate adds a means of doing so which is presumably uniquely human - and at the same time it vastly increases the human capacity to make mistakes. Successfully adjusting one's interests to the world through inner debate entails recognising contradictions and consequences. A statement feeds back into the system as though someone else had said it, so that contradictions will be felt and consequences will follow. Consequences and contradictions which follow directly from the structure of language - i.e. from the assumptions which are encoded in the language system - will do so with near automaticity. But although a claim to truth in principle constitutes a call for assessment in the light of everything one knows, in practice variable accessibility distorts the picture. Hopes, desires, fears, whatever interests are most highly aroused at the moment of input, those will be favoured in its interpretation. Language creates an illusion of perfect objectivity and warranted certainty at the same time as it provides a means of struggling towards those ideals. As Dewey put it, "Perfect certainty is what man wants."[14] Tough. Only death and the extinction of interest brings that.

* * *

We saw in the last chapter that the notion of 'drive' was a cousin of the idea of interest as an activating force that has been so central in this thesis. Until, roughly, the late fifties - early sixties, some relation of interest - drive, instinct, desire, motive, intention, etc. - had been of central importance

in all twentieth-century psychological theories. Indeed, Bartlett - who is famous nowadays for his promotion of the idea of a schema - certainly took himself to be emphasising the crucial role of interests as such, in both the formation and the later recall of memories (see Bartlett, 1954). What happened? Where did interests et al disappear to? Briefly, the demise of behaviourism discredited its whole theoretical apparatus and the notion of drive along with it. As 'cognitive' psychology took over from behaviourism, there seemed no longer to be any place or need for the concepts of drive or energy (see, e.g., the introduction to Neisser's Cognitive Psychology, 1963). More refined versions of the notion survive, however - its most common appearance nowadays being in a box labelled Propositional Attitudes: Desires (see, e.g., Fodor 1981, Stich, 1983). Apart from the goal oriented theories which we have seen are too narrow, only in clinical, i.e. practical, psychology have concepts of this ilk continuously held a central theoretical role. I shall take up some comparisons with psychoanalytic theory below. Before doing so, I look at some of the consequences and problems for desires-in-a-box / beliefs-in-a-box theories.

The style of thinking in which desires, hope, fears, etc., and beliefs are all treated as objects of propositional attitudes - i.e., as propositions to which an attitude of desire, hope, etc., is taken - derives from philosophy, and permeates what its exponents refer to as 'cognitive science'. (Stich's contention (op. cit.) that this view holds a central place in 'folk psychology' strikes me as bizarre.) Stating a desire, etc., or

handling it in a logic, standardly involves representing it as a proposition embedded in a that clause - as in "I hope (that) John doesn't get measles". There is nothing wrong in doing so, of course, the problems arise when conclusions about the mind or about the nature of desire, etc., are derived from this manner of representation. (Quite apart from the absurdity that it may seem to make sense to think of people's minds as stuffed with boxes full of propositions [- a different box for every attitude?].) My main reason for eschewing this mode of speech is that it treats desires, etc., as of the same order as beliefs, qua assumptions.

'Belief' has two senses which I believe it important not to confound. One is what I have been referring to as 'assumption' i.e., what one takes for granted, proceeds on the basis of, what has made a difference to the outcome of an interest. 'Belief' in the other sense is more or less equivalent to what I have been calling an 'issue'; that is, it is not taken for granted, but on the contrary, is uncertain, in doubt - awaiting resolution. Once an issue has become settled, a definite outcome achieved, then that outcome becomes a 'belief' of the other sort, namely an assumption. An issue is a creative force in its own right, a place at which energy peaks, given a direction by prior assumptions, i.e., it is an interest. But prior assumptions function to direct energy towards satisfying the interest they inform. Because they are taken for granted, we waste no energy on assumptions; while issues require an investment of energy leading to expression, from which the benefits of feedback may be reaped.

Issues and questions are the peculiar case of interests which lead directly to the establishment of changes in the form of the network, rather than leading directly to (real or imaginary) changes in the world.[15] So they are central to the business of the (rational) 'fixation of belief' that Fodor pinpoints as the proper subject matter of cognitive psychology. The move from issue to assumption is the move from tentative unsettled belief to fixated belief. Here is Fodor musing on what it is to acquire an assumption: "I assume that every yes-box [i.e. assumption 'box'] is connected with an elaborate mechanism of wheels, pulleys, relays, and so forth, such that putting a mental representation token in the box has a correspondingly elaborate variety of causal consequences, both for the behaviour of the organism and for the distribution of other mental representation tokens. Moreover, which causal connections you get from putting a given token into the yes-box depends, again elaborately, on (non-semantic) properties of the token. (Maybe it depends on its weight, or its shape, or its electrical conductivity.)" (1983, p. 28; my emph.)

So, the picture is of a 'yes box' full of bits of information with specifiable interrelationships, which 'elaborately' depend in some way on the form of representation of each individual bit. Furthermore, each of these assumptions must have a representation the form of which will make sure that the "wheels and pulleys" which will be set in motion have the right consequences for the behaviour of the organism. Whether acquiring new assumptions or bringing old ones to bear, the problem is to get the right beliefs to inform the right desires

so that appropriate actions and responses will result. Since our every action proceeds on the basis of relevant assumptions, and so does our every word and every thought, different sets of assumptions must constantly be coming into play. If they are to be picked out afresh by some mechanism each time, then they must display unmistakably the characteristics which fit them for the role they are about to play.

How a model or theory of cognition distributes the information load when representing a cognitive system will be determined partly by what theory it is, and what its exponent is trying to do, and partly by aesthetic considerations such as elegance and simplicity. At one extreme, one could conceive of all information being jumbled together in no particular order in the 'yes-box'. Or one could assume, for example, that assumptions which have certain identical consequences were grouped together, forming a subsystem of the whole. Then by marking an assumption for its appropriate subsystem(s), the information load which it carries could be greatly simplified. The view I have been expounding is at the furthest extreme from the yes-box rag-bag, which leaves all the elaboration to the individual 'belief representation token'. From my standpoint, a cognitive system is primarily a collection of diverse and densely interconnected interests, each of which is a locus of information, both in respect of acquisition and of access. So the organisation of assumptions is derived from the interests they cluster round and give form to. For the vast majority of one's actions one doesn't have to assemble and organise the relevant beliefs freshly each time - they're all there, ready for

use, taken for granted, constituting together the schema of this action or that.

From this perspective it appears that almost the whole consequentiality of a newly acquired assumption will depend on the spot(s) it hits in the system, i.e. the interest(s) it affects. What spots it hits will be a result of the interaction between prior levels of interest in the arousal system, and the unique character of what is being learnt, i.e. it will depend on what it is relevant to, for N, at a given moment. Once an assumption has been accepted and absorbed, has had its effect on some part or parts of the system, it may subsist thereafter purely in the consequences it has had - which remain. The content of an assumption as it comes in is a matter of where activity occurs in the system in consequence of its absorption. Its lasting 'representation' is as pure contentless form: unenergised, taken for granted, effect on future proceedings.

So, far from being in separate boxes, assumptions and interests (which of course include desires, hopes, and all those others) are complementary, interlocked, interdependent, but distinct, aspects of the same phenomenon: namely, the interest system as a whole, through which all expenditure of any intelligent being's energies is directed, organised, informed. In storage an assumption is utterly context-dependent, being simply part of the network, its meaning deriving from which part. In order to communicate, discuss, or examine an assumption, one must a) energise it, i.e. reimbue it with content, and b) give it expression in some relatively context-free form, i.e., in a language. Articulating a previously inarticulate assumption is

not a matter of taking an obfuscating step away from the private lucidity of a proposition couched in a universal, context-free "Language of Thought" (see Fodor, 1975); rather, it is a matter of making a clarifying move towards an ideal of objectivity and explicitness.

A school of contemporary thought about cognition which has clear points of similarity with my own views as presented above is that of "spreading activation theory". Initially concerned with questions of retrieval in a semantic network (see Collins and Quillian, 1969), with Norman's picture of a network of interrelated schemas (1981) through which activation spreads, the scope of the theory has widened well beyond the strictly semantic. Since, as I have been arguing, 'my' network of interests can also be viewed as a network of schemas, 'my' network and Norman's are presumably in some sense equivalent. However, as far as I know, it has not occurred to current spreading activationists that their networks might be primarily assemblages of interrelated interests at different degrees of arousal.

How far and in what direction activation spreads on a given occasion must be a result of interaction between how highly aroused various interests are at that moment, how reliable the connections between them are, and how striking the input is. Therefore, the network's being derived from the set of a person's interests and displaying the characteristics of interests, is not irrelevant to spreading activation theory. (It means, for example, that it cannot be a theory from which precise predictions flow, because it can never be possible to know in

advance what particular interests are aroused in a given individual.) However, as we shall see below, network theories, and certain of my own findings, have been anticipated in the area of psychoanalysis. And in that area, of course, energy (as 'libido') has remained a central theoretic concept.

By far the most striking parallels between my own views (of the interest system) and those which currently prevail in cognitive science, are with a school of Artificial Intelligence researchers known as the 'New Connectionists' (see Dennett, 1984). It is worth quoting extensively from Dennett's summary of the typical characteristics of their models. "The most obvious and familiar abstract feature shared by most of these models is a high degree of parallel processing." In my terms, 'parallel processing' means that numerous different interests can be to some degree active - digesting information - at the same time. Connections go on being made by activated interests whilst other interests are taking in fresh information and being activated by it in turn. "Another typical feature is 'distributed' memory ... in which disambiguation occurs only 'globally'". I take this to be a function of the context-dependency of assumptions. There is "No central control, but rather a partially anarchic system of rather competitive elements". As long as we discount the potential feedback from cognitive space, that describes the interest system perfectly. There is "no complex message-passing between modules or subsystems". Instead, computing information is a matter of being "appropriately connected to large numbers of similar units". On this point new connectionism is close to spreading activation theory. And lastly, there is "the

relatively mindless and inefficient making and unmaking of many partial pathways of solutions, until the system settles down after a while not on the (predesignated or predesignatable) 'right' solution, but only with whatever 'solution' or 'solutions' 'feel right' to the system" (my emph.). For the closeness of this last point to my view, see Chapters 5 and 6 above, but for some qualifications, see below.

Curiously enough, the New Connectionists, far from taking themselves to be modelling a system of interests, take themselves to be modelling features of the brain itself, and take the neuronal hardware as their starting point. I don't know what to make of this (see points raised in Chapter 9). However, I do know that, just as neither New Connectionists, nor Spreading Activationists see themselves as dealing with interests (or desires or any other member of that family), nor do they complete the picture of the mind as I have done, by bringing in cognitive space. Although I think there is a good deal that is right in the views about the nature of information systems or networks just discussed, those views are not sufficient either to constitute a cognitive theory in their own right, nor to effectively undermine other schools of thought. Among the points which require to be addressed are these: how come other approaches, resting on apparently radically different assumptions, can provide illuminating analyses of cognitive phenomena? If what one assumes is so attached to its context of interest, how is it possible to get two and two together, i.e. assemble information from different areas? And what about cognitive space?

So, first point: the proof of the pudding is in the eating.

Let us take a look at a conspicuously successful contribution to this area which makes different assumptions about assumptions, namely Sperber and Wilson's 'Theory of Relevance'. It is, in effect, a theory of that area of natural reasoning which concerns the assessment of consequences and contradictions which follow propositional input. Although such input may have propositional form and a truth value, I have been arguing that the assumptions which constitute the background are not independent objects of that kind at all. And yet, in practice, the interaction between the input and the background yields results which (with certain provisos, see Sperber and Wilson, 1985, and also below) conform with the inference rules of deductive logic, just as though background and input were in fact made up of propositional forms with truth values.

How does this appearance of good logic square with the non-truth-functional character of assumptions, and the 'making and unmaking of random pathways' that I have been asserting? In the first place, the distribution of information in the network is not random, but is clustered around interests so that information one has aready perceived as related is automatically stored together. Therefore, even when brand new pathways are being made from a certain starting point, they stand a good chance of making a useful connection. Furthermore, useful connections are more likely to come into play another time, and re-utilised connections will last: there are very many well-travelled pathways in the network, and most of them are well-motivated. And, of course, when the input is linguistic its primary destination is within the best organised, best motivated,

and best kept part of the information network.

Any given input meshes only with a limited sub-set of the total of a person's assumptions. The appearance of natural reasoning matching logic crucially depends on its being assessed for consequentiality only against a restricted background.[16] Anyone who wishes to model a particular case of natural reasoning must find an intelligible way of displaying the interrelationships between the input and the set of relevant assumptions: language is essential. The standards by which a model will be assessed as a good or fair representation of a set of assumptions will be according to how well the same input to either yields the same results. That is, the model should predict the conclusions the natural reasoner will reach. The set of linguistic propositions which represent the set of relevant assumptions must in some sense be mappable onto those assumptions. (See Kowalski, 1985, for the technical equivalence of 'procedural' and 'declarative' representation.) The propositional forms which appear in the model are theoretical constructs without which there could be no explication of the phenomenon in hand. Therefore, a theory which is designed to explicate natural reasoning is exempt from criticism on the grounds that assumptions are not in fact proposition-like objects with truth values. So, the actual psychological status of assumptions appears to be irrelevant in choosing between theories at this level of description.

From my angle, what distinguishes Sperber and Wilson's account from others which assume an identity between propositions and assumptions, is its dual emphasis on economy on the one hand,

and change on the other. Their principle of relevance is quite explicitly a principle of economy, which makes the conservation of energy a central consideration. It is, roughly speaking, that one will seek to be maximally relevant in return for minimum processing costs. The relevance of an input depends on the changes it makes to the background of assumptions. So conforming to the principle of relevance involves trying to make as many changes as possible in return for minimum expenditure of energy. That is a principle which must be implicit in any account, such as mine, which treats the mind as an economy, a system for the distribution and organisation of energy. Processing costs will themselves depend on the accessibility of the material background. The fact of variable accessibility, which is thus built into their account, has been a central theme of this thesis. Despite what may or may not be profound differences as to the status of assumptions, at ground level it appears that Sperber and Wilson's theory of relevance and my own are complementary rather than contradictory.

So, I believe I have answered the question as to how certain positions which seem to differ radically from my own as to the psychological status of assumptions can get good results if views like mine are in fact correct. The next question was, If assumptions are so context-bound, how is it possible to get two together with two, that is, to assemble information from different parts of the network? Anyone who has been reading this thesis so far will know that my answer to that is, by creating representations - especially m-utterances - in cognitive space (see especially Chapters 5 and 6). As far as I know, neither

Spreading Activationists nor New Connectionists have any answer to this problem, and nor do they assign any role to anything like cognitive space. 'Imagination' or 'consciousness' though not ruled out by such theories, is simply ignored by them, presumably because it is not regarded as relevant to what goes on in the network.[17] But, as we have seen, what goes on in cognitive space feeds back into the network, and can affect its form. It seems that one has to go back to Bartlett for recognition of this phenomenon within the area of cognitive theory. "Consciousness reflects the mechanism that enables the person to subject his own schemata to scrutiny"; "a way of turning round on the organism's schemata and making them the objects of its reactions"; "the past is being continually remade" (1054). So, the answer to the last question I listed above, i.e., What about cognitive space? is, generally, it's been forgotten or ignored.

As far as I know, within contemporary psychology, only psychoanalytic theory insists, as I do, on there being two radically different components of 'mind', within which entirely different mental processes occur. How closely does the psychoanalytic distinction between the 'unconscious' and the 'conscious' match my distinction between the network of interests and cognitive space? As Freud explained it, the unconscious is that part of mental life which must be postulated to account for observable behaviours but cannot itself be observed directly. In that respect, the unconscious is identical with the network of interests I postulate. Furthermore, the picture is of a pool of 'mobile energy' within which areas of excitation compete for expression with one another, and lead to the production of what

may be consciously observed.[18] Despite these parallels, Freud's account of the different processes that occur on one side or another of the boundary is incompatible with mine. Here is Charles Rycroft's helpful summary of Freud's views on the subject:

> Primary process thinking is 'characteristic of unconscious mental activity', 'images tend to to become fused and can readily replace and symbolise each other, it uses mobile energy, ignores the categories of space and time and reduces the unpleasure of instinctual tension by hallucinatory wish-fulfilment.' Secondary process thinking is 'characteristic of conscious thinking', obeys the laws of grammar and formal logic, uses bound energy, and is governed by the reality principle, i.e. reduces the unpleasure of instinctual tension by adaptive behaviour.' (Rycroft, 1968)

In my view, a number of distinctions are confounded here. One is the difference between activation within the interest system (cf 'mobile energy') which we know only by its consequences, and expressed energy which creates objects of perception in real and in cognitive space (cf 'bound energy'). That is where I believe the boundary strictly belongs. But in that case, images have no place among the 'primary processes': the difference between verbal and non-verbal activities in cognitive space has been confused with the boundary just drawn. I believe this confusion derives, at least in part, from the ambiguity of the term 'consciousness', which I touched on in Chapter 6. As well as being used to refer to what I'm calling 'cognitive space', the word also gets used to refer to that which is attended to.

In Chapter 5 we found that imagery was ruled by the individual's whims, with meanings derived from the individual's interest system. And - perhaps just because there are no rules for the objective repetition of an image - we found that images were relatively evanescent, unfixed, and hard to attend to. Much of our imagery is 'unconscious' in the sense of unattended, and its relative unfixity is the same thing as that tendency to fuse and be replaced that Freud stresses. However, these are characteristics not of activity within the interest network, but within cognitive space. M-utterances, those quasi-aural activities in cognitive space, contrast with images in being relatively fixed, precisely repeatable and easy to attend to. They use meanings which are not subject to the individual's whim but are derived from society's interest system, and must be combined in a rule-governed fashion. As far as I could tell (in Chapter 6), even unattended or virtually unattended m-utterances are in good grammatical, linguistic, and logical order. In short, many of the characteristics which Freud attributes to secondary, 'conscious', processes are derived from the characteristics of verbalisation.

The last distinction that I think should be taken independently into account here is that between id fancy at one extreme, and serious thinking at the other. Both 'hallucinatory wish-fulfilment' and 'adaptive' m-behaviour take place in cognitive space. The difference between them is not a crisp distinction, but is one of degree. It is not a matter of processes of different kinds, but of different interests being in play, which entail different standards being applied in the judgement of their output in cognitive space. Imagery may be

particularly potent in the wish-fulfilment line, but there are also m-utterances in my daydreams, and there is imagery in my serious thinking. And verbal, grammatical, rule-governed and attended to behaviour in cognitive space may be as fanciful as one pleases, so long as one does not take it to be true.

To sum up, I believe Freud was fundamentally mistaken in equating, as he in effect does, 'unconscious' processes with unreasonable, 'maladaptive' ways of thinking, and 'conscious' ones with the opposite. In effect, even his psychology, in which desires appear to play so central a role, is infected by that separation between 'rational beliefs' on the one hand and desires on the other, which blights views of cognition current today.[19] The contents of cognitive space need not be cool, orderly and reasonable, and the set of interests - as we have been seeing in this thesis - is the reverse of 'maladaptive', seethingly active and disorderly though it may be. One crucial respect in which the network of interests is not 'irrational' is in having as a subpart the semantic system of the person's language. That is, part of the network is in fact orderly and self-consistent on a large scale.

> "The facts of linkage between words and morphemes, which make the categories and patterns in which linguistic meaning dwells ... correspond to neural processes of a nonmotor type, silent, invisible, and individually unobservable. The structure of the matrix relations ... can only be determined by a penetrating study of the language spoken by the individual whose thinking processes we are concerned with, and it will be found to be fundamentally different for individuals whose languages are of fundamentally different types." (Whorf, '1956', probably written in 1937 - his emphases).

Whorf is not claiming here that all meaning is linguistic, nor that the semantic 'matrix' he describes exhausts the individual's capacity to perceive patterns and make connections. Rather, he is asserting that the connections and discriminations which are automatically and unreflectingly made by individuals who speak different languages will differ as much as the semantic systems of those languages differ. Do the findings of this thesis support that contention? Briefly, No. Individuals may take for granted any number of connections and discriminations not laid down prearranged in their languages. There can therefore be no guarantee that two speakers of 'fundamentally different' languages[20] may not, as it happens, make all the same connections and distinctions - ones laid down in A's language which differ from those laid down in B's coinciding with ones B has made as a result of her own interests, and vice versa. Although such a coincidence is highly unlikely, it cannot be ruled out a priori. Nevertheless, the findings of this thesis do support a weaker version of 'the Sapir-Whorf hypothesis'. Namely, that we "unconsciously project [our language's] implicit expectations into the field of experience" (Sapir, 1931).

Insofar as an interest system amounts to a world view, then in being partly comprised of the semantic system of a language, what language that is will affect the individual's world view. Is it legitimate, then, to speak - as Sapir does - of "the tyrannical hold that linguistic form has upon our orientation in the world" (op. cit.)? For several reasons, the answer once more is, No. For one thing, the semantic system of a language is not self-contained but is just a relatively fixed, self-consistent,

and readily available sub-system of the total network of interests for any individual, and the points made in the last paragraph apply. Everyone can and does make connections and draw distinctions which are not language-given.

A further reason why linguistic form cannot be held to exercise any 'tyranny' is that it is itself derived from common interests in some community or communities of interest, and encodes the shared assumptions which give those interests their form (see Chapter 8 above). Therefore what has a 'hold' over our 'orientation in the world' is not truly the linguistic form but rather the community or communities of interest in which a given language is used, i.e. it is 'society'. So, although a semantic system is not dependent on any individual's whim, its meanings for any individual will depend on the communities of interest within which that person plays a role. Precisely because speech, i.e. expressions in a language, whether uttered, m-uttered or written down, uses meanings which are independent of any individual's interests, it can be used as a tool for manipulating interest systems and therefore ultimately changing the available meanings within a community of interest. Language does not tyrannise thought.

Speech frees my thoughts from the interests which give rise to them, and enables them to affect any other part of any interest system within some language community, including my own. By affecting interest systems, my thoughts have effects on how energies are spent in the world. So speech gives my thoughts power in the world. As a self-consistent, always reliable and available system of well-attested assumptions on which I need

waste no thought, language frees my energies for tackling the new. It gives me maximum information in return for minimum effort. The price I pay is that of inescapably sharing many aspects of my world view with my fellows - but that is surely gain not loss; and of having part of me always on call, unable not to respond at least with understanding to input in my own language. But that power which same-speakers have of playing on my interests is a power I also enjoy over theirs, and we each have over our own.

Understanding the relation between language and thought has been one of my most pressing concerns for very many years now, and its fruit is this thesis. For myself I feel that appetite satisfied at last, and I hope yours is too, dear reader.

FOOTNOTES

Mis-Words

1. I invented this expression to cover a range of examples from actual slips of the tongue to covertly noted slips to the 'tip of the tongue phenomenon'. This is a revised and greatly shortened version of a paper given at L.S.E. in 1978. In almost exactly this form it appeared as Murray (1979).

2. Freud himself uses a narrower (though no easier to apply) definition of the class of errors which interests him: 'The disturbance could result from interference from outside this word, sentence, or context.' (1911, 56.).

3. Here, and throughout the thesis, I use 'expression' and 'express' to include the formulated but not necessarily uttered.

4. Cp. Dell and Reich (1977).

Chapter One: Introduction

1. I had thought of preserving the first drafts embedded in explanatory text, as an illustration of the natural messiness and unpredictability of thought. I have to thank A. Cormack for persuading me that would have been more ego trip than reader's aid.

2. I owe this phrase to Mike Lesser.

Chapter Two: Conversational Rules

1. See Levinson (1983) and references therein.

2. For discussion, see Kempson (1975, 1977).

3. Unfortunately, this antedates Sperber and Wilson's perceptive analysis of irony (1981).

4. See his comments on 'therefore' in the paper I am discussing.

5. See Schegloff (1977) for a very perceptive discussion of ambiguity.

6. See Halliday (1978) and Gumperz (1968).

7. Re the orderly succession of turns, see Sacks, Schegloff and Jefferson (1974) and Schegloff (1968).

8. This issue is not as clear-cut as I thought when I wrote this chapter. As a recipient, misinformation and sound information are all the same to me, so long as I assume they are true. It is as an actor or creator that the difference between good information and bad information comes into play and makes a difference to my life.

Chapter Three: Relevance

1. In Grice (1975) and see last chapter, above.

2. See Sperber and Wilson (1984, 1985).

3. See Minsky (1975), Neisser (1976), Anderson (1976), and Bartlett (1954).

4. See Sperber and Wilson (1985).

Chapter Four: Issues and Questions

1. Successive drafts of this chapter have been improved by the comments of Deirdre Wilson, Dick Hudson, Ruth Kempson, Colin McGinn and Gill Brown; to all of whom, my thanks, and the usual exoneration from blame. In almost exactly this form, this chapter appeared as Murray (1983).

2. Keenan and Schieffelin (1976) and Reichman (1978) both make issues central to the analysis of conversation. But the notion of issue they work with differ, I think, both from each other a from mine. Similarly, some of Schank's work in artificial intelligence has given "interestingness" a central place (1977; esp. 1978). While his focus is on facts or events, and he uses interestingness as a value marker, for me an interest is something which is pursued by a person or persons, and has its own inherent - though not fixed - value.

3. Bosley (1975) has a very interesting discussion of "the art of forcing conclusions" (see especially chap. 1).

4. This was written down very soon after it occurred; it was sufficiently striking to be easily memorable. Regrettably, the was no recording equipment on the spot, so if there are errors my recollection, we will never know them. But I am morally certain that any inaccuracies there may be do not affect the interest-structure of the discourse, which is what chiefly concerns me here.

5. This useful phrase was coined by Halliday (1975, 1978). S also Chapter 8.

6. This is part of a 90-minute conversation which I recorded (with the parties' knowledge) and transcribed in full.

7. Goodwin (1979) makes a detailed analysis of shifts of gaze in the light of what is new to whom. He remarks that Sacks had noted a preference for the new. Labov and Fanshel find that assertions about events known to all parties tend to produce a minimal response (1977: 101). See also Schank on "interestingness" (1978).

8. That Reinhart finds "aboutness" and "old information" as recurrent, apparently distinct uses of "topic" is not surprising. It is a case of the context-defining interest (aboutness) and its taken-for-granted background displaying their inherent reciprocity.

Chapter Five: Interests, Obsessions, and Images

1. For a view which is in certain respects very close to mine, see Schwarz and Schiller: "Below threshold excitations give rise to conscious experience when they receive an energic supplement from within the mental apparatus" (p. 4; 1970).

2. On 'imagery' see Stern (1932), Merleau-Ponty (1962), Borkovec (1976) ed. Pope and Singer (1978), Goodman (1982), and see also the discussion in Wingfield and Byrnes (1981) and the references therein.

3. Since writing this, I have been struck by the frequent cinematic use of this image.

4. On the somatic effects of thought, see, e.g., James (1890), Pribram (1976), Barthes (1979).

5. See Wingfield and Byrnes (1981) for a very thorough discussion of the issues involved, which in many respects supports the position I am taking here.

6. "A thought limited to existing for itself, independently of the constraints of speech and communication, would no sooner appear than it would sink into the unconscious ... we present our thought to ourselves through internal or external speech. It does indeed move forward with the instant and, as it were, in flashes, but we are then left to lay hands on it, and it is through verbal expression that we make it our own." Merleau-Ponty, p. 177, 1962.

7. For discussion see Sperber and Wilson (1985).

8. See footnote 5.

Chapter Six: M-utterances

1. For discussion of the idea of a 'stream of consciousness' (or 'thought') see Chapter IX of James's Principles of Psychology (1890), also Titchener (1909), and the collection o this topic edited by Pope and Singer (1978). Joyce's Ulysses (1922) and Canetti's Auto da Fe (1946) are two of the best among a host of literary examples.

2. For some relevant discussion, see Goodman (1982) and Dennett (1982); Meichenbaum (1976) and the references therein provide substantial empirical support for some of the contentions and conclusions of this chapter. For the notion of 'Inner speech' see Vygotsky (1962), and Sokolov (1972), though I think that school tends to underestimate the importance of other people 'in the head'.

3. Cp. Dennett, "we learned to milk each others' (and then our own) minds in certain ways" (1982).

4. For an illuminating analysis of challenges, commands, etc., see Labov and Fanshel (1975). And 2

See Sperber and Wilson (1985).

.7 On 'irrational' beliefs see Sperber (1982).

.6 For some contemporary discussions of 'consciousness' see ed. Underwood (1982), ed. Schwartz and Shapiro (1976), ed. Davidson, Schwartz and Shapiro (1983), and ed. Josephson and Ramachandran (1980). Ramachandran describes "an imaginary world which we construct in our minds. We then see ourselves as active agents striving to do things in this imaginary world" (1980). On a more general level, from their different perspectives, the papers by Underwood and by Shotter (1980) and that by Pribram (1976), put forward views very close to my own.

Chapter Seven: Meaning I

1. Grice (1957) also approaches questions of linguistic meaning in the light of more general aspects of meaning. But our conclusions differ considerably. In particular, in my view, successful linguistic communication is not a matter of getting other people to 'recognise one's intentions'. Rather, it is a matter of creating with one's fellows interests as like one's own as possible, by choosing the most effective ways of expressing them.

2. See Anscombe (1957) on mud pies.

3. Ryle is quoted in Geach (1965) as expounding a view of if-then which seems close to my own.

4. For an excellent general discussion of the issues of reference, etc., see Lyons (1977), ch. 7, and for the question of pointing in particular, see Wittgenstein (1958).

5. For some different perspectives on speech acts, see Lyons (1977), ch. 16.1, Levinson (1983) and Sperber and Wilson (1985).

6. For a very helpful discussion of 'sentence' meaning, from a linguistic perspective, see Lyons (1977), esp. Chs 11 and 12; see also Barwise and Perry (1983).

Chapter Eight: Meaning II

1. To my ear, the crucial explanatory character of interests appears to be a central and even dominating theme of Putnam's Meaning and the Moral Sciences (1978).

2. For the notion of 'methodological solipsism', see Fodor (1981).

3. Kripke's comparable views (1972) are not susceptible to the criticisms I level against Putnam's arguments here - though they are not immune from other criticisms.

4. cp. "We take as given the idea of distinction and the idea of indication and that we cannot make an indication without drawing a distinction ... There can be no distinction without motive, and there can be no motive unless contents are seen to differ in value" p.1. - G.S. Brown, 1969.

5. See Davidson (1967).

6. For discussion of 'revisability' see Quine (1953).

7. For discussion of 'family resemblances', and for views very similar to those expounded here, see Wittgenstein (1953).

8. For a view of the psychological status of language which is in many respects similar to mine, and is accompanied by many detailed linguistic arguments, see Givon (1982).

9. For the notion of 'meaning potential' see Halliday (1975).

Chapter Nine: Arousal

1. During periods of dreaming, one's interests from wishes, hopes, fears, to issues and questions, are clearly active and sending material into cognitive space. The crucial difference between dreaming and the waking material in cognitive space would then be that dreams are aimed solely at oneself, cp. Rycroft (1979).

2. For some of the differences between individual rats, see Brown (1952) and references therein.

3. This 'demise' can perhaps be dated to Chomsky's devastating attack on Skinner in 1959.

4. See Baars (1976) and Posner and Abelson (1977).

5. For schemas, etc. as a way of handling goal oriented research see references in fn. 4, also Neisser (1976) and Anderson (1976).

6. For the non-identity of brain activity and mental events, see Davidson (1970).

7. See Kahnemann (1973) and Wingfield and Byrnes (1981) for discussions of the phenomenon of divided attention - which is, on some level, just another name for 'parallel processing'.

8. This typical tale of insect inflexibility came from a child's book of bizarre facts, which we seem to have lost or destroyed. I apologise to the reader for my inability to supply a reference.

9. I trust this phenomenon of being permanently 'primed' for one's own name is sufficiently universally recognised that evidence from experimental psychology is superfluous.

10. See footnote 7.

11. See, e.g., Schwarz and Schiller (1970), Collins and Loftus (1975) and Dell and Reich (1977).

Chapter Ten: Conclusions

1. For a fairly extreme statement of this position, see Fodor, 1981.

2. Because interests are interconnected, and any part of the interest network, from the very narrow to the very broad, may be aroused, there can be no general principle for discriminating and counting them. 'An' interest is simply any subpart of the whole.

3. Shimon Tzabar's invaluable mycological expertise helped me pick this example.

4. Any input/output node of the network has a form given by its place in the network. The forms of its interlocking connected schemas constitute the network, and those schemas are what determine the form of the energies which pass through them. For the general idea of a schema (or 'frame') see Bartlett (1954), Minsky (1975), Anderson (1976) and Neisser (1976).

5. For the notion of a dictionary as a 'disguised encyclopaedia' see Eco (1984) and for further discussion and references see Haiman (1980).

6. Because of the considerations outlined in Chapter 8, the notion of 'a' language is an analytic fiction, essential for some purposes (see Chomsky, 1965) but distorting for others (see, e.g., Gumperz, 1968).

7. This is a corollary of the facts of 'priming' touched on in Mis-Words, above, and presupposed in all 'spreading activation' theory, see e.g., Dell and Reich (1977), Collins and Loftus (1975), Schwarz and Schiller (1970). For much helpful discussion, and further references, see Wingfield and Byrnes (1981).

8. For discussion of language and community see Gumperz (1968), and Halliday's penetrating analysis in Language as a Social Semiotic (1978).

9. For the notion of language as a self-consistent system, see Quine (1953 and 1960) and see Givon (1982).

10. For the notion of a language as a generator of endlessly many new expressions see Chomsky (1965).

11. For a brave empirical investigation of daydreaming - which generally supports my findings here - see Singer (1966).

12. See Sperber and Wilson (1985).

13. For beliefs as 'settled conjecture' see Grene (1974) and Polanyi (1960).

14. See Dewey's The Quest for Certainty (1930), in which many pertinent points are made.

15. See footnote 13.

16. On restricting the context, see Sperber and Wilson (1985 and 1985b).

17. For discussions of imagination and consciousness, see ed. Underwood (1982), ed. Pope and Singer (1978), ed. Schwarz and Shapiro (1976), and ed. Josephson and Ramachandran (1980).

18. Psychoanalytic theory developed a rich and explicit network theory quite independently of and partly prior to Quillian, etc., see e.g., Schwartz and Rouse (1961), Schwarz and Schiller (1970).

19. Cf. Nietzsche, "The misunderstanding of passion and reason, as if the latter were an independent entity and not rather a system of relations between various passions and desires; and as if every passion did not contain its quantum of reason" (first published 1901-6, this translation, 1967). From an admittedly slight acquaintance with the work of Nietzsche, it strikes me as likely that anything of value in this thesis was probably anticipated by him.

20. The very idea of a 'fundamental difference' between two languages is dubious, partly for reasons touched on in footnote 6, partly because so many similar interests occupy people the length and breadth of the world, and those interests will typically be reflected in the languages they speak.

Bibliography

Anderson, J.R., 1976. Language, Memory, and Thought. Hillsdale, N.J.: Lawrence Erlbaum

Anscombe, G.E.M., 1957. Intention. Proceedings of the Aristotelian Society.

Baars, B.J. & Motley, M.T., 1976. 'Semantic Bias effects on the outcomes of verbal slips.' Cognition 4: 177-187

Bartlett, F.C., 1954. Remembering. Cambridge: C.U.P.

Barwise, J. and Perry, J., 1983. Situations and Attitudes. Cambridge, MA: MIT Press.

Barthes, R., 1979. A Lover's Discourse. London: Cape.

Bindra, D. and Stewart J. (eds), 1966. Motivation. Harmondsworth: Penguin Books.

Borkovec, T.D., 1976. 'Physiological and cognitive processes in the regulation of anxiety'. In: G.E. Schwarz and D. Shapiro (eds.)

Bosley, R., 1975. Aspects of Aristotle's Logic. Assen: Van Gorcum.

Brown, G.S., 1969. Laws of Form. London: George Allen and Unwin.

Brown, J.S., 1953. 'Problems presented by the concept of acquired drives'. In: D. Bindra and J. Stewart (eds.).

Canetti, E., 1946. Auto da Fe. New York: Stein and Day. Also in Penguin.

Chomsky, N., 1957. 'Review of Skinner, B., Verbal Behavior'. Language 35: 26-58.

Chomsky, N., 1965. Aspects of the Theory of Syntax. Cambridge, MA: MIT Press.

Clark, H. and Marshall, C., 1981. 'Definite reference and mutual knowledge'. In: A.K. Joshi, B. Webber and I. Sag (eds.), Elements of Discourse Understanding. Cambridge: C.U.P.

Cole, P., 1981. Radical Pragmatics. New York: Academic Press.

Collins, A.M. and Loftus, E.F., 1975. 'A spreading-activation theory of semantic processing'. Psychological Review. 82: 407-428.

Collins, A.M. and Quillian, M.R., 1969. 'Retrieval time from semantic memory'. Journal of Verbal Learning and Verbal Behavior, 8: 240-47.

Davidson, D., 1967. 'Truth and meaning'. Synthese VII: 304-23.

Davidson, D., 1970. 'Mental events'. In: L. Foster and J.W. Swanson (eds.), Experience and Theory. Amherst: U. of Mass. Press.

Davidson, R.J., Schwarz, G.E. and Shapiro, D., 1983. Consciousness and Self-regulation, vol. 3. Chichester: John Wiley and Sons.

Dell, G.S. and Reich, P.A., 1977. 'A model of slips of the tongue'. In E.L. Blansitt and R.J. Di Pietro (eds.), The Third LACUS Forum. Columbia, S.C.: Hornbeam Press.

Dennett, D.C., 1982. 'Why we think what we do about why we think what we do: Discussion on Goodman's "On thoughts without words"'. Cognition 12: 219-227.

Dennett, D.C., 1984. 'Computer models and the mind - a view from the East Pole'. The Times Literary Supplement, Dec. 14.

Dewey, J., 1930. The Quest for Certainty; a study of the relation of knowledge and action. London: George Allen & Unwin, Ltd.

Donchin, E., McCarthy, G., Kutas, M. and Ritter, W.E., 1983. 'Event-related brain potentials in the study of consciousness'. In: R.J. Davidson, G.E. Schwarz and D. Shapiro (eds.).

Eco, U., 1984. Semiotics and the Philosophy of Language. London: Macmillan.

Ellis, A.M., 1977. 'On the Freudian theory of speech errors'. In V. Fromkin (ed.).

Fodor, J.A., 1976. The Language of Thought. Hassocks: Harvester Press.

Fodor, J.A., 1981. Representations. Hassocks: Harvester Press.

Freud, S., 1911. The Psychopathology of Everyday Life. Trans. A. Tyson. Hammondsworth: Penguin, 1975.

Garfinkel, H., 1967. Studies in Ethnomethodology. Englewood Cliffs, NJ: Prentice Hall.

Geach, P., 1965. 'Assertion'. The Philosophical Review LXXIV, No. 4.

Givon, T., 1978. 'Negation in language: pragmatic function ontology'. In: P. Cole (ed.), Syntax and Semantics, 9: Pragmatics, New York: Academic Press.

Givon, T., 1982. 'Logic vs Pragmatics, with human language as the referee.! Journal of Pragmatics, vol. 6: 81-133.

Goffman, I., 1971. Relations in Public. New York: Basic Books.

Goffman, I., 1975. Frame Analysis. New York: Free Press.

Goodman, N., 1982. 'On thoughts without words'. Cognition 12: 211-17.

Grene, M., 1974. The Knower and the Known. California: U. of California Press.

Grice, H.P. 1957. 'Meaning'. The Philosophical Review LXVI: 377-88.

Grice, H.P., 1975. 'Logic and conversation'. In: P. Cole and J.L. Morgan (eds.) Syntax and Semantics 3: Speech Acts. New York: Academic Press.

Gumperz, J., 1968. 'The speech community'. In: International Encyclopaedia of the Social Sciences. London: Macmillan.

Haiman, G., 1980. 'Dictionaries and encyclopaedias'. Lingua 50.

Halliday, M.A.K., 1975. Learning How to Mean. London: Arnold.

Halliday, M.A.K., 1978. Language as a Social Semiotic; The social interpretation of language and meaning. London: Arnold.

Hamlyn, D., 1978. Experience and the Growth of Understanding. London: Routledge and Kegan Paul.

Hebb, D.O., 1955. 'Drives and the C.N.S. (Conceptual Nervous System)'. In: ed. D. Bindra and J. Stewart.

James, W., 1890. Principles of Psychology. London: Macmillan.

Josephson, B.D. and Ramachandran, V.S. (eds.), 1980. Consciousness and the Physical World. Oxford: Pergamon Press.

Joyce, J., 1922. Ulysses. Paris: Egoist Press. Also in Penguin.

Kahneman, D., 1973. Attention and Effort. New Jersey: Prentice-Hall Inc.

Keenan, E. and Schieffelin, B., 1976. 'Topic as a discourse notion: a study of topic in the conversation of children and adults'. In: C.N. Li (ed.) Subject and Topic. New York: Academic Press.

Kempson, R.M., 1975. Presupposition and the Delimitation of Semantics. Cambridge: C.U.P.

Kempson, R.M., 1977. Semantic Theory. Cambridge: C.U.P.

Kowalski, R., 1985. Inaugural lecture, Imperial College, London.

Kripke, S., 1972. 'Naming and necessity'. In: D. Davidson and G. Harman (eds.), Semantics of Natural Languages. NJ: Atlantic Highlands.

Labov, W., and Fanshel, D., 1977. Therapeutic Discourse. New York: Academic Press.

Lang, P.J., Miller, G.A. and Levin, D.N., 1983. 'Anxiety and fear: central processing and peripheral physiology'. In: R.J. Davidson, G.E. Schwarz and D. Shapiro (eds.).

Levinson, S., 1983. Pragmatics. Cambridge: C.U.P.

Lewis, D., 1969. Convention: A Philosophical Study. Cambridge MA: Harvard University Press.

Loftus, E.F., 1973. 'Activation of semantic memory'. American Journal of Psychology 86: 331-7.

Luria, A.R., 1975. The Working Brain. Harmondsworth: Penguin.

Lyons, John, 1977. Semantics 1 and 2. Cambridge: C.U.P.

Meichenbaum, D., 1976. 'Toward a cognitive theory of self-control'. In: G.E. Schwarz and D. Shapiro (eds.).

Merleau-Ponty, M., 1962. The Phenomenology of Perception. London: Routledge and Kegan Paul.

Minsky, M., 1975. 'A framework for representing knowledge'. In: P.H. Winston (ed.), The Psychology of Computer Vision. New York: McGraw-Hill.

Moruzzi, G. and Magoun, H.W., 1949. 'Brain stem reticular formation and activation of the EEG'. EEG clin. Neurophysiol. 1, 455-73.

Murray, D.K.C., 1979. 'Mis-Words'. In: Working Papers of the London Psycholinguistic Research Group 1: 60-68.

Murray, D.K.C., 1983. 'Conversational concerns: issues'. Journal of Pragmatics 7: 1-15.

Neisser, U., 1963. Cognitive Psychology. New York: Appleton-Century-Crofts.

Neisser, U., 1976. Cognition and Reality. San Francisco: W.H. Freeman.

Nietzsche, F., 1967. (1901-1906). The Will to Power, trans. Kaufmann, W. New York: Random House.

Nissen, H.W., 1953. 'Instinct as seen by a psychologist'. In: ed. D. Bindra and J. Stewart.

Norman, D.A., 1981. 'Categorization of action slips'. Psychological Review 88: 1-15.

Polanyi, M., 1960. Personal Knowledge. Chicago: Univ. of Chicago Press.

Pope, K.S., and Singer, J.L., 1978. The Stream of Consciousness. New York: John Wiley and Sons.

Pribram, K.H., 1976. 'Self-consciousness and intentionality'. In: G.E. Schwarz and D. Shapino (eds.).

Prince, E., 1981. 'Towards a taxonomy of given-new information'. In: P. Cole (ed.).

Psathas, G. (ed.) 1979. Everyday Language: Studies in ethnomethodology. New York: Irvington.

Putnam, H., 1975. 'The 'meaning'of "meaning"'. In H. Putnam, 1975b:

Putnam, H., 1975b. Mind, Language and Reality. London and New York: C.U.P.

Putnam, H., 1978. Meaning and the Moral Sciences. London: Routledge and Kegan Paul.

Reichman, R., 1978. Conversational Coherency. Cambridge MA: Bolt Beranek and Newman.

Reinhart, T., 1981. 'Pragmatics and linguistics: an analysis of sentence topics'.

Rommetveit, R., 1974. On Message Structure. London: John Wiley and Sons.

Pycroft, C., 1968. A critical Dictionary of Psychoanalysis. Harmondsworth: Penguin.

Rycroft, C., 1979. The Innocence of Dreams. Oxford: Hogarth Press.

Sacks, H. and Schegloff, E.A., 1979. 'Two preferences in the organisation of reference to persons in conversation, and their interaction'. In: G. Psathas (ed.).

Sacks, H., Schegloff, E.A. and Jefferson, G., 1974. 'A simplest systematics for the organization of turn-taking for conversation'. Language 50: 696-735.

Schank, R., 1977. 'Rules and topics in conversation'. Cognitive Science 1: [illegible].

Schank, R., 1978. 'Interestingness: controlling inferences'. New Haven: Yale University Department of Computer Science, Research Report 145.

Schank, R., and Abelson, R.P., 1977. Scripts, Plans, Goals and Understanding: an inquiry into human knowledge structures. [illegible]

Schegloff, E., 1968. 'Sequencing in conversational openings'. American Anthropology [illegible].

Schegloff, E., 1972. 'Notes on a conversational practice: formulating place'. In: D. Sudnow (ed.). Studies in Social Interaction. New York: Free Press.

Schegloff, E., 1977. 'Some questions and ambiguities in conversation'. In: G. Psathas (ed.).

Schiffer, S., 1972. Meaning. Oxford: Clarendon Press.

Schwarz, F. and Rouse, R.O., [illegible]. 'The [illegible] and recovery of asso ciations'. [illegible] Monograph 9, IUP.

Schwarz, F. and Schiller, P.H., [illegible]. 'A psychoanalytic model of attention and learning'. [illegible] Monograph 23. IUP.

Shotter, J., [illegible]. [illegible] games, and alternative [illegible]'. In: [illegible] (ed.).

Singer, J.L., 1966. Daydreaming. New York: Random House.

Sokolov, A.N., 1972. Inner Speech and Thought. New York: Plennon Press.

Sperber, D., 1982. 'Apparently irrational beliefs'. In: S. Lukes and M. Hollis (eds.), Rationality and Relativism. Oxford: Blackwell.

Sperber, D. and Wilson, D., 1980 (1982). 'Mutual knowledge and relevance'. In: N.V. Smith (ed.), Mutual Knowledge.

Sperber, D. and Wilson, D., 1981. 'Irony and the use-mention distinction'. In: P. Cole (ed.).

Sperber, D. and Wilson, D., 1984. 'On defining relevance'. In: R. Grandy, and R. Warner (eds.) Philosophical Grounds of Rationality: Intensions, categories, ends.

Sperber, D. and Wilson, D., 1985a = MSa. 'On choosing the context for utterance interpretation'. In: J. Allwood and E. Hjelmquist (eds.), Foregrounding Background. Doxa: Lund.

Sperber, D. and Wilson D., 1985 = MSb. Relevance. Oxford: Blackwell.

Stern, G., 1932. Meaning and Change of Meaning. Goteborgs: Hogskolas Arsskrift XXXVIII.

Stich, S.P., 1983. From Folk Psychology to Cognitive Science. Cambridge, Mass.: Bradford Books, MIT.

Titchener, E.B. 1929. Systematic Psychology: Prolegomena. Ithaca: Cornell U. P..

Tyler, S.P., 1978. The Said and the Unsaid. New York: Academic Press.

Underwood, G. (ed.), 1982. Aspects of Consciousness, Vol 3: Awareness and Self-Awareness. New York: Academic Press.

Underwood, G., 1982b. 'Attention and awareness'. In: G. Underwood (ed.).

Vygotsky, L.S., 1962. Thought and Language. New York: MIT and John Wiley and Sons.

Wingfield, A. and Byrnes, D.L., 1981. The Psychology of Human Memory. New York: Academic Press.

Wittgenstein, L., 1958. Philosophical Investigations. Oxford: Blackwell.