Part I

Introduction

1

Introduction

1.1 The Philosophy of Language

The philosophy of language raises issues both important and difficult.

The importance of language to human life is obvious. All human societies are language using, as are all their more or less normal members. Language acquisition is one of the few cognitive skills that is, near enough, both common and peculiar to humans. This skill gives the human species an enormous advantage over others: language is a quick and painless way of passing on the discoveries of one generation to the next. Some theorists see language as *the* most central characteristic of the human species.

The obvious importance of language makes its study worthwhile, but does not prepare one for the ascendancy that this study has achieved in philosophy. Over the last century or so the philosophy of language has become the pivotal area of philosophy, particularly within the English-speaking tradition. Opinions about language have been thought to settle traditional philosophical problems in epistemology, metaphysics, and ethics. The very nature of philosophy itself has been linked by many to the study of language. In our view, much of this goes too far: philosophy of language has become too big for its boots. We shall return to this often (Parts IV and V).

The philosophy of language is certainly difficult. In part this is because of our closeness to language: we find it hard to get a proper perspective on it. In any case, there is a vigorous, wide-ranging, and bewildering debate

on foundational and conceptual issues in the philosophy of language. Many competing theories are on offer, yet it is often hard to see what these are theories *of*. What problems are they trying to solve? What counts as *evidence* for or against a theory? Are different theories really concerned with the same problems, so that they are really in competition? If not – and it often seems not – how do they relate to one another? Finally, the status of philosophical theories of language is obscure and controversial, as we shall begin to see in section 1.3.

You may wonder why language is studied *in philosophy* at all. Why is the study not left to linguistics which is, after all, "the science of language"? The three main branches of linguistics are phonology, concerned with *sounds*; syntax, concerned with *grammatical structure*; and semantics, concerned with *meaning*. It is in semantics, or "the theory of meaning", that the theoretical and conceptual chaos mentioned above is most striking. There are also some similar problems in syntax. Philosophy is typically concerned with the most intractable and conceptually difficult parts of various disciplines. So it is not surprising that it is deeply concerned with semantics and has some concern with syntax. In contrast, it is not concerned with phonology at all.

Finally, a word of warning. What follows is not an auctioneer's catalogue of the theories and ideas in this controversial subject. It cannot therefore be either neutral or comprehensive. Our approach reflects our ideas of what is central and peripheral; of the main issues and of the blind and side alleys. The approach is not idiosyncratic, but it is not shared by all, or even most, philosophers.

1.2 What Is the Problem?

What is the main problem in the study of language? What are the phenomena that pose the problem? To get an appropriate distance from the phenomena, it helps to imagine that we are Martians visiting Earth for the first time. What linguistic phenomena strike us? We observe humans producing sounds and inscriptions: they are talking and writing to each other. We see that these activities play extraordinarily important roles in human life. We wonder what properties the sounds and inscriptions have that enable them to play these roles. We decide to call these properties "meanings". The problem then is to describe and explain these meanings: to say what they are and to say what makes it the case that something has one. In brief, the problem is to give a theory of meaning or a semantics.

This problem, in turn, gives rise to a second, psychological, problem. What features of the human mind make it possible for humans to produce, and react, to these sounds appropriately? How do we manage to use

language? So, as we shall see, the problem of explaining meaning is linked to the problem of explaining linguistic competence and understanding.

We have identified the semantic problem by referring to the "important roles" of language in human life. What roles are these? We think that the folk saying, "language expresses thought", captures the central role: a human language is *a system for expressing or communicating thought*. Many thoughts are informational, being about the social or physical environment. But thoughts need not be informational, or purely so: language is used to greet, question, command, joke, offend, abuse, intimidate, and so on.

To say more about the importance of language we have to consider the roles of the thoughts that language expresses. There are two roles that are obviously of interest to us. First, thoughts cause people to do things: a person who believes that it is raining is likely to take her umbrella. So, knowing what others are thinking tells us *what they are likely to do*. We depend intimately on others, so this knowledge is important to us. Language is the main way we gain that knowledge. Second, many thoughts carry information or misinformation about the world. So, knowing what a person thinks often *tells us about the world*: if the person who believes that it is raining is reliable about such matters we can learn from her about the weather without looking outside. Much of our knowledge of the world comes from the linguistic expression of beliefs. This is why language gives us the great advantage over other species noted at the beginning of the chapter: it enables each of us to benefit richly from others' experience.

In sum, the central role of language is to express thoughts. Derivatively, it has at least these two roles: explaining behavior and informing us about the world. Meanings are the properties that enable it to play these roles.

We do not pretend that these brief remarks are sufficient to identify the semantic problem. Implicitly, the problem concerns the "natural" languages of humans; for example, English or Swahili. Such languages are not our only communication systems; think of flag signals or "body-language", for example. And many animals clearly have communication systems. Doubtless all of these systems can be used to explain behavior and learn about the world. So, what is special about our natural languages?

Consider also "animal language" experiments. Alex, an African grey parrot trained by Irene Pepperberg, produces intelligible and contextually appropriate English-like sounds. If one shows Alex a red plastic object and asks "what color?" Alex says "red". If one shows Alex a wooden peg, in response to a parallel question about composition, Alex says "wood". Dolphins can respond to sometimes quite complex and novel verbal commands. Most famously of all, a number of primate species

have learned to interact with their trainers through gesture or the use of arbitrary plastic tokens. All these "utterances" could be used to explain behavior and inform us about reality. Yet controversy rages as to whether these animals are really speaking a human language. What might the animals be doing, or failing to do, that would settle the matter? What is it to use "red" as an English word rather than, say, just as a means of extracting reward from a trainer?

We think that progess with these questions, hence with identifying our problem, is chiefly to be made by looking more closely at the relation between human language and thought. We shall do so later (7.4). Meanwhile, we make some progress with identifying our problem by setting out some of the salient features that make human natural languages and their meanings so special:

a. Stimulus independent

In most circumstances, as full a description as you like of a person's physical environment does not enable you to predict her next utterance. The contrast with animal communication systems, for example, is notable.

Briefly, animal communication systems seem to be of two sorts. First, birds (and apparently nonhuman primates) have a fixed and fairly small repertoire of distinct signals, each of which has a set function: flight call, alarm call, and the like. A particular environment elicits the appropriate response. Human language does not consist in such a small fixed repertoire of predictable responses. Second, consider bees. A bee returning from a distant food source dances a message. The positioning of the dance and its pattern indicate the direction and distance of the food source. This remarkably efficient system of communication differs from those of birds in having an unlimited number of signals: the length and the pattern are capable of indefinitely many variations. Nevertheless, the bee's system is not flexible in the way human language is. Each response is environmentally fixed: if you know where the bee has been, and if you know the coding system, you can predict the pattern of the dance. In contrast, if a person comes from a food source – a good restaurant, for example – you cannot predict her words. Her food description – indeed, whether she talks about food at all – is stimulus independent.

b. Abstract

A sentence may abstract from many details of a situation, focusing on just one. Thus, 'Orson weighs 130 kg' tells you nothing about Orson other than his mass. Symbols in many other systems cannot be quite so abstract; a photograph or sketch of Orson will tell of many of his properties.

Even the bee's dance cannot be silent on the distance of the food source while revealing its direction.

c. Arbitrary

In general, linguistic symbols have no intrinsic or necessary connection with their referents. The inscription, 'Ronald Reagan', happens to refer to a certain president of the USA, yet it is in an important sense arbitrary that it does so. That inscription could have been used to refer to Bob Hope; and Reagan could have been called 'Hopalong Cassidy'. This arbitrariness is nicely illustrated by the English logician, Charles Dodgson, better known as Lewis Carroll:

> " . . . – and that shows that there are three hundred and sixty-four days when you might get un-birthday presents – "
>
> "Certainly," said Alice.
>
> "And only *one* for birthday presents you know. There's glory for you!"
>
> "I don't know what you mean by 'glory,'" Alice said.
>
> Humpty Dumpty smiled contemptuously. "Of course you don't – till I tell you. I meant 'there's a nice knock-down argument for you!'"
>
> "But 'glory' doesn't mean 'a nice knock-down argument,' " Alice objected.
>
> "When *I* use a word," Humpty Dumpty said, in rather a scornful tone, "it means just what I choose it to mean neither more nor less."
>
> "The question is," said Alice, "whether you *can* make words mean so many different things."
>
> "The question is," said Humpty Dumpty, "which is to be master – that's all." (Carroll 1962: 247)

What Humpty Dumpty is emphasizing is the arbitrariness of language: we can call anything anything. (This is not to say that Alice does not have a point too. This could be brought out using the distinction between speaker meaning and conventional meaning: section 7.4.)

d. Learned

Many animal communications systems are complex mixes of the learned and the innate. The same is probably true of human language. Noam Chomksy has famously argued that the most important features of human syntax are innate (8.10). Yet it is uncontroversial that the large vocabulary of every human language must be learned. So the amount of learning involved in mastering any human language greatly exceeds the learning involved in nonhuman systems of communication.

e. Medium independent

Linguistic communication can be effected in speech, writing, braille, gesture, and so on. There seems to be no limit to the media we could use for, say, English. The birds and the bees are more limited.

f. Systematic

The matching of each signal with its meaning is not something that we learn signal by signal. We learn the elements of signals – words – together with a recipe for making complete signals – sentences – out of the elements. Thus your knowledge of a few words and the constructions of English enables you to understand 'Andropov liquidated the Hungarians' and 'The Hungarians liquidated Andropov' even though you may never have come across these sentences before. (The bee's system has *some* systematicity but note how simple its generating recipe is.)

g. Power

One very special thing about language is its power and versatility. It can serve the many purposes of communication. It enables us to deal with the past and the future, the present and the absent. We talk of an enormous range of topics: of "tables, people, molecules, light rays, retinas, air waves, prime numbers, infinite classes, joy and sorrow, good and evil" (Quine 1966: 215). Contrast this with the bee's monomania.

It is clear that language gets its power from many of the other special features we have mentioned.

We can use the arbitrary symbol 'Mandela' to convey information, or misinformation, about a politician who is thousands of miles away. We can use the arbitrary symbol 'Thales' to convey information, or misinformation, about a Greek philosopher who has been dead for two millennia. Linguistic symbols have properties, "meanings", that make these feats possible. These meanings enable the symbols to play extraordinarily important roles in human life, particularly roles in explaining and predicting behavior and informing us about the world. Symbols have their meanings in the context of language: a uniquely powerful communication system that is stimulus- and medium-independent, abstract, arbitrary, learned, and productive. We think philosophy of language is confronted with two related main problems. One is to describe and explain the properties of symbols in virtue of which they play the central role they do in our lives; we call this the problem of explaining meaning. The related problem is to describe and explain linguistic competence; the features of people's minds in virtue of which they

can use and understand symbols and the symbol system we call language.

In Part II, our focus will be on meaning. We follow custom in identifying meanings with the properties that make linguistic symbols important to us, but that should not engender the idea that there is general agreement about the nature of the problems. The ordinary use of the term 'meaning' is loose, and theoretical uses are varied. Partly as a result of this, there is obscurity and controversy over the very problems in the philosophy of language. This is the first sign of the difficulties mentioned in the opening section. In Part II we will, so far as possible, set aside issues to do with the mind's relation to language. However, it will prove impossible to ignore the issue of competence. In Part III, the mind's relation to language is at centre stage.

1.3 What Is a Theory of Language?

There is obscurity and controversy not only over the problems for which we need theories of language but also over the status of the theories themselves. This issue of status is highly abstract: it requires a theory of theories of language, a "meta-theory". It would be nice to ignore the meta-theory and get on with the theory, but that is a luxury we cannot afford. We think that many mistakes in the theory of language arise from a mistaken meta-theory. Further, we think that these mistakes are often facilitated by a failure to be explicit about the meta-theory: once the implicit meta-theory is exposed, it can be seen to be implausible and unsupportable. So we shall start by laying our cards on the table, with little in the way of a defense. We shall offer some sort of defense in Part V. We shall return to meta-theoretical issues often in our criticism of other views.

Our approach to the problem of language is *naturalistic* in two respects. The first respect is *epistemological* (concerned with our way of knowing). A theory of linguistic phenomena has just the same status as a theory of any other phenomena: it is empirical and conjectural; it is known *a posteriori* ("justified by experience") not *a priori* ("justified independently of experience"). We are confronted with a mysterious and complex world and have developed theories to explain and render tractable these complexities: theories in physics, biology, the social sciences and the like. The theory of language is just another such theory; another part of our total theory of nature.

One can be misled into treating the theory of language differently by the fact that, at bottom, much of what it has to say at present is "common sense". People divide linguistic phenomena into sentences and words; they divide words into nouns, verbs, etc. They think that expressions are meaningful and have meanings. They think that words refer to parts of the

world. They think that sentences express thoughts; that some are true and some false, but none are both; that some are questions and others commands; and so on. These gems of common sense, taken over by linguistics, may seem to have some special status. It is easy to succumb to the illusion that whereas science, especially its abstract and speculative branches like cosmology and particle physics, is conjectural, empirical, and fallible, common sense is not. So, it is thought to be a *theory* that the earth and moon are linked by gravitational force, but a *fact* that sentences have meanings. This illusion is engendered by the familiarity of common sense.

Common sense is best seen as a mix of *folk theories* or, if talk of theories seems too pretentious here, *folk opinions*. These, like scientific theories, help people better understand and explain the phenomena that confront them. So the above gems are best seen as articulations of folk linguistics' response to features of language of the kind discussed in the last section. Folk theories differ from scientific ones in being immature: they are less precise, systematic, and explicit; they lack a methodology for development. More seriously, they differ in being believed uncritically.

A glance at the past shows that folk theory has no special warrant. Early European folk geography, folk meteorology, and folk medicine have been comprehensively rejected. We no longer think that the Earth is flat, nor that the winds are under the direction of supernatural agents. We no longer explain health in terms of the humors of the blood. Nevertheless, where a folk theory is working well, just as where a scientific theory is working well, it is *unlikely* to be wholly wrong. So, it is reasonable to suppose, for instance, that folk psychology, with its long history of fairly successful use (we have *quite* a good understanding of one another), has a lot of truth in it. Similarly, perhaps, folk linguistics. Nonetheless, even the best folk theories stand in need of supplementation and revision.

The second respect in which our approach is naturalistic is *metaphysical* (concerned with what there is and with what it is like). How do facts about language relate to facts about people and to other facts about the world? The answer, in our view, is given by *physicalism*.

We think that people are best seen as part of the natural world. They are not special except in detail and complexity. (i) They are part of animate nature; part of the biological world. (ii) The biological differs from the inanimate only in complexity: no vital essence distinguishes the living from the nonliving. To be living is only to have a special, if complex, chemistry.

Physicalism is intrinsically plausible. It has excellent scientific support from evolutionary theory, biology, and biochemistry. These sciences underscore the biochemical and physiological continuities between humans and the rest of nature. There are, we believe, no good arguments *against* this perspective.

Our theory of language must, therefore, be physicalistic. Any linguistic

facts there are must be, ultimately, physical. Semantic notions like meaning, truth, and reference can be used only if they can be explained in nonlinguistic terms; they are not primitive. Biologists were not satisfied to leave the notion of gene as primitive: they wanted to understand the mechanism by which inheritable characteristics are encoded in a cell. Their search led to the discovery of the structure of DNA. Similarly, we seek a deeper explanation of semantic notions. We might, for example, hope to explain them in psychological terms; then, hope to explain the psychological in neuro-anatomical and biochemical terms; then, explain those in physical and chemical terms.

Some think this hope is a vain one. In particular, the famous Harvard philosopher, W. V. Quine, doubts that we can explain a robust notion of meaning physicalistically. This leads him into semantic "eliminativism": the familiar notions of folk semantics have no place in a developed theory of humans and their languages. We are not eliminativist, but we accept his moral: if notions like meaning and truth cannot be explained in nonsemantic, naturalistic terms, we should do without them in our theory of language.

We return to the defense of naturalism in Part V.

**1.4 The Menu

(Passages and references between asterisks (**) are difficult and could well be ignored in an initial approach.)

We finish this chapter with a preview.

Part II is centered on the problem of meaning, on explaining those properties of linguistic symbols that enable them to play their distinctive role in our lives. In chapter 2 we propose a "representationalist" view of meaning. On this view, the core of a sentence's meaning is its truth condition; that is, the property of a sentence which, together with the world, make it true or false. We suggest that the truth condition of a sentence depends on the referential properties of its elements together with its syntactic structure. So, in chapters 3 to 5 we discuss reference and in chapter 6 we discuss structure.

Description theories of reference may work for some terms but we argue that they fail for proper names and natural kind terms. We initially favored pure-causal theories for these but think that we may have to settle for hybrid descriptive-causal theories. We explore other possible theories of reference for a variety of terms. In discussing structure, we draw on "transformational generative grammar".

In Part III, our focus changes from symbols and their meaning to an intimately related but nonetheless distinct arena, the mind.

In chapter 7 we suggest that the relation between thought and language is somewhat symbiotic. On the one hand, we argue for the "language-of-thought" hypothesis, the view that thought is typically language-like in character. Indeed, there is probably a close relationship between the language a person thinks in and her public language. On the other hand, the development of that public language and the system of conventions it embodies must depend on the achievement of cognitive pioneers in having thoughts not then expressible in their public language.

Chapter 8 is devoted to the vexed and difficult problem of linguistic competence. We develop, to some extent, a view of competence, but the basic thrust of the chapter is critical. We reject certain intellectualist theories favored by linguists and philosophers. We are dubious of the idea that grammatical rules play a significant role in language processing. We look critically at the view that some grammatical rules or principles are innate.

Chapter 9 briefly defends our sort of representationalism against rivals. One rival is the view that the meaning of an expression should be explained in terms of its "functional role" in the mind. Another is a "two-factor" theory that explains meaning partly in terms of functional role and partly in terms of representational properties. Finally, we consider the skeptical argument against representational meaning that Saul Kripke has found in the work of Wittgenstein.

In Chapter 10 we take a decidely critical view of linguistic relativism: the view that, (i) your general picture of the world is influenced and constrained by the language you speak; and that (ii) languages differ enough to produce incommensurable world views. We concede a little to this view but, in general, find it exciting only when false.

Part IV considers the relation between theories of language and the metaphysical doctrine of realism. Commonsense realism is the view that the ordinary furniture of our environment – cats, trees, stones, etc. – exist independently of us and our thoughts on the matter. Scientific realism takes a similar view of the objects of science. Many theorists have inferred anti-realist views from their favored theories of language. In some usually ill-specified sense, the world is said to depend for its existence or nature on us.

We consider several of these theorists, taking a mostly critical view of their theories of language, and a very critical view of the metaphysical views they derive from them. However, our main point is that a realist metaphysics has more secure epistemic foundations than *any* theory of language. So, the appropriate strategy is to construct theories of language from that perspective, not to construct metaphysical views from the perspective of one's favorite theory of language. If a theory of language contradicts our best overall picture of the world, so much the worse for the theory.

In chapter 11 we consider the verificationism of the logical positivists

and Michael Dummett. In chapter 12 we consider "constructivism", the view that, one way or another, different groups make different worlds by imposing their views. We start with Benjamin Lee Whorf, who thinks that we construct realities with languages. We go on to the radical philosophers of science, who think that we do the job with scientific theories. Finally, we consider the former realist, Hilary Putnam. Chapter 13 is devoted to structuralism (or semiotics, as it is frequently known). This movement, centered in France, is an extreme form of constructivism.

The book ends, in Part V, with a discussion of philosophy itself. Considerations about language have dominated much of twentieth-century Anglo-American philosophy. In chapter 14 we recount that dominance, tentatively diagnose it, and reject it. Our stance is, as always, naturalistic. In chapter 15 we reject another challenge to naturalistic philosophy: rational psychology. That approach takes our ordinary views of people and their language – folk psychology and folk linguistics – to be outside science; they are thought to supply knowledge of a different sort altogether.**

Suggested Reading

1.1

On the centrality of language to the human species, and for its role in making us the creatures that we are, see: Bickerton 1991, *Language and Species*; Corballis 1991, *The Lopsided Ape*; Donald 1991, *Origins of the Modern Mind*; Lieberman 1991, *Uniquely Human*; Pinker 1994, *The Language Instinct*; Noble and Davidson 1996, *Human Evolution, Language and Mind*.

1.2

For a good survey of the animal communication literature from the standpoint of someone who thinks it shows a rich cognitive life, see: Griffin 1992, *Animal Minds*, especially chs. 8–11. For somewhat more measured views, see: Byrne 1995, *The Thinking Ape*, ch. 11; Roitblat and Meyer 1995, *Comparative Approaches to Cognitive Science*, Part V; and especially Hauser 1996, *The Evolution of Communication*. For a defense of the "language-likeness" of ape language, see: Savage-Rumbaugh 1986, *Ape Language*; Savage-Rumbaugh and Lewin 1994, *Kanzi* (1994). For a very skeptical view of the significance of these projects, see Pinker 1994, ch. 11.

For more on the nature of the semantic problem, see Devitt 1996, *Coming to Our Senses*, ch. 2.

1.3

In a series of works Patricia and Paul Churchland defend the possibility of massive revision of folk theory. They argue for eliminativism about the mind which, because of the links between mind and language to be explored later (Part III), is closely related to semantic eliminativism. See Paul Churchland 1993, "Evaluating Our Self Conception", for a nice summary of their position; also his excellent introductory text, *Matter and Consciousness* (1988). Part IV of Lycan 1990, *Mind and Cognition* has some classic papers on eliminativism.

Many philosophers argue that eliminativism is *incoherent*. See Hannan 1993, "Don't Stop Believing" for a sympathetic presentation of the argument. For criticisms see Sterelny 1993, "Refuting Eliminativism on the Cheap?", and Devitt 1996: 249–52.

Explanation of one set of facts in terms of another often involves quite complex relations between the two. See Fodor 1975, *The Language of Thought*, introduction. Boyd, Gasper, and Trout 1991, *The Philosophy of Science*, section III, is a good selection of papers on this issue.

Quine's classic assault on meaning and its relatives is "Two Dogmas of Empiricism", in *From a Logical Point of View* (1961). **This eliminativism was continued in *Word and Object* (1960) with his famous, but difficult, argument for the indeterminacy of translation: ch. 2.** When reading Quine it should be kept in mind that a behavioristic conception of language constrains his views. See Gibson 1982, *The Philosophy of W. V. O. Quine* for more on Quine's behaviorism. The most readable account of his position is in *Philosophy of Logic* (1970), ch. 1. A sympathetic and interesting examination of Quine's views is to be found in Romanos 1983, *Quine and Analytic Philosophy*. Antony 1987, "Naturalized Epistemology and the Study of Language", is a nice critical piece.

Just as linguistics should be naturalized so also should epistemology. The quotation from Quine in section 1.2 is from the opening paragraph of a marvellous essay arguing for a naturalized epistemology, "The Scope and Language of Science" in *The Ways of Paradox* (1966). Kornblith 1994, *Naturalizing Epistemology* is a very helpful collection.

1.4

There are a number of anthologies on philosophy of language, to which we shall often refer in "Suggested Reading". Two very helpful general ones are Martinich 1996, *The Philosophy of Language*, and Ludlow 1997, *Readings in the Philosophy of Language*.

Schwartz 1977, *Naming, Necessity, and Natural Kinds*, is a good collection focusing on the problems of reference. Davis 1991, *Pragmatics*,

focuses on the pragmatic aspects of language. Searle 1971, *The Philosophy of Language*, attends to this too but its main focus is on the implications of contemporary linguistics.

Block 1981, *Readings in the Philosophy of Psychology, Volume 2*, is an excellent collection on issues raised in Part III. It is organized into parts, each of which has a helpful introduction. Stich and Warfield 1994, *Mental Representation*, is also helpful on Part III issues. It has a good selection of papers offering naturalistic theories of meaning. Lycan 1990 contains a number of papers that bear on Part III issues. So also do parts IV and V of Rosenthal 1991, *The Nature of Mind*.

Geirsson and Losonsky 1996, *Readings in Language and Mind*, brings out the close ties between theories of language and mind. It is an interdisciplinary collection, reflecting the rise of cognitive science.