Introduction

Philosophy, as Gilbert Ryle[1] noted long ago, deals characteristically in dilemmas: their exploration and (sometimes) their resolution. Ryle was clearly right. Philosophical puzzlement very often originates in a question – which for some reason seems to us momentous – either answer to which commits us to unpalatable or implausible consequences.

So it is with the question whether "language" in the abstract, language taken as a semantic order, a system of meanings, "mirrors the world": whether the categories, concepts, structures with which it furnishes us, far from being inventions of the human mind, simply transcribe categories and structures already inscribed in Nature, or Reality. If we answer "yes," we surely discount, or at least minimise to an implausible degree, the part played by human ingenuity in the constitution of meaning in actual languages. If we answer "no," by contrast, we seem to be denying the possibility of truth and objectivity. For how are we to describe anything truly, if the terms in which language forces us to frame all that can be said are set, not by the nature of what is to be described, but by linguistic or social convention?

The dilemma is a characteristically philosophical one; one, certainly, which has occasioned the spilling of much ink by philosophers. But its implications transcend the bounds of philosophy, at least philosophy narrowly considered as what goes on in philosophy departments. In linguistics, literary studies and the social sciences, many of the debates of the past thirty years have turned on the issue of the "referentiality," or otherwise, of language. A range of influential writers, including Derrida, Saussure, Lévi-Strauss, Roland Barthes, Julia Kristeva, among many others, have argued that the constitution of meaning within a language is neither constrained nor validated by anything external to language. Those opposed to these developments have tended to see them as promoting forms of relativism hostile to the very possibility of objective truth.

In this book we shall opt for neither wing of this ramifying and occasionally acerbic debate. Instead we shall argue that the debate itself is

misconceived, because the choice, between relativism and referentiality, which appears to most of its participants to exhaust the available options, does not in fact exhaust them. There is, we shall show, a way of understanding the constitution of meaning in natural language that will allow one both to deny the existence of any extralinguistic correlate of meaning and yet, perfectly consistently, to affirm the possibility of objective truth.

The path we shall pursue falls half within and half outside the familiar terrain explored by analytic philosophy since Frege. Because our conclusions are, if correct, rather surprising ones, it is perhaps hardly surprising that the argument should be not only long and complex, but at times quite unfamiliar, not only in a number of its crucial moves, but even in much of its detail. We have tried to address the difficulty by dividing the stages of the argument between short, numbered and subheaded, and we hope reasonably clearly related, sections of text, easily relocatable by reference to an unusually detailed table of contents. A preliminary map of the stages of the argument, with some indication, however rough and preliminary, of their content and interrelationships, may nevertheless prove helpful. To that we now turn, with the proviso that what is offered in the next few pages is a bare and highly schematic outline of the argument, leaving out, along with most of its detail, most of what might make it – we trust – persuasive.

The argument of the book has a main thread running from beginning to end, to which are attached, at various points, a number of essential but subordinate discussions. The business of the main argument is the refutation, and replacement, of the doctrine introduced in Chapter 2 under the label *Referential Realism*. The Referential Realist holds that meaning is introduced into a language by the association of some class of meaning-bearing elements of the language with some class of real-world entities whose existence and nature owe nothing to linguistic convention. The case for Referential Realism, a powerful and enduring one, is developed in Chapters 1–2 and §i of Chapter 4. It rests essentially with the thought that unless the members of some class of elements of language derive their meaning simply from association with the members of some class of elements of "the world," language becomes hermetically self-referential, a prison made for itself by the mind, rather than the means of articulating thoughts concerning a mind-independent reality.

In Chapter 4 §i, we develop what is in effect a *reductio* argument against Referential Realism. It follows from Referential Realism, we argue, that, in general, we can know whether a string of words expresses a thought, in the Fregean sense of a content capable of being assessed for its truth or falsity, only if we know some other proposition to be true: for example, the proposition that each of the names in the proposition possesses, "out there in the world," a bearer. But we cannot set about assessing the truth of any proposition until we know what it asserts, as until we know that, to put it bluntly, there is nothing to submit to such assessment. It follows that all

questions concerning the assertoric content of propositions must be settled in advance of raising the question whether any proposition is true or not, because in advance of those questions being settled, there is nothing about which to raise the question. And from that it follows that Referential Realism, as it entails the contrary, must be false.

This argument can be exhumed, with minimal exegetical effort, from discussions between Russell and Wittgenstein during the period from Wittgenstein's first meeting with Russell in 1911 to the composition of what was to become the *Tractatus Logico-Philosophicus*. It is, we suggest, the argument that gives force to a remark of Wittgenstein's that we label *Wittgenstein's Slogan* and that constitutes the leitmotiv of the *Tractatus* and the *Notebooks 1914–16:* "Logic must take care of itself" (*Die Logic muss für sich selber sorgen*). Coming at the argument by this route, given that a majority of philosophers now regard the *Tractatus* and its arguments as of purely historical interest, will to many eyes give Chapter 4 a quaint air of philosophical palaeontology. Can these antique speculations have much bearing on more recent writers, such as Kripke, Davidson, Dummett, or Gareth Evans?

One answer is that introducing the argument in its historical context renders not only its provenance but also its motivation and implications considerably clearer. Another is that contemporary philosophy of language is sufficiently Russellian in its assumptions, and even in much of its content, for Wittgenstein's early dissenting voice to have, as we shall see as the argument develops in detail, a sharper resonance today than might at first sight seem likely.

At first sight Wittgenstein's argument might seem indeed to lead into a dead end. On the one hand, it seems perfectly sound. On the other, it is difficult to see how its conclusion can be correct. If the constitution of linguistic meaning must logically precede the establishment of any contingent truth about the world, even the truth that a given name has a bearer, it is hard to see how language can get off the ground. Without some connection with reality, it would seem, language can be nothing but an hermetic game played with contentless counters; but how could reality enter the process of meaning-constitution, except by way of our grasp of some body of contingent truths, if only the truth that the noise "Mama" designates Mama?

The goal of the book is, in effect, to answer this question: to construct an account of meaning in natural language in tune with the implications of Wittgenstein's Slogan. This enterprise proceeds in three stages, the first roughly coextensive with Chapter 3, the second with Chapters 5–7, the third with Chapters 9–12. The first of these sections proposes an outline solution. It is, in effect, that we stop attempting to represent "the relationship between language and the world" as a relationship between *meaning-bearing elements of language and some class of entities envisaged as corresponding elements of the world.* The alternative proposed is that we think of the relationship as a two-stage one, in which world and meaning-bearing elements of language are related

to one another not directly, but only via their relationship to the third constituent of the relationship: practices. This move, in effect, separates two questions commonly supposed conterminous: "How is language related to the world?," and "How do linguistic expressions acquire meaning?" Once these questions are seen in this way to be distinct, we are in a position to avoid much philosophical muddle arising from the attempt to answer them as if they merely expressed different aspects or versions of one and the same question. The answer to the second, we suggest, is that linguistic expressions acquire meaning through their involvement in a wide variety of practices. The answer to the first is that the practices through which linguistic expressions acquire meaning are not, for the most part, practices of symbol-manipulation (although some, card-games for instance, are). For the most part their point, and their utility in our lives, stems from the fact that they involve (as, for instance, the practices of measurement, or the recording of music in terms of the tonic scale involve) the manipulation of things not constituted by human convention: actually existing elements of the sensory field. It follows that the meaning-bearing elements of a natural language cannot be said, in their content and structure, to "mirror the world." If their content and structure "mirror" anything, it is the content and structure of the practices through involvement in which they have acquired whatever meaning we have bestowed on them. But it does not follow from that, that language is an hermetic play of signs, for the practices in which linguistic signs participate are not (or not all), as Locke would say, themselves "occupied about" signs, but about real things: things "real" in the sense of things existing prior to, and independently of, human convention.

For this outline solution to stand as a tenable account of meaning in the spirit of Wittgenstein's Slogan, however, it needs to be shown in detail how it can be developed to account for "meaning" of at least the following two kinds: on the one hand, the kind that consists in the relationship of a proper name to its bearer, and, on the other, the kind ("sentential meaning," "assertoric content") that renders at least some sentential signs fit to be assessed for truth or falsity. The first of these issues occupies Chapters 5–7, the second Chapters 9–12.

Philosophical dispute about proper names has addressed two closely connected questions. The first concerns the issue of what it is for a proper name to possess a meaning. Are we to say that a proper name has a Fregean sense, usually equated with an identifying description, or are we to say that it is a purely – or "directly" – referring expression, whose meaning (Fregean *Bedeutung*) is to be identified with its bearer? The second question concerns the conditions that have to be met in order that a speaker may be said to be in a position to refer by means of a proper name. To both we offer answers that fall outside the range of options offered by contemporary debate.

On the second question, there is agreement that a speaker cannot be in a position to refer by means of a proper name unless he or she stands

in some special relationship to its bearer. Opinions differ, however, as to the nature of the required relationship. One body of opinion holds the necessary relationship to be forged internally to the mind, and to consist in the possession by each competent speaker of an identifying description, or something of the sort. Another holds, with Kripke, that the connection is forged externally to the mind, and consists in the existence of a causal chain linking present uses of the name to past ones, and ultimately back to an original baptism.

The view presented in Chapters 5–7 might be seen from one perspective as a version of externalism, but one that appeals not to a causal history, but to the integrity of a system of practices, namely all those practices that involve the painstaking recording of proper names for an indefinite variety of purposes from registers of birth and deaths to library catalogues, records of shipping, maps, legal documents, and so on. Such practices, we suggest, make up an intricately crossed-referencing matrix through which individuals of many kinds, including human beings, ships, townships, farms, and so on, may be traced or tracked by the traces left by recorded uses of their names. We call this matrix, created by the observance of a multitude of social practices involving names, the *Name-Tracking Network*. The path from a name to its bearer is traced through this network, much as, in Kripke's account, it is traced back down a causal chain of uses to an original baptism; and as either route proceeds externally to the mind of any individual speaker, our view could well be seen as, in a similar sense to Kripke's, "externalist" in character.

But in another way, however, we might be supposed to hold a version of internalism. A speaker's knowledge of his own language will include, it is to be supposed, familiarity with some large subset of the practices that enter into the constitution and maintenance of the Name-Tracking Network. So he will be in a position to infer, merely from the occurrence of a name in a context appropriate to one or more of them, to the actual existence of a bearer of that name, even though he or she not only lacks an *identifying* description of that individual but also any description *of that individual* whatsoever! At the same time, a description of the circumstances of occurrence of a name is a *description*, even if it isn't a description of the individual denoted by the name! So we are proposing, it might appear, an account of the conditions for successful reference by means of a proper name that is, absurdly, both "internalist" and "externalist" in character, albeit in odd senses of those terms!

In fact, our position is neither "externalist" nor "internalist" in the usual senses attached to these terms in current theorizing, as the introduction of the notion of a Name-Tracking Network brings about a crucial shift in the terms of the discussion. It does so by allowing us to dispense with the founding assumption of current discussion, noted earlier, namely, the assumption that a speaker cannot be in a position to refer by means of a proper name

unless he stands in some relationship (whose precise characterisation provides the main matter of subsequent debate) to its bearer. According to us, the conditions for reference can be met just in case a speaker stands in the sort of relationship to the Name-Tracking Network that consists in familiarity with some large subset of its practices, augmented by knowledge of some set of circumstances of occurrence of the name in question. He or she does not need to stand in any relationship to the bearer of the name, either via a description or via a causal chain, because the task of locating the bearer of the name is performed, relative to a given set of circumstances of name-use, by the Name-Tracking Network (and so, *a fortiori*, by a means external to the mind of the speaker).

Settling the question of the conditions for successful reference in this way allows us to take a related line on the question of the meaning of a proper name. Here also current debate rests on a shared assumption, namely, that for a name to have a meaning is also for it to stand in a certain relationship to its bearer, in this case the relationship of *reference*, in one sense of that term. Our counterclaim can, once again, be viewed from two aspects. On the one hand, we propose that a name has meaning, not because of an occult relationship linking it to some individual, but because it is *in use* as one of the verbal counters deployed in the process of conducting one or more of the practices that make up the Name-Tracking Network. On the other hand, this move fails to sever the link between name and bearer, because it merely transforms it from a *direct* link (one whose very directness renders it occult because *sui generis*) into an indirect (and hence naturalistically explicable) one, established as the resultant of the distinct and very different relationships in which, respectively, name and bearer stand to a fabric of practices. Of course the name is introduced to its role *as a name* through being *associated*, baptismally or otherwise, with some individual. But (and this is the essential point) the act of so associating it would not confer on it the status of a name – would be a mere empty ceremony – if it were not for the "background" provided by the web of socially instituted and maintained practices within which it will subsequently find a use. Absent those practices indeed, no sense could be attached to the notion of *the bearer of a name*, because it is only by reference to those practices that we can make clear what it is to *be* the bearer of a name. To what, then, can we say that a name such as "Odysseus" or "Saul Kripke" *refers*? *Of what*, exactly, is it the name? It is, we shall suggest, the name of a *name-bearership*: a role or status as defined relative to a set of practices. That role is, of course, occupied in the case of those and other names, by particular persons, but there is in each such case no *single* "relationship of reference" that *both* links a phonemic string to a person *and*, by so doing, constitutes it *as a name*. Rather, there is a double, or two-stage relationship: on the one hand, the conferring on the phonemic string of the status of a name through its involvement as a counter in some set of practices, on the other, the accession of an individual to the role of

bearer-of-a-name through the bestowal on him, in the context of those same practices, of that string *as* his name.

We thus have a theory of naming consistent with the demands of Wittgenstein's Slogan. It follows from it, contrary to the requirements of Referential Realism, that, in order to attach a meaning to a statement, it is not necessary to know any contingent fact concerning the bearers of the names that figure in it. All that is necessary is familiarity with some reasonably large subset of the practices making up the Name-Tracking Network. It follows also, consistently with the outline solution proposed in Chapter 3, that there is no relationship of name-reference, conceived as a relationship between a meaning-bearing element of language and some corresponding element of the world, but rather a two-stage relationship, linking both language-element and world-element to practice, and only through those links relating them to one another, of exactly the sort proposed in Chapter 3. Finally, Chapters 5–7 introduce a further claim central to the argument, namely that insofar as a language-element can be said to possess a referent, its referent is invariably some object constituted in part relative to linguistic convention, or as we put it (Chapter 5 §i) a *nomothetic object*. Among the various extraordinary doctrines recommended in this book, this is the one most likely to stick firmly in the craw of anyone of decently Realist philosophical predilections. It will therefore give us particular pleasure to show it, in due course, to be consistent with every variety of Realism to which a decent Realist should wish to subscribe.

We now come to the third and final stage of the main spine of argument in the book; the one mentioned earlier as occupying Chapters 9–12. Here, what centrally concerns us (once again this outline summary of the argument excludes much detail) is the genesis of assertoric content, or to put it nongenetically, what it is that makes certain strings of words ("This table is a metre long") susceptible or truth or falsity, whereas others ("James Peter John") are not thus susceptible. It is characteristic of Referential Realism to hold that the capacity for truth and falsity is not *intrinsically* conferred on linguistic expressions in consequence of the operation of any set of linguistic conventions, although the explanation of the truth-conditions of particular sentences may often require reference to such conventions. On the contrary, according to the Referential Realist, it is in principle possible for the truth-conditions of an utterance, an utterance, that is, taken merely as a semantically unmarked phonemic string, to be explained simply by associating it with some perceptually salient set of environmental conditions. Most Referential Realists also have held that this is precisely the way in which the simplest sentences in a language, those having a purely sensory content and reference, are in fact explained.

In Chapter 10 §§i–v, we deploy against these claims a second *reductio* argument that, although it will strike most readers as wholly unfamiliar, is in fact to be found in Wittgenstein, notably in the early sections of the *Philosophical*

Grammar and the *Philosophical Remarks*. The argument comes in two parts, or stages. Stage 1 runs as follows. Referential Realism implies the in-principle possibility of communicating a grasp of the truth-conditions of a phonemic string S, simply by indicating environmental circumstances in which S takes the value "true," together with circumstances in which it takes the value "false." But that will be possible only if there is some criterion by appeal to which aspects of environmental circumstance relevant to the truth or falsity of S *in virtue of what S means, or asserts*, can be distinguished from aspects that either have no bearing on the truth or falsity of S, or the kind of purely contingent bearing dependent, as Quine would put it, on "collateral information." And, manifestly, no such criterion is, or could be, furnished by the environmental circumstances themselves. Stage 2 of the argument addresses the question what could furnish such a criterion. We suggest, again following Wittgenstein, that possessing the required criterion is equivalent to seeing an intrinsic connection of some sort between the natural circumstances that, according to native speakers, justify the assertion of S and those that, according to native speakers, justify its denial. For to be unable to see any such connection is, precisely, to lack any means of distinguishing between natural circumstances that the native speaker takes as excluding the value "true" for S merely in virtue of the meaning (the assertoric force) of S, and natural circumstances that either have no bearing on the truth or falsity of S, or some connection apparent to the native speaker, but unguessable to the learner because merely contingent in character. Finally, if we now ask what could supply the required intrinsic connection between what is asserted by S and what is asserted by its denial (by "S is false"), the only possible answer (one given by Wittgenstein) appears to be that the assertoric contents in question are related to one another as alternative possible outcomes of the application of some practice. Thus to grasp what aspects of an indicated object are relevant to the truth of S when S is "O is 3 cm. long," it is essential to know what would be asserted by the denial of S, namely that O is *some other length in centimeters*. And grasping this is – can only be – a matter of grasping that statements of length are intrinsically related to one another as expressing alternative outcomes of applying a measuring stick according to the terms of some native system of measurement.

It follows that what makes certain phonemic strings rather than others apt for truth and falsity cannot be the mapping of those strings on to anything assertoric "out there" in the world: anything, that is, in the nature of Russell's "facts." There *is* nothing assertoric in the world external to language. On the contrary, aptitude for truth can only be conferred internally to language, by the manner in which we choose to position sentences as the verbal markers of alternative outcomes in the operation of one or another kind of practice.

We thus see, conformably to the spirit of Wittgenstein's Slogan, how meaning in the sense of aptitude for truth or falsity can be determined prior to the affirmation of any contingent truth. It is determined in determining the

place occupied by specific sentences relative to the conduct of one or more of the multiplicity of practices, from measuring to sorting colour-samples relative to one another, or classifying animals or plants with respect to some chosen principle of classification, which underlie language. Meaning, in conformity with the outline solution proposed in Chapter 3, is a relationship between language-elements and practices, whereas the relationship between language and reality is reconstrued as a relationship between those practices and the aspects and elements of the extralinguistic world on which they operate.

So much, then, by way of a very bare sketch of the main articulations of the book's central argument. The remaining sections and chapters are designed to counter possible lines of objection, to demonstrate the power of the argument to dissolve certain varieties of meaning-scepticism, and to draw out its implications for various phases of the evergreen dispute between relativists and Realists.

The reception of Wittgenstein's thought has long been bedevilled by the accusation that his position in the later work amounts to a version of verificationalism or operationalism, distinguished from others only by the tedious length and obscurity of its exposition. That reading has been helped on its way by the overt verificationism of a good deal of would-be "Wittgensteinian" writing in the field, particularly in the era of so-called Ordinary Language Philosophy. Be that as it may, one of our aims in this book has been to demonstrate, at least in outline, the possibility of a reading of Wittgenstein that reveals the true extent of the gulf separating him from Vienna Circle Positivism, thus incidentally making sense of his own claim never at any time to have subscribed to a verificationist account of meaning. This phase of the argument comes to a head in §§xii–xiii of Chapter 10.

A second line, or rather two opposed lines, of objection to the views proposed here concern the notion of a rule. As will now be apparent, the theory of meaning proposed here relies heavily on the notion of a practice. It might reasonably be assumed that practices are – have to be regarded as – systems of conduct governed by rules. But that thought would seem to expose us to two objections, of equal destructive power, although of diametrically opposed and mutually inconsistent purport, coming from opposite poles of current philosophical debate. On the one hand, stand philosophers such as Dummett, who defend the possibility of a theory of meaning making explicit the rules, a grasp of which, according to them, constitutes mastery of a language. It is an essential part of Dummett's position that a theory of meaning must represent mastery of a language as a species of theoretical knowledge. If this is correct, it must follow that the position recommended here is internally incoherent, as we claim both that meaning is a matter of the involvement of linguistic expressions in practices, and also, in the spirit of Wittgenstein's Slogan, that mastery of language is logically prior to knowledge of any contingent truth, and thus intrinsically nonepistemic. On

the other hand, there are all those philosophers who argue, with Kripke, that there is no sense to be made of the notion of a rule. Kripke holds, with Dummett and most of the recent writers on rules, that the notion of obedience to a rule is essentially epistemic in character. An agent is applying a rule correctly on a given occasion if, and only if, his conduct is *guided* by what he did on past occasions when he applied the rule. But his conduct on those past occasions always can be reinterpreted in such a way as to certify *anything* he may do on the present occasion as obedience to the rule. It follows that there can be no fact about the agent's mental life capable of guaranteeing the correctness of his interpretation of the rule, and hence that the required guidance can only come from the willingness of the language-community in general to accept what he does on each occasion as correct. We – the holders of the view recommended in these pages – thus seem to be left with an unappetising choice between a Dummettian cognitivism that leaves us holding an incoherent theory, and a Kripkean anti-cognitivism that commits any theory of meaning based on the notion of a practice to relativism and meaning-scepticism in their most radical forms.

These issues are addressed in Chapter 8. The reason for placing them here, at the end of the discussion of naming but before the opening of Part III, is, of course, that the questions they raise are pivotal, and cannot be left hanging once the general drift of our proposals has become clear. Answering them requires us to offer some suitable account of what it is to understand and participate in a practice, and this we do in Chapter 8 §vi. According to that account, a speaker is participating in a practice if, and only if, he shares a certain pattern of habits of response and initiation of behaviour with other members of the linguistic community, and exercises those habitual patterns of conduct in such a way as to indicate that he is aware of what advantages are to be gained by doing so, and proposes to gain those advantages. The advantage of such an account is that it allows us to dispense with the notion that conduct can be understood as intelligent only if it is "guided" by appeal to some piece of knowledge, either knowledge of a "rule," or knowledge of some "fact" about past conduct. It opens the way to a *nonepistemic* account of linguistic competence. It enables us to acknowledge the force of the argument underlying Wittgenstein's Slogan, in other words, while avoiding the Kripkean interpretation of Wittgenstein, together with its terminus in a combination of meaning-scepticism and social relativism. The quasi-Humean problem perceived by Kripke, of stating the grounds that justify a speaker's belief that his conduct of a practice, such as counting, is "correct," fails to arise if no such justification is required. And no such justification is required. The competent speaker simply activates learned habits of response that mesh with those of other speakers. Nor, according to us, could there be "grounds" external to the practice to which appeal could be made, to establish the "correctness" of its implementation by an individual speaker. No "fact" about an individual speaker could serve this