

Scepticism and Justification

A. C. Grayling

I

The study of scepticism might be said to define epistemology. As the enquiry into the nature and sources of knowledge, epistemology's two fold concern is to identify and explicate the conditions whose satisfaction will amount to knowledge. Familiarly, one of these crucial conditions is justification. The problems facing the justification of knowledge-claims can best and most powerfully be described by framing them as sceptical challenges, meeting which – if possible – will certify that we are at least sometimes indeed entitled to make claims to knowledge.

Given the centrality of the question of justification in epistemology, and given that the work required of justification is defined by sceptical challenges to our claims to know, it is therefore essential to get the nature of scepticism itself right. My task in what follows is to describe the anatomy of scepticism correctly. At the end I remark that the anatomisation I offer suggests what form a response to scepticism should take.

II

Despite traditional appearances, scepticism is not well described as doubt or denial, nor is it properly understood without limitation of subject matter. Rather, it is best and most sharply characterised as a motivated challenge, in a specified area of discourse, to the makers of epistemic claims in that discourse. The challenge is to defend the grounds offered in support of those claims so that the concerns embodied in the sceptic's motivation for issuing the challenge will be met. His motivation consists in the battery of familiar sceptical considerations which, in the tradition of debate on these matters, have come to be called sceptical 'arguments'. One of the main points I urge is that this is a misdescription. Unravelling this characterisation gives us our anatomy of scepticism.

First, it is a mistake to think of scepticism as consisting in an agnology, that is, a thesis to the effect that we are ignorant either globally or in some region of enquiry. Certain early forms of scepticism (notably the Pyrrhonian) appeared to take this form, but the briefest reflection shows that global agnology is trivially self-defeating (if we know nothing, then we do not know that we know nothing), while local agnologies must themselves consist in positive claims to the effect that we are ignorant in the given sphere, and any positive claim can itself be challenged for its justification. Of course, weak forms of local agnology – which remind us that our knowledge in given regions of enquiry is incomplete, or provisional, and that a healthy attitude of open-minded scrutiny must greet each new claimed advance in them – are perfectly acceptable (and perhaps reflect moderate Academic scepticism, advanced in antiquity in opposition to Pyrrhonism on the grounds that life must be lived). But they do not amount to scepticism in the sense important in epistemology; in this guise they amount merely to injunctions to proportion assent to grounds – in short, to be rational.

But not only is scepticism not well described as the thesis that we are ignorant, it is not even well described as an attitude of doubt. Such an attitude would be premised on the view that there is something inherently suspect about our epistemic practices, a presumption which, when it does not verge on being self-defeating after the manner of global agnology, loads the dice against enquiry before it has offered what it can claim in its support. There is a colloquial use of 'sceptic' to denote one who is very hard to persuade even about the most evident matters – a stance that commits the opposite sin to credulity or too ready assent – which this characterisation conveys. But there is as little reason to think in advance that knowledge claims are by their nature largely doomed as to think that they are all justified. Our interest lies in separating the wheat from the chaff; and this we expect an examination of properly characterised scepticism to help us do.

It does so when we recognise scepticism in a given domain of enquiry as a motivated challenge to make out the justification offered for epistemic claims made in that domain. The best sceptic does not himself claim anything; he asks for a defence of our justificatory practices in the light of certain important considerations relating to them. These important considerations concern contingencies affecting our ways of getting, testing, employing and reasoning about our beliefs. The contingencies in question are familiar from the traditional epistemological debate: they relate to the nature of perception, the normal human vulnerability to error, and the existence of states of mind – dreaming, hallucinating, being deluded, and the like – which can be subjectively indistinguishable from states that we normally take to be appropriate for reliably forming and employing the beliefs in question. By invoking these considerations the sceptic motivates his request to see the support that can be adduced on behalf of claims made in the course of standard epistemic practice. His invocation of them does not support an agnology, nor does it license doubt beyond the reasonable norm in enquiry. If the sceptic's aim were truly to establish either of these two positions, adducing these contingent facts about perception, vulnerability to error, and the rest, would not succeed in doing so without further strong argument to the effect that these knowledge-defeating contingencies are universal, unavoidable, undetectable (and so on), which would be an extremely hard case to make, if only for the reason that it would require of us that we never know what we are talking about when we use such expressions as 'knowledge', 'justification', 'truth', and the like.

But putting matters this way also shows that scepticism is not to be rebutted by saying that we should begin with confidence in our standard epistemic practices, and require the sceptic to show special reason why on any given occasion he himself has justification for mounting a challenge. (Such might be the strategy, for example, of a 'relevant alternatives' rebuttal of scepticism.) The point about the contingent facts which the sceptic can adduce is that they illuminate the (at least *prima facie*) defeasibility of our epistemic practices, and thereby show that there is an onus on us to rule them out when we advance knowledge-claims, or at

least to accommodate them, in ways which do not defeat the project of getting knowledge or at least, well-justified belief.

III

It is important to see how the traditional sceptical arguments conform to this diagnosis of scepticism as challenge rather than claim. One characteristic pattern of such argument is drawn from a set of considerations about error, delusions and dreams. Another trades on facts about perceptual relativities and the fact that our cognitive capacities play a constitutive role in the nature of our experience. Empiricist views suffer particular problems from considerations relating to the nature and limitations of perception, and here the pattern is at its clearest. The best current empirical account of perception tells us a highly complex causal story, beginning with impingements by the environment on our sensory surfaces and ending with the full richness of phenomenal experience and its sequelae in thought and memory. How this remarkable transaction occurs is still mysterious to science and philosophy. But occur it does; and the causal complexity of the process appears richly to invite sceptical challenge. From the subject's viewpoint there might be no way of telling the difference between normally and abnormally caused experience.

The pattern is: if one knows that *p*, then nothing is acting to defeat one's justification to claim knowledge that *p*. But one can seem to oneself fully entitled to claim to know some *p* and in fact lack that entitlement, as the foregoing considerations show. An alternative characterisation is to say that we can have the best grounds for claiming to know that *p*, and yet *p* can be false: the conjunction of the set of propositions asserting the grounds for *p* with the negation of *p* is not a contradiction. Employing this latter idiom captures the underlying logical structure of the defeasibility of epistemic practice, and the contingency of what makes it so. So our claims to knowledge are in need of better grounds than we standardly take ourselves to have. We must, in short, find a way of defeating the defeaters.

IV

One immediate result of grasping the challenge pattern of scepticism as exemplified in these ways, is to note that sceptical considerations are not correctly described as 'arguments'. Sceptical mooting of the familiar considerations is much more like Wittgenstein's 'assembling reminders', by itself enough to show that there is work to be done in justification of epistemic practices. And it follows immediately that if sceptical considerations are neither claims of an agnological tendency, nor arguments purporting to establish an agnology or even just enquiry-undermining infective doubt, then obviously it is a mistake to respond to scepticism as if it were either. In particular, it is a mistake to try to respond to sceptical challenge piecemeal, taking each of the considerations a sceptic might adduce one by one and offering a demonstration that it does not unseat the epistemic project – for the considerations adduced by the sceptic have merely *maieutic* status, and it is a disjunction of them – or any one of them by itself – which prompts the thought that epistemology has positive work to do in support of justification.

The point can be well illustrated by considering (just as an example) Gilbert Ryle's well-known attempt to refute the argument from error by using a 'polar concept' argument in response. There cannot be counterfeit coins, Ryle observed, unless there are genuine ones, nor crooked paths unless there are straight paths, nor tall men unless there are short ones – and so on. Many concepts come in such polarities, a feature of which is that one cannot grasp either pole unless one simultaneously grasps its opposite. Now, 'error' and 'getting it right' are conceptual polarities. If one understands the concept of error, one understands the concept of getting things right. But to understand this latter concept is to be able to apply it. So our very grasp of the concept of error implies that we sometimes get things right.

Ryle assumed that the sceptic is claiming that, for all we know, we might always be in error. Accordingly his argument – that if we understand the concept of error, we sometimes get things right – is aimed at refuting the intelligibility of claiming that we might always be wrong. But of course the sceptic is not claiming this. He is simply asking how, given that we sometimes make mistakes, we can rule out the possibility of being in error on any given occasion of judgment – say, at this present moment. But the sceptic need not concede the more general claim Ryle makes, namely, that for any conceptual polarity, both poles must be understood, and further and even more tendentially – to understand a concept is to know how to apply it, and for it to be applicable is for it to be applied (or to have been applied). This last move is question-begging enough, but so is the claim about conceptual polarities itself. For the sceptic can readily cite cases of conceptual polarities – 'perfect-imperfect', 'mortal-immortal', 'finite-infinite' – where it is by no means clear that the more exotic poles apply to anything, or even that we really understand them. Taking a term and attaching a negative prefix to it does not guarantee that we have thereby grasped an intelligible concept.

However, like Descartes' 'dreaming argument' or any of the other familiar sceptical considerations, it is not required that they themselves be sustainable or defensible in any specially robust way. They are not claims; they are suggestions, considerations, examples, adduced to motivate the challenge.

These remarks imply that sceptical considerations, even if singly they appear implausible, jointly invite a serious response; which is what, centrally, epistemology seeks to offer.

VI

One reflex way of describing the anti-sceptical task of epistemology is (as both Descartes and Russell formulated it in classic texts) to discover the grounds of 'certainty'. This too is misconceived. Certainty is a psychological state that one can be in independently of whether or not one knows or justifiably believes that *p*. The falsity of *p* is no barrier to one's feeling certain that *p* is the case. Religious beliefs, the conviction that

a given horse will win the next race, and many similar instances, amply demonstrate this. The original conflation of the subjective psychological state of certainty with the possession of a secure basis of knowledge is an artefact of Descartes' way of constructing the epistemological task. In his view, the task was to provide a secure route to knowledge from a subjective origin, viz. the origin or viewpoint of the knowing subject's own private data of consciousness. From these subjective data a route had to be found, supported by the right kind of epistemological collateral, to a public world existing outside and independently of that consciousness. Since the knowing subject himself had to be able to sort the contents of his consciousness into those that merited being called knowledge, from those that did not merit the label – that is: since the knowing subject had to be able to be sure that these items were knowledge, whereas those were (say) merely dreams – the mark of knowledge from the subjective viewpoint was thus the nature of the psychological mode of its entertainment, viz. a feeling of certainty.

Nearly all of Descartes successors in epistemology, up to and including Russell and Ayer, shared this view that the starting point for epistemology lies in subjective experience, and therefore faced the same difficulty in securing the epistemic guarantees for at least some of that experience which would rebut sceptical challenges. As with Descartes himself, doubtless the conceptual elision between certainty as a psychological attitude and certainty as a property of foundational or otherwise indubitable truths played its part in this. But in any case, certainty is not the target in epistemology, and if it appears in the residue of showing how to meet sceptical challenges it will be because the kind of epistemic assurance thus provided has, as an expected epiphenomenon, that attendant psychological attitude.

VII

As noted, sceptical challenges inform us that we suffer an epistemic plight, namely, that we can have the best evidence for claiming to know that *p*, and yet be wrong. Scepticism thus demonstrates the existence of a gap between the grounds a putative knower has for some claim, and the claim itself. Traditional responses to scepticism take the form of attempts either to bridge or close the gap. The standard perceptual model, in which beliefs are formed by sensory interaction with the world, together with ratiocination upon the data thus gleaned, postulates a causal bridge across the gap as the basis of knowledge. Given the vulnerability of the bridge to sceptical sabotage, the causal story requires support of some kind.

So do other bridges over the gap – inferential ones, as (again for a prime example) in the Cartesian epistemological tradition. Descartes identified the epistemology's task as the need to specify a guarantee – call it *X* for the moment – which, added to our subjective grounds for our beliefs, protects them against sceptical challenge and thereby elevates them into knowledge. His candidate for *X* was the goodness (the goodness, not merely the existence) of a deity – for a good deity's goodness would ensure that it would not wish us to be misled by what appears to be evidence – so long as we use our epistemic endowments responsibly. The mere existence of a deity is insufficient; it might, as Descartes' own evil demon hypothesis suggests, be a mischievous deity, thereby guaranteeing the unreliability of our beliefs instead.

Some of those of Descartes' successors who accepted his starting point (henceforth, the Cartesian starting-point) likewise sought an *X* to support the inferences from subjective experience to an independent and objective world, but not all of them felt able to invoke a divinity for the task. Yet others sought not to bridge the gap but to close it; this was the strategy of Berkeley and of the phenomenologists. In Berkeley's view, problems arise from thinking that behind or beyond experience there is a material world, where 'material' is a technical philosophical term denoting an empirically undetectable substance existentially supporting the sensible properties of things. Berkeley rejected the concept of matter (he did not reject the existence of the physical world) on the grounds of its empirical undetectability and the fact that sensible properties can, qua experienced entities, only have as their substance what is capable of experiencing them, viz. mind – and thereby closed the putative gap between experience and the world, for the latter turns out to consist in the former.

The phenomenologists made a similar move, with the interesting and more complicated difference that whereas for Berkeley all existence is actual (because everything is always actually perceived by the divine mind), most existence for the phenomenologists is merely possible – for the world is a (logical, in Russell) construction out of sensibilia, by which is meant actual and possible sense-data. Thus all those things not currently perceived by anyone exist as possibilities of perception (in Mill: a physical object is a permanent possibility of sensation). In being committed to the bare truth of an infinity of counterfactual conditionals, phenomenology is therefore less tidy than Berkeleyan idealism, in which all counterfactuals (which have a use only for finite minds: ³if I were in my study I would see my books²) are cashed in terms of what is actually the case from the infinite mind's point of view. (It was once thought that one gets phenomenology from Berkeley merely by subtracting god; this is incorrect – one has to add an ontology of possibilities and with it a commitment to the bare truth of counterfactuals.)

These gap-bridging or gap-closing endeavours all assume, more or less directly, the Cartesian starting-point, and their familiar failure to provide satisfactory responses to the sceptical challenge accordingly suggests, among other things, that there is something deeply suspect about that starting-point. In their different ways Wittgenstein and Dewey both argued that the epistemological enterprise should start in the public domain, not in the privacy of individual consciousness – Dewey nominates the participant perspective, Wittgenstein's private language argument appears to subvert the notion of a Cartesian starting-point altogether. It does so because a private language is what a Cartesian subject requires in order to discourse about the inner contents of his mental life. A private language, in Wittgenstein's sense, is one that is available only to one speaker, not as a contingent but a logical fact: no-one can share the language with that speaker even in principle. Wittgenstein's argument to the incoherence of this notion is this: language is a rule-governed activity, and one only succeeds in speaking a language if one follows the rules for the use of its expressions.

But a solitary would-be language-user would not be able to tell the difference between actually following the rules and merely believing that he is doing so; so the language he speaks cannot be logically private to himself. It must be shared – and can indeed only be acquired, in the first place – in a public setting.

The anti-sceptical possibilities of the private language argument did not seem to be apparent to Wittgenstein himself, for later, in his notes 'On Certainty', he employs a more traditional response to scepticism, somewhat in the tradition of Hume and Kant, by saying that there are some things we simply have to accept in order to get on with our ordinary ways of thinking and speaking. Such propositions as that there is an external world, or that the world came into existence a long time ago, are not open to doubt; it is not an option for us to question them, because they constitute the framework of the discourse within which more particular claims of knowledge and expressions of doubt make sense. Wittgenstein calls them 'grammatical' propositions – thereby according them a constitutive, sense-fixing or practice-fixing role – and described them figuratively as the 'scaffolding' of our ordinary thought and talk; or, varying his metaphor, as the bed and banks of the river of our discourse.

VIII

Whatever one makes of the quasi-Foundationalism of Wittgenstein's approach in 'On Certainty', its similarity to its forerunners in Hume and Kant is suggestive, and prompts one to remember that part of Kant's aim in arguing that our cognitive structures play an essential constitutive role in shaping how the world appears to us, is precisely to address the 'scandal of philosophy', which as he saw it was its inability hitherto to reply adequately to sceptical challenges.

Prescinding from the details of Kant's endeavour, it is arguable that the strategy has much to recommend it, and it has been worked out in more detail both by the present writer and P. F. Strawson. Strawson was criticised by Stroud and others for having smuggled a verification principle into his version of the Kantian transcendental argument; my defence of the strategy consists in describing the foundational beliefs identified by such argument as undischARGEABLE assumptions of our conceptual scheme. If that manoeuvre is right, it shifts the sceptic's position from a challenge about justification to a challenge over the question whether the conceptual scheme is the only one possible – for if it is not, then ultimate security for the practice of making knowledge claims is still lacking since the claims in question are only justified parochially to the given scheme. The task then is to show that it does indeed make sense to talk only of one conceptual scheme (on Davidson's grounds – perhaps, indeed, that 'the very idea of a conceptual scheme' is itself empty if there can only be one). And this is to say that arguing against scepticism thus, and finally, takes the form of arguing against relativism.

IX

So the claim being made here is this: properly anatomised, scepticism is a challenge, in a given area of discourse, to justify the epistemic grounds we assume or employ in that discourse. The motivations (the sceptical considerations adduced) identify where the positive task of justification is to be carried out. To shift attention to relativism as the last resort of the sceptic, it is necessary to make out more fully the claim that invoking facts about our frameworks of discourse (the conceptual scheme presupposed to given types of epistemic activities) settles the sceptical challenge at the level at which it has been traditionally pitched.

The intuition from which to begin is that one cannot know or believe just one thing. A commonplace belief about some object or state of affairs in the world comes as a component of a network of beliefs between which there are complex relations of support and dependency. Questions about these relations, especially about those which provide justification for particular knowledge claims, are among the most important in epistemology. An idea worth examining, therefore, is that the network of beliefs constitutes an implicit inference-licensing scheme, in which specifiable general beliefs play something like a foundational role and in which a particular pattern of inference (I shall suggest that it is deductive inference on something like the covering-law model) is dominant. Each of the points here requires examination.

The expression 'foundational role' has just occurred. A characteristic shortcoming of foundationalist theories of any kind is their failure to yield satisfactory explanations of the relation between what they respectively identify as basis and superstructure in the epistemic edifice. Across a range of proposals there is little persuasive detail about the logical mechanisms by which these different candidates for the role of conceptual support play their part. A promising model of this relationship is suggested by the concept of a 'covering law', the idea being that an assertion about some particular matter is legitimate when its being inferable from a description of its grounds is a result of that inference's being licensed by our conceptual commitments for that region of interest, in the form of one or more covering generalisations. More fully, the idea is that certain assumptions serve as inference-licensors which stand to particular inferences either in the direct relation of a major premise, or as setting the terms in which particular arguments are permitted to count as sound in the standard logical sense. This intuition seems promising; the task is to make it out more fully.

The strategy is to argue that finitary constraints on our capacities as investigators prompt the need for a conceptual framework which enables us to mitigate the restrictions they impose. The two tasks confronting us at the outset are (a) to specify the constitution of the framework, and (b) to describe the relations between the framework and particular beliefs deployed in our ordinary epistemic practice. Because of the enormous practical difficulties of carrying out these tasks, an appropriate starting-point is to explore some possibilities for constructing a model which conservatively satisfies (a) and (b) together. The model is conservative because it attempts to specify the relationships at issue in (b) in standard logical terms, and because the commitments at issue in (a) are thought of, as uncontroversially as possible, in this familiar way: as the belief that there are causally-interactive particulars (and events involving them) occupying space and,

whether or not they are objects of thought or experience while doing so, persisting through stretches of time. Because of the causal character of the relations between the elements of this ontology, we can be regarded as assuming also that the physical realm is nomological in character, a fact which allows us to place a high degree of confidence in the regularities we take it to display. This picture is straightforwardly realist, the key to its being so lying in our commitment to the independent existence and character of the realm over part of which we take our thought and experience to range. The rest we take to be concealed from view, its constituents inaccessible to us either in fact or in principle, depending upon the manner in which they transcend our investigatory capacities.

At this level of generality the description is neutral with respect to finer-grained metaphysical issues, in particular the questions whether there are other things beside concrete spatio-temporal items, and whether some or all of these latter should be thought of as events rather than particulars. On the question of causal laws, however, something more definite is implied by the suggestions to follow.

Certain commonplaces about our limitations as individual epistemic agents, together with certain other commonplaces about our powers in the same respect, offer to explain the role of the first order realist assumptions just nominated as constituting the (model) framework; as follows.

X

It is an all too familiar fact that the epistemic capacities of humans are finite. We suffer the 'finitary predicament'; our empirical resources for acquiring and testing information about the world are limited, and so are our powers of inference, analysis, recognition and memory.

A standard way of dramatising this predicament is to reflect on the circumstances of a lifelong solitary and in particular the question of what he might construct for himself in the way of a world-view with nothing available to him beyond his native cognitive equipment. There are several such models in the recent literature of epistemology, including influential suggestions by Russell, Strawson and Ayer. But one must guard against the failure to distinguish between different notions here: that of the egocentric predicament (the epistemically solipsistic predicament) which is what a true solitary would suffer, and the finitary predicament, which is suffered even by members of a community who share a language and pool their epistemic resources. If the notion of an egocentric predicament makes sense, it does so as a species of finitary predicament, for although what is central to its being effectively a form of solipsism is the individual's isolation, the problem which this renders insurmountable is the limited character of the individual's native cognitive resources, that is, his epistemic finitude. But there is good reason to think that talk of an egocentric predicament fails to make sense. This is because it turns on the idea of a wholly subjective perspective, in which the subject is supposed to recognise his perspective as his own without having any way of locating it in a setting of other perspectives, since these *ex hypothesi* do not exist. And this is controversial: for it would seem that the notion of a sense of self, or at least of a sustained centre of experience which in some sense recognises that experience as its own, cannot be rendered intelligible independently of systematic relations to other such perspectives - other selves - and this implies that to be a self is necessarily to be a member of a community of such things. Among other things, membership of such a community seems to be a necessary condition for acquisition of the scheme, best thought of as embodied in language as the base theory its semantics requires, and by reference to which such experience is enjoyed.

But recognising that the notion of an egocentric predicament is incoherent does not diminish the demands made on our epistemology by the fact of epistemic finitude. Here what is crucial is that even as members of a co-operative we suffer sharp finitary constraints on our epistemic capacities, and the dramatisation of that predicament which the egocentric case affords serves only to identify a gap in need of filling: the gap between what any individual might be imagined capable of constructing by way of a world-view on his own account, and the contents of the conceptual scheme we in fact possess.

Recognising both the existence and character of the gap is what, as noted, forces the abandonment of the Cartesian perspective in epistemology. According to that perspective our privileged access to the data of our own consciousnesses, and their incorrigibility, is what is supposed in large part to underwrite our confidence in what they convey. The sceptical challenges - based on what opens the gap, and renders it unbridgeable; namely, the considerations about perceptual relativities, psychological contingencies, test cases like deception by an evil demon, and the mediate and inferential nature of perception itself - are responsible for this; they are in effect fatal to Cartesian epistemology.

But there are still lessons to be learned from examining the gap which Cartesian epistemologists tried in so many ingenious ways to bridge or close. The lessons flow from trying to answer the question posed above, and which in fuller form runs: if we take seriously the fact that the cognitive capacities of individuals are limited, what follows for an understanding of the global theories we formulate and employ concerning the objects of those capacities? That is: how, given their epistemic limitations, do individuals come to have and use a shared theory of the world with putatively inclusive ambitions, namely, our common-sense conceptual scheme as adjusted and supplemented by science? What status can we suppose that scheme to have, given its radical underdetermination by the evidence which subscribers to it can acquire in the course of activities bearing upon the verification or falsification of the beliefs constituting it?

There is a platitudinous but rather vague answer, already suggested. It is that our sharing a conceptual scheme which none of us individually could have generated, is a function of our belonging to a community whose chief instrument of community is language. Communal possession of language plays the major role in enabling community members to apportion epistemic tasks, to process and record the results, and to put them to use. There is no suggestion here that there cannot be any sort of shared conceptual scheme in the

absence of language, for manifestly there are good naturalistic accounts to be given of languageless creatures displaying concept-applying behaviour in common with others of their kind, frequently in ways indispensable to their co-operative interactions. But we have no reason to suppose that a non-language-mediated conceptual scheme has complexity above a certain level, and it is barely controversial that the one we (humans) possess would be impossible without language.

It would not strictly be incorrect to describe the finitary predicament in terms of the deficiency or incompleteness of our information relative to our practical epistemic needs, in particular our having to choose courses of action. This is a formulation from which discussions of the ampliative character of induction often begin, and it serves as a statement of the dilemma posed by the joint fact of our pressing need for techniques of ampliative inference and their imperfection relative to deductive standards. But it would or at least could be misleading to begin this way, because the problem which confronts us is not so much the deficiency of our state of information about the world as, at a quite different level, its completeness. At the level of detail—of particular matters—our knowledge is radically deficient. But at a highly general level there is a background of assumptions, some of a structural nature, against which our ordinary thought about the world proceeds and which makes it possible; and these jointly take the character of an overarching picture of the world which we hold steady as the interpretative frame for the first, deficient, state of knowledge. It is the character of the overarching picture—the framework—and its relation to the latter, which invites attention here.

XI

So: we are taking seriously the fact that the cognitive capacities of individuals are finite, and asking what follows from this for an understanding of the global theories we apply to the domains over which we exercise those capacities. This is not a question about how, in the light of epistemic finitude, individuals come to have and use a putatively inclusive explanatory theory of the world; rather, it is a question of what work that theory – that conceptual scheme – does. Privileging the second question over the first is something we have been taught by Kant to do: he pointed out that the crucial question concerns not how we get our concepts, but what role they play.

The question can be formulated in alternative ways to bring out others of the concerns which the fact of epistemic finitude prompts. For one important example, we can approach the task by asking what status we can suppose our conceptual scheme to have, given its radical underdetermination by the evidence which the subscribers to that scheme can acquire in the course of the activities which bear upon the verification or falsification of the commitments (the beliefs and theories) in which the scheme consists.

It is evident from the finitary character of individual cognitive powers that the conceptual scheme we employ would be at least extremely difficult to acquire—arguably: impossible to acquire—if it were left to individuals to construct it for themselves from their own resources—if the notion of such a proceeding were, in the light of the Private Language Argument, intelligible. The range of an individual's powers is restricted to his current perceptual environment and whatever of past experience and future expectation his limited powers of recall and inference can provide. Without supplement, these powers would (whether or not equally) at best very weakly support a large number of widely divergent interpretations of what they give their possessor access (or supposed access) to.

This shows the special interest of the discrepancy between the finitary predicament of any of us taken individually, and the richness of the conceptual scheme we each in fact employ. What makes possession of such a scheme possible? One ready and persuasive answer comes from reminding ourselves that we are not isolated individuals, but members of an epistemic community whose chief instrument of community is language. Language enables members of the community to share epistemic tasks and to process, record and utilise the results. There are certain obvious ways in which that process can be portrayed: Popper, for example, with a certain literal-mindedness, thinks of libraries (or, more generally, data banks of various kinds) as embodying the outcomes of the community's joint epistemic activities over time. This must be partly right. But what is more to the present point is that the central role of linguistic competence, in making possible the difference between a rich conceptual scheme and individual finitude, suggests that the key lies in what goes into possession of that competence. And the thought must be that linguistic competence essentially involves or instantiates a theory of the world which enables any speaker of the language to interpret, or indeed to have, experience of the world in that way. Another way to put this is to view linguistic competence as a sixth epistemic modality, where by 'epistemic modality' I mean a means of acquiring, interpreting, processing, storing and transmitting the information yielded by experience and reflection on experience. So considered, linguistic competence is a vastly more powerful epistemic modality than the five other, strictly sensory, modalities; it is what provides them with their framework of operation. What is needed is an account of the scheme instantiated by application of linguistic competence.

XII

One thing we immediately recognise about the scheme is that, as already noted, it is realist in character. Despite appearances and much misleading debate, realism is an epistemological thesis asserting the independence of the objects of discourse from discourse itself. More precisely, it asserts that relations between thought and its objects, perception and its targets, knowledge and what is known, acts of referring and referents—call them 'mind-world' relations, although in fact they are all different if intimately connected—are external or contingent ones.

We can recognise, as fundamental to understanding the way in which our discourse works, assumptions of the kind listed earlier about the world being a law-governed realm of causally interacting spatio-temporal

particulars, and events or other entities typically individuated by reference to these. Moreover, commitments in these respects reveal why ordinary discourse invites accounts of reference and truth which are distinctively realist, for a view of the foregoing kind about the domain over which our discourse ranges is very naturally interpretable in terms of those familiar views about the links between referring devices and things, and between sentences and objectively obtaining truth-conditions, which are presupposed to much recent discussion in this area: namely, that reference works on a naming-paradigm in causally direct ways, and that truth in some sense consists in fit between what we say or think about the domain over which discourse ranges, and the domain itself.

Whatever particular difficulties affect giving an exact account of the ontology and the semantics invited by this realist picture, it is at least clear that it constitutes a simple and powerful view which on the whole successfully sustains the demands made on it by experience—a strong pragmatic justification for it. That is a fact which is independent of debates about whether the realist commitments of the scheme are literally true or not, a question upon which much turns; but for present purposes they can be left aside, because we need only note that we are construing the commitments weakly as assumptions of the scheme. Whether they are taken as literally true or merely as assumptions, the scheme's justificatory character remains. It becomes a matter for the second task, identified earlier, to settle this question of 'literal truth', that is, how these first-order facts about the scheme are to be interpreted in the light of the sceptical problem which gives that second task its content.

XIII

Collecting the suggestions already made, we can venture the following as a first approximation of what the framework of ordinary epistemic practice looks like, treated as a justificatory scheme. Such a scheme is in effect an inferential scheme, representable as providing security for familiar practices of basing judgements on evidence. In standard thinking about these matters, empirical judgements are thought of as inductively based on the evidence for them, and as being defeasible to the degree that the evidence is partial. But we have just noted that the scheme consists of a set of assumptions about the nature of the world over which our experience ranges, and we have further noted that these assumptions include some to the effect that the world is lawlike and independent. Add these assumptions to statements of evidence as supporting premises, and the logical picture changes: we see that the form of reasoning being employed is enthymematic deduction on the covering-law model.

At its roughest, the picture is something like this. A judgement about some particular matter of fact is inferred from the evidence for that judgement (reported by evidential premises) in the presence of more general premises about the kinds of things in question, and even more general standing premises about the world (background premises). As such the form of reasoning is representable as deductive: the conjunction of evidential and background premises entails the judgement. But, of course, empirical judgements are defeasible, which appears to conflict with the idea that inferences to them take deductive form. The answer lies in noting, firstly, that background premises have to carry *ceteris paribus* clauses, or clauses about normal conditions; and secondly, that evidential premises are only as good as the evidence they report, and here the usual finitary constraints apply. Accordingly we can be, and often enough we are, wrong in our judgements. But we can often measure the degree of confidence that we repose in our judgements, by taking into account the relevant defeating possibilities inherent in either or both the evidence that evidential premises report, and the stability of the normal conditions assumed in background premises. This is where this picture saves what is persuasive about conceptions of probability, and in effect 'solves the problem of induction' by suggesting that all reasoning is always deductive in form. (For example, inferences by analogy assume uniformity of nature grounds—and so for other non-enumerative inductive forms.)

It is illuminating to think of Aristotelian classification as, obliquely, among the forerunners of this idea of an inferential scheme. Two reasons why there is only heuristic value in remarking the connection are these: the Aristotelian system of classification by genus and species is too neat (too simplistic) even for the domain where it has most plausibility, namely, the biological domain; and secondly, discussion of it came to be distracted, perhaps not unnaturally, into discussion of definition, and in that guise the objections to it are many and obvious. Some are tellingly summarised by Locke in the *Essay* (not every term can be precisely explained by two other terms giving genus and differentia; and some words cannot be lexically defined, on pain among other things of regress and failure to constrain their meanings by reference to extralinguistic considerations). But there is much that is suggestive there (and in later logic: see the Kneales on medieval theories of assumption). Note for example the striking resemblance between the Tree of Porphyry and the structure of reasoning employed in the game of Twenty Questions, in which moderately skilled questioners can identify any individual spatio-temporal object usually in fewer than twenty steps, exploiting classificatory conventions governing our picture of the world together with appropriate cognitive strategies. What both seem to capture is the sense that inferences about matters of fact proceed to their conclusions by way of the deductive inferential structure outlined: premises of relatively great generality are conjoined with premises of relatively lesser generality, including particular ones, in any number of steps within practical constraints, to yield, in a highly reliable way even in the face of defeasibility considerations, judgements about matters of fact.

XIV

These remarks are merely schematic, and only gesture in the direction of a research programme. But the implications of such an approach are clear. The quest for an account of justification is satisfied by this picture, as explaining how it consists in inherence in a scheme or framework. Full statements of justification proceed, via a report of the relevant evidence in the case, to appeal to the scheme as a whole: 'this', we might in the end say in such a transaction, 'is how (we think) the world is', and that has to satisfy the sceptic at that level

of enquiry. He might then—he should—raise his sights to the question of the justification of the scheme as a whole, but that, as noted is a different and further matter. (Some, like Carnap and Wittgenstein in their different ways, would take it that there can be no such higher task.) But then the problem of scepticism comes to be seen as arguably it ought to be seen: as the problem of relativism.

The standard difficulties concerning justification can be taken implicitly to specify desiderata that have to be satisfied by any adequate theory of justification. They are not best satisfied by attempting to defeat the defeaters proposed by sceptical arguments of the familiar sort, as earlier epistemology often tried to do, not untypically on a blow-by-blow basis. Rather, they are satisfied by a positive theory which shows how justification is secured. A theory like the present one, which postulates ultimate justification by reference to the assumptions of the scheme, serving as foundational premises from which, together with other premises, judgements of lesser generality are deduced, accordingly satisfies these desiderata. It does so conditionally upon resolution of the higher order question about the overall justification of the scheme itself, of course, which is where the determining connection of the first to the second order enterprises becomes manifest: it shows that something like a transcendental deduction of the assumptions of the scheme, together with an argument that any alternative scheme can only be a variant of this one (the modality is seriously intended), is required to deal with scepticism fully. Once again, this higher-order task is sometimes claimed to have been effected already in the lower-order one, on the grounds that only the lower-order one is possible anyway; but that is precisely one of the central controversies of recent epistemology.

Among much else that might be said about the strategy outlined here, I shall mention in conclusion just one. In earlier debates about empirical knowledge it was pointed out that reasoning about matters of fact does not take the form of conjunctions of evidential and background premises entailing judgements. Judgements, especially perceptual ones, are typically immediate, appearing to consist in the exercise of well-rehearsed, experience-based recognitional capacities. As a description of how empirical judgement feels (so to speak), this is surely correct; but it involves a confusion of psychological and logical facts about its structure. True, we do not as a matter of psychological fact usually go through processes of inference in such cases; but we can see that if challenged to justify a judgement we would have to state our grounds and, if pressed, the background assumptions that constitute them as grounds. In a full story of this kind we would have the conceptual scheme qua inferential framework fully present, and would be able to trace inferential routes to the judgement itself.

XV

The nub of the claim here is that we can redescribe the problem of justification as the problem of epistemic finitude, and, by seeing how such finitude is overcome—viz. by our possession and employment of a realistic conceptual scheme which serves as a geographical-historical explanatory framework designed to make experience coherent, serving as a framework of inference, specifically deductive in form, in which the general assumptions of the scheme play an undischageable role—we thereby see how we come by justification for our workaday epistemic judgements. And we thereby also identify where the major philosophical task in this region lies: namely, in justifying the scheme itself, which is the same thing as refuting scepticism in its most interesting and substantial guise, namely, as relativism. This is a different problem of justification, but it only comes to the fore when epistemology's traditional problem of justification has been dealt with in the way suggested here.