

**Do We Really Need to Know
That We Know?**

Giorgio Volpe

M.Phil. Thesis

**The London School of Economics and
Political Science**

UMI Number: U055823

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U055823

Published by ProQuest LLC 2014. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

THESES

F

7042



*211596611

Abstract

A popular argument against realism rests on the claim that the correspondence theory of truth cannot possibly play any significant role in our epistemology (Putnam, Fine) or theory of meaning (Dummett, Wright) because of the epistemic gap it opens between evidence and truth.

This thesis develops a fallibilist reply to this anti-realist argument and provides an assessment of the epistemological consequences one must be prepared to accept if that reply is to prove successful.

Chapter 1 attempts to clarify what the issue of realism is about. The ontological, epistemological and semantic aspects of realism are sorted out, and the relationship between realism about the external world and various theories of truth is explored. Endorsement of a correspondence theory of truth (as opposed to an epistemic one) is argued to be of great importance to a successful defence of realism about the external world, though not identical with it. Chapter 2 argues that commitment to the epistemic accessibility of what Putnam calls 'the God's eye point of view' is inessential to the correspondence theory of truth. Saying that the goal of our cognitive efforts is to come up with sentences which are correspondence-true does not commit one to hold that human beings can reach an absolute standpoint from which they can view the world as it is in itself and compare it to their own representations. A discussion of the recent debate between internalist and externalist epistemologies leads to the proposal of a naturalistic view of knowledge according to which knowledge of the things in themselves needn't be knowledge from a God's eye point of view. Chapter 3 discusses how the rejection of the claim of the epistemic accessibility of the God's eye point of view affects the epistemological problem of assessing the truth-conduciveness of competing sets of methodological rules. The objection that a realism 'without a God's eye view' is in fact only a concealed form of scepticism is rejected as a legacy of the internalist standpoint of traditional epistemology.

Contents

Introduction

What do We Need to Have Epistemic Access to?	5
---	----------

Chapter 1

Realism and Truth

1.1	Is Realism a Doctrine About Meaning?	10
1.2	Commonsense, Scientific, and External World Realism	13
1.3	Scientific Realism and Correspondence Truth	17
1.4	External World Realism and Truth	19
1.5	Correspondence Truth and the God's Eye Point of View	24

Chapter 2

Realism Without a God's Eye View

2.1	The Quest for Certainty	31
2.2	Three Kinds of Internalist (Externalist) Doctrines	37
2.3	The Vertical Epistemic Regress Problem	41
2.4	Nozick's Externalist 'Refutation' of the Sceptic	43
2.5	Externalist Knowledge and Internalist Justification	49
2.6	Is Epistemic Justification Essential to Knowledge?	51
2.7	The Divorce of Certainty and Truth	56
2.8	Is There any Internal Link Between Justification and Truth?	59

Chapter 3

Do We Really Need to Know That We Know?

3.1	The KK and Objectivity Principles	65
3.2	Why Should We Value Internalist Justification?	68
3.3	Justification and Successful Argumentative Interaction	71
3.4	Styles of Reasoning	73
3.5	Internal and External Rationality	77
3.6	Epistemological Contextualism	78
3.7	The Context-Relativity of Epistemic Justification	81

Conclusion

Is Fallibilism Merely Scepticism in Disguise?	85
--	-----------

References	88
-------------------	-----------

What do We Need to Have Epistemic Access To?

The problem of the epistemic accessibility of truth has been one of the major worries of modern philosophy, but only in relatively recent times some people have come to suspect that the main source of the problem might lie in a poor concept of truth rather than - as Descartes and many philosophers after him have thought - in an inadequate methodology. If truth appears epistemically inaccessible - so goes the argument - that must be the effect of a mistaken concept of what it is for a sentence to be true or false. If a sentence is deemed to be true (false) when it corresponds (fails to correspond) to how the world is in itself, no wonder that epistemic access to truth will appear quite difficult. For how can one possibly know how the world is in itself in the first place, and then compare it with one's own representations? What we really need is not a better methodology, but a better notion of truth, that is, an *epistemic* (or *pragmatic*) notion of truth¹.

I am not impressed by this sort of argument. The most general view underlying the present work is that correspondence truth may after all be epistemically more accessible than most (reasonable) brands of epistemic truth currently available on the philosophical market. If it is possible to show that epistemic access to correspondence truth needn't be so problematic as an exceedingly intellectualistic epistemology would have it, the most powerful motivation for accepting an epistemic or pragmatic theory of truth will be undermined.

The argumentative strategy explored in this work is to separate commitment to the correspondence theory of truth from commitment to the epistemic accessibility of what Hilary Putnam calls 'a God's eye point of view'. If epistemic access to correspondence

¹ Epistemic (or pragmatic) doctrines of truth understand truth as a function of such epistemic (or pragmatic) concepts as verifiability, warranted assertability, rational acceptability, etc.

truth implies access to a God's eye view, it can easily be contended that insuperable difficulties will prevent human beings from gaining access to correspondence truth. But I shall try to argue that we needn't be committed to an epistemology that makes of the epistemic access to an absolute perspective a necessary condition for the pursuit of truth. In other words, I shall try to show that, provided we are willing to endorse a naturalistic picture of the cognitive activities of human beings, *knowledge of the world as it is in itself needn't be knowledge from the God's eye point of view*. Hence its achievement may prove less arduous than it is usually thought. I propose to explore the implications of the idea that epistemic access to correspondence truth is achieved by developing representations of reality that correspond to the world as it is in itself, rather than by gaining access to an absolute standpoint from which those representations could be compared to that world.

To gain access to the Senate House of the University of London is to be allowed to enter its building. One needn't be *aware* of entering the Senate House to gain access to its building. In the present work I shall try to spell out the epistemological consequences of the view that epistemic access to the world resembles in significant respects physical access to buildings. According to this view, to gain epistemic access to the world is just to develop beliefs which are *sensitive* to how the world is. One can gain epistemic access to the world without ever being able to take that absolute perspective which could enable one to *see* if one's beliefs match the structure of the world.

I oppose the view that epistemic access to reality requires that one's beliefs be *justified as true*. It is under this view that endorsement of correspondence truth will imply either the epistemic accessibility of the God's eye point of view or surrender to the arguments of the sceptic. Facing a similar choice, one may well be tempted to define truth epistemically, so that no gap can arise between evidence and truth. But if it can be shown that the epistemic gap between evidence and correspondence truth does not preclude the epistemic accessibility of truth, then we shall not be confronted with that choice in the first place. This does not however rule out that the gap may affect the possibility of proving that our beliefs are true.

The problem of a correct characterization of what it is for a sentence to be true or false does not merely affect the issue of the epistemic accessibility of truth. It also affects the issue of what sort of 'things' we can have epistemic access to. I do not believe

that the traditional controversies between realist and anti-realist philosophers can be reduced to controversies between evidence-transcendent and epistemic theories of truth. Still, it is a fact that some theories of truth can be more easily combined with a realist conception of the objects of knowledge than others. In particular, a successful defence that the correspondence theory of truth does not make truth inaccessible will provide considerable support for the realist claim that the world we can have epistemic access to is not merely a 'phenomenal' world (as Kant believed it should be, if human knowledge is to be objective), but the world as it is in itself. A successful defence of the correspondence theory of truth will thus provide indirect support for the thesis that *there is* a world in itself to which our representations can correspond or fail to correspond.

By providing an argument in defence of the correspondence theory of truth, one undermines that train of thought that from acceptance of epistemic truth leads to the belief that full-blown realism can no longer be a viable metaphysical option. But precisely to what kind of realist mill can my argument, if sound, bring grist?

The way the issue of realism is currently addressed by English-speaking philosophers makes it appear as if a very long time has elapsed since J.L. Austin first pointed out that it is the *negative* use of the word 'real' that 'wears the trousers'. By this he meant that 'the function of "real" is not to contribute positively to the characterization of anything, but to exclude possible ways of being *not* real' (Austin 1962, 70). I do not think that philosophers owe a particular allegiance to ordinary language. Provided they have some interesting philosophical question to discuss, I cannot see why they should not be allowed to use the word 'real' as best suits their needs. However, I do think that Austin's remark is valuable advice. Here is a first clue to the kind of realism I shall deal with in this work.

The realism I am interested in contradicts the claim that the (only) world we can know about, and (hence?) the only world that can exist, is a *phenomenal* world, a world which is *constituted* by our knowledge. This realism is opposed to that particular brand of 'knowledge-phenomenalism' which denies that trees, cats, or persons can *really* exist as trees, cats, or persons (and not simply as unknowable things in themselves) independently of our knowledge, language, or conceptual schemes. In this philosophical sense of 'real', there is indeed some common way in which trees, cats, and persons can fail to be real: they can all fail to be real entities in the world (as opposed to mere

phenomenal entities), if they are thought to be constituted as trees, cats, and persons by their being actual or potential objects of knowledge².

If it is possible to resist the conclusion that correspondence truth must be epistemically inaccessible by rejecting the claim that epistemic access to correspondence truth presupposes epistemic access to the God's eye point of view, it is conceivable that trees, cats, and persons can be known by human beings without being constituted by human knowledge.

Here is an outline of the structure of the present work. In the first chapter I attempt to clarify the relationship between realism and different theories of truth. In this chapter I take realism to be a metaphysical doctrine about what there is rather than a semantic doctrine about the meaning of certain sets of sentences. I also present what I call the anti-realist arguments from the idleness of correspondence truth, and explain why a realist unwilling to endorse a deflationary conception of truth ought to reject the claim that endorsement of correspondence truth makes it impossible to understand how human beings could acquire any knowledge of the world.

In the second chapter I provide an example of the dogmatic epistemology I believe a realist needn't be committed to and present my tentative reply to the claim that endorsement of correspondence truth makes it impossible to understand how human beings could acquire knowledge of the world. Internalism and externalism in the theories of knowledge and of epistemic justification are dealt with at some length. The claim that epistemic justification is not essential to knowledge, but to knowledge claims, provides the foundation for a naturalistic view of knowing according to which correspondence truth may prove even more accessible than epistemic truth.

In the third chapter I spell out some further epistemological consequences of my tentative defence of correspondence truth. I discuss the KK principle ('If S knows that p, then S knows - or at least can know - that she does') and the relevance of the claim that 'One can never know that one knows that p' for the assessment of competing sets of methodological rules. A 'pragmatic' view of epistemic justification is put forward as an

² The ideas of Kant, Kuhn, Goodman, and Putnam illustrate several ways of denying that trees, cats and persons may be real parts of the world as it is in itself.

attempt to support the thesis that it can be rational to develop and apply criteria of epistemic justification which cannot be proved to be truth-conducive.

The last section of the work weighs the pros and cons of the argumentative strategy developed in the text and ventures a tentative conclusion about the tenability of a realism without a God's eye view.

Realism and Truth

1.1 Is Realism a Doctrine About Meaning?

Realism seems to be a very popular topic among English-speaking philosophers. Many different views in such heterogeneous disciplines as metaphysics, epistemology, semantics, the philosophy of science, and the philosophy of mathematics are currently labelled as 'realist' and contrasted with 'anti-realist' views. One can hardly see what the features shared by those realist and anti-realist views respectively are.

The most influential bid for an overall approach to the various issues raised in current (and past) debates about realism is probably that of Michael Dummett. He claims to have discovered the common features shared by many different debates commonly regarded as involving the issue of realism (Dummett 1978, 45). In his view, the realism-nominalism controversy in the theory of universals, the realism-idealism controversy in metaphysics, and many other philosophical controversies, like those about the reality of the past and the existence of mathematical entities, all are essentially *semantic* in character. Any of these controversies concerns the interpretation to be given to the statements in some given class. In each case, the hallmark of realism is endorsement of the thesis that statements in the relevant class are determinately either true or false independently of our being capable to recognize their truth-value.

If one accepts the widespread belief (which goes back to Frege) that to know the meaning of a sentence is to know the conditions under which that sentence is true, a consequence of Dummett's definition is that realism will be committed to a truth-

conditional theory of meaning where truth is to be conceived as (possibly) evidence-transcendent³.

In short, Dummett's contention is that realism and anti-realism are primarily theses about *meaning*, and only derivatively theses about *what there is* (Dummett 1978, xl). Semantics is regarded as more fundamental than metaphysics. Dummett claims a greater generality for his semantic approach, on the ground that it encompasses controversies (e.g. about the future and about ethics) which can hardly be construed as metaphysical controversies about the existence of entities of a given kind (Dummett 1982, 55).

Later on, I will say a few words on Dummett's anti-realist arguments from the acquisition and manifestation of linguistic competence. In this section, I only want to address the question of the adequacy of his semantic characterization of the realism issue.

Dummett's definition conflates the issue of realism with the issue of truth. In his wake, it has indeed become fairly common in recent years to think of realism as a commitment to an evidence-transcendent, absolutist notion of truth (see, e.g., Haack 1987, 276-284). Nevertheless, there appear to be good reasons for keeping the two issues clearly distinct.

Dummett remarks that the commitment to an evidence-transcendent notion of truth entails the existence, roughly, of a knowledge-independent reality in virtue of which the statements in some given class are either true or false (Dummett 1982, 55). This entailment is supposed to explain why the doctrine of realism is naturally cashed in metaphysical terms. But talk about the existence of some knowledge-independent reality, Dummett suggests, is merely metaphorical (Dummett 1978, xxv-xxvi, 229), and can be usefully dismissed in favour of his semantic approach.

Two points can be made against Dummett's contention. First, metaphysical talk of a knowledge-independent reality appears perfectly meaningful in its own right, and not parasitic upon semantic talk as Dummett suggests. Second, the semantic thesis about the evidence-transcendence of truth and the metaphysical thesis about the existence of a

³ A conception of truth is said to be evidence-transcendent when it places some or all of the sentences of our language beyond our capacity to recognize whether they are true or false. In other words, endorsement of an evidence-transcendent notion of truth means commitment to count as true-or-false even those sentences whose truth-values we have no way of settling.

knowledge-independent reality are logically independent: the former does not entail the latter, nor the latter entails the former.

As for the first point, Dummett himself has often presented his anti-realism as a generalisation of mathematical intuitionism. Thus Michael Devitt is likely to be right when he points out that Dummett's thesis on the metaphorical character of metaphysical talk about reality arises from his philosophy of mathematics (Devitt 1991, 264-266). In that particular field, it may be sensible to suppose that platonic talk about mathematical objects is purely metaphorical, as the real problem appears to be, in Kreisel's words, 'not the existence of mathematical objects, but the objectivity of mathematical statements' (quoted in Dummett 1978, xxviii). However, there is no reason to believe that metaphysical talk about the existence of physical bodies should inherit the supposed metaphorical character of platonic talk about mathematical entities. After all, physical bodies, unlike platonic entities, can bring about causal effects on human beings!

The thesis that metaphysical talk about the existence of a knowledge-independent world is not to be taken as parasitic upon the semantic talk about the interpretation of the statements in some given class is further supported by our second point. In effect, if the semantic and the metaphysical theses about realism are logically independent, that will mean that the former enables us to say things that are not expressed by the latter.

Does the evidence-transcendence of truth entail the existence of a knowledge-independent world? Apparently it doesn't. For one can think of a version of coherentism according to which truth is in a strong sense evidence-transcendent without this implying the existence of a knowledge-independent world⁴.

Does the existence of a knowledge-independent world entail the evidence-transcendence of truth? Surely it doesn't, for one can easily conjoin the metaphysical thesis about the existence of a knowledge-independent world with a deflationary (disquotational) theory of truth. And a deflationary theory of truth needn't say anything

⁴ For such a version of coherentism, see Dancy 1985, 138 f.: if the notion of a 'fully coherent set of propositions' is construed in a suitably idealized form, so that 'at any time, or timelessly, there is, for each set [of propositions], a larger and more coherent set', then 'for any such set, no matter how large, there remains the possibility that it is false'.

But note that Dancy identifies commitment to the existence of evidence-transcendent truths with *realism* (Dancy 1985, 19), and thus takes his epistemology to provide a *realist* version of coherentism. However, Dancy's coherentism needn't come out as a realist position on the basis of a non-semantic characterisation of realism and anti-realism as the one I am trying to defend.

at all about the problem of the epistemic accessibility of truth (compare Devitt 1991, 41 and Grayling 1992, 52). Alternatively, one could adopt some form of dogmatic epistemology involving the claim that we have immediate, intuitive access to some set of axiomatic truths, and demonstrative knowledge of their consequences. If the information so acquired is adequately comprehensive, such a dogmatic epistemology will be consistent with the metaphysical thesis of the existence of a knowledge-independent world without being committed to the possibility of evidence-transcendent truths.

If these arguments are sound, then there seems to be a logically independent, metaphysical thesis, that Dummett's semantic approach cannot possibly capture. True, some philosophical positions which are classifiable as realist on the basis of Dummett's criterion turn out to be totally independent of any such metaphysical thesis. Shall we say that the metaphysical thesis is not essential to realism? Anthony Grayling sensibly suggests that 'what we should say about those "realisms" which are not readily classifiable in terms of some reality or realm of entities is, simply, [...] that they are not *realisms*' (Grayling 1992, 54).

So there appears to be some point in trying to devise a precise formulation of the metaphysical thesis of realism as distinct from the semantic thesis of the evidence-transcendence of truth, if only to spell out in greater detail how these theses are related to each other.

1.2 Commonsense, Scientific, and External World Realism

In the next three paragraphs, we shall focus our attention on the metaphysical thesis of realism in its utmost generality. However, it's now time to make clear that what I am ultimately concerned with is the traditional thesis of external world realism. The following characterization of the general thesis of realism is thus to be seen as a first step towards an appropriate definition of the more specific claim of external world realism.

Devitt points out that the metaphysical thesis of realism implies two kinds of claims: *independence* claims and *existence* claims (Devitt 1991, 14-22). He then characterizes the first kind of claims as claims about the knowledge- and mind-independence of those entities which are said to exist (he actually says claims about the

'objectivity' and the 'non-mental' character of those entities). He builds the mind-independence requirement into the independence claim because he is exclusively concerned with the reality of the external, physical world. That's what I am also primarily interested in. But in defining the metaphysical thesis of realism, one needn't confine oneself from the beginning to external world realism. Thus, I propose to say that the metaphysical doctrine of realism makes an *ontological* claim about what entities exist and an *epistemological* claim about their independence from anybody's knowledge of them⁵.

The *epistemological* claim is indeed essential to realism: those entities which are asserted to exist must exist independently of any actual or possible knowledge of them. They needn't be epistemically accessible to anyone (empirical, transcendental or godlike subject) in order to exist - which is not to say that they *cannot* be. (Of course artifacts owe their forms to their manufacturers' knowledge, as well as some kind of knowledge-dependence is involved in the existence of tools and social entities; but neither artifacts, nor tools or social entities, bear an *epistemological* dependence on the knowledge of their knowers⁶).

The *ontological* claim specifies what *kind* of entities are supposed to exist. Devitt notes that for a person who rejects, as several philosophers are today inclined to do, both the 'in corrigibility thesis' and the 'self-intimation thesis' about the mental⁷, the question whether or not a mental event exists will be an altogether objective question (Devitt 1991, 15). According to our proposal, if such a person does believe in the existence of mental events, as far as these events are concerned she will be classified as a realist. External

⁵ Another reason why I prefer talk of knowledge-independence to (the once more common) talk of mind-independence is that one may envisage cases of knowledge-dependence which are not, strictly speaking, cases of mind-dependence. Instinctive behavioural dispositions programmed by natural selection in the nervous system of several kinds of living creatures and formal constraints ascribed by neo-kantian philosophers to transcendental subjects of a very abstract and impersonal sort may give rise to instances of knowledge-dependence which are only in a very loose sense instances of mind-dependence.

⁶ The kind of knowledge-dependence the realist wants to deny has nothing to do with the trivial causal relations by which such mental states as beliefs and desires can affect the physical world. What she wants to deny is rather that kind of *epistemological* knowledge-dependence where the knowing subject *constitutes* the object of her knowledge and can have epistemic access only to such a *phenomenal* object as opposed to the object as it is in itself.

⁷ The 'in corrigibility thesis' is that a person cannot be wrong about, and the 'self-intimation thesis' that she cannot be ignorant of, her own mental states.

world realism will require a further step, that is, the assertion that entities *of a non-mental kind* exist.

Speaking of the knowledge-independent existence of entities of a non-mental kind is of course very vague: 'This commits realism only to an undifferentiated, uncategorised, external world, a Kantian "thing-in-itself"' (Devitt 1991, 17). An external world realist will want to commit herself to more than that. To provide a suitable definition for something like common-sense or scientific realism, one needs to be more specific about the kinds of entities one claims to exist. These cannot be mere Kantian noumena. Beware: I am not saying that common-sense or scientific realism entail any particular claim about the epistemic accessibility of reality. I am just saying that they are committed to the knowledge-independent existence of entities which are not Kantian noumena. If one is a realist about trees, all that one is committed to is the truth of the statement:

$$(Ex)(Tx \ \& \ ((t)(\neg Ey)(Kyxt) \Box \rightarrow Tx))$$

where 'Tx' stands for 'x is a tree', 'Kyxt' stands for 'y knows at t that x is a tree', and the box-arrow symbolises the subjunctive conditional 'If it were the case that..., then it would be the case that...'. (To be sure, one should also specify at what time the entities in question are claimed to exist; but for our purposes we can omit such details).

Making that statement, one doesn't say anything about one's grounds for making it or the epistemic accessibility of trees (although making that statement may *implicate* the commitment to provide some justification for its utterance). In particular, one does not say that she is always, mostly, or usually right about what she takes to be trees⁸. Realism as I am trying to define it entails an epistemological claim about the knowledge-independence of the entities which are said to exist, but leaves completely unsettled the question of their accessibility to human knowledge.

It is clear that we cannot phrase realism about chairs, paintings and poems as we have phrased realism about trees: the existence of such things presupposes the existence of their manufacturers. However, in order to avoid unnecessary complications, I will

⁸ This is clearly incompatible with those theories of meaning which make successful reference depend upon the correctness of our beliefs about the object in question. But today we have viable alternatives to such theories (see, e.g., Putnam 1975).

restrict my definitions of common-sense and scientific realism to what we might call 'natural kinds' (though I am aware that several scientific 'phenomena' are human products in a very close sense to tables and chairs; but see note 3 above). I don't think this restriction can significantly affect the ensuing arguments. After all, what I am mainly concerned with is external world realism in general, not realism about social or cultural wholes such as institutions or text-meanings.

This is how I propose to understand the doctrines of common-sense and scientific realism:

Common-Sense Realism: Tokens of most current common-sense physical types exist independently of anybody's knowledge of them.

Scientific Realism: Tokens of most current unobservable scientific physical types exist independently of anybody's knowledge of them.

These definitions are modelled on those of Devitt (1991, 24), but the independence requirement is here purely epistemological. For economical reasons, they involve quantification over types, but this is not an essential feature of them (see Devitt 1991, 20 f.).

External world realism can be identified either with common-sense realism or with the conjunction of common-sense and scientific realism⁹. We may agree to call 'modest' external world realism the former version, and 'comprehensive' external world realism the latter. In what follows, I shall confine my arguments to the modest version, which can be more easily defended (or so it appears to me). Whenever I speak of 'external world realism', the phrase should be understood as referring to the modest version. Indeed, I shall mostly write 'Realism', *sic et simpliciter*, with the capital 'R', as a shorthand for 'external world realism'.

⁹ Grover Maxwell and Wilfrid Sellars have argued that common-sense and scientific realism are in fact incompatible and the former should give way to the latter. I think there is a plurality of descriptive levels at which reality can legitimately be depicted, any of them possessing a peculiar explanatory role in our picture of the world. But I won't argue this point, because it is completely inessential to my purposes.

1.3 Scientific Realism and Correspondence Truth

What can be said about our proposed definitions? It will be convenient to start from *scientific* realism, which has been the focus of extensive debate in recent years.

According to W.H. Newton-Smith, the word 'realism' has been used 'to cover a multitude of positions in the philosophy of science, all of which, however, involve the assumption that scientific propositions are true or false where truth is understood in terms of a cleaned-up version of the correspondence theory of truth' (Newton-Smith 1981, 28 f.). A cleaned-up version of the correspondence theory of truth is characterized by Newton-Smith as the claim that 'to be true (false) is to be true (false) in virtue of how the world is independently of ourselves' (Newton-Smith 1981, 29). I wouldn't say, as Newton-Smith does, that *all* the different versions of scientific realism which have been proposed in recent years involve the acceptance of some sort of correspondence theory of truth. It is true, however, that *many* of them do, and it is also true that many of them draw their deeper inspiration precisely from that acceptance.

Another claim which currently goes under the label of 'scientific realism' is the claim that 'the historically generated sequence of theories of a mature science is a sequence of theories which are improving in regard to how approximately true they are' (Newton-Smith 1981, 39). Talk of 'convergent' or 'cumulative' realism hints at this feature of scientific realism.

Apparently, then, the content of our definition of scientific realism is much narrower than one would be likely to expect on the basis of the current debate. Ought we to have built the semantic thesis about the nature of truth and the epistemological thesis about the increasing verisimilitude of successor theories in mature science into our concept of scientific realism? I believe we do well not to.

Obviously, anybody is free to define scientific realism as they like best: nothing hangs on the labels we attach to our concepts. Yet, I prefer to treat (1) the ontological claim about the knowledge-independent existence of the unobservables posits of science, (2) the correspondence theory of truth, and (3) the convergency thesis as separate issues. Even if it is true that (1), (2), and (3) appear to many philosophers as the elements of a whole, cogent picture (see, e.g., Popper 1963, ch. 10, and Newton-Smith 1981, 28-43), nevertheless they are logically independent, and an increasing number of self-declared

scientific realists would not be prepared to accept them all. Brian Ellis, for example, has argued that 'scientific realism is incompatible with any form of the correspondence theory of truth' (Ellis 1990, 160 f.), because the ontology of science cannot possibly accommodate the entities to which that theory assigns the role of the *bearers* of truth (Ellis 1990, ch. 5). In the opposite camp, Bas van Fraassen, whose 'constructive empiricism' provides one of the most compelling forms of anti-realism currently available, appears non-equivocally committed to the correspondence theory of truth (van Fraassen 1980, 90, 197). Indeed, were it not for his commitment to the correspondence theory of truth, his argument could not even establish its point.

I am ready to admit that scientific realism as I have proposed to define it, that is, as a purely metaphysical claim, is not a very interesting position. While Realism would be an interesting thesis even if our singular claims about the tokens which fall under the kinds we take to be instantiated in the furniture of the world were mostly wrong, the same is not true when applied to scientific realism. If there were in fact a few knowledge-independent electrons around, but the singular claims about electrons implied by our present scientific information were mostly wrong, then our scientific knowledge would be remarkably poor. So I am keen to concede that any interesting formulation of scientific realism will have to involve a claim about the approximate correctness of our theoretical picture of electrons (or, as Ian Hacking has suggested, about our knowing how to use electrons to bring about various kinds of physical phenomena). That is not to say, however, that any interesting formulation of scientific realism will have to consist in some sort of *convergent* realism. The statement that the singular claims implied by our *present* scientific information are largely correct will satisfactorily do the job. Popper's original 'conjectural realism' (see Worrall 1989, 110 f.) provides an example of a resolute (if somehow qualified) scientific realism not committed to the increasing verisimilitude of the theories of mature science.

What I am most interested in, though, is a clear distinction between (1) and (2). In section 1.1, I argued that one is allowed to formulate non-metaphorical and perfectly meaningful *metaphysical* theses about what there is. What I now want to emphasise is that current arguments in favour of scientific realism do not provide any grounds to believe that our formulation of its ontological claim is hollow unless it is supplemented by the

acceptance of the correspondence theory of truth. Our next step will be an exploration of the links between Realism and different theories of truth.

1.4 External World Realism and Truth

External world realism is usually taken to bear a privileged relationship to the correspondence theory of truth. So here is my definition of correspondence truth, which is a variant of Devitt's definition (see Devitt 1991, 29):

Correspondence Truth: Sentences of type X are true or false in virtue of: (1) their structure; (2) the referential relations between their parts and reality; (3) how reality is independently of anybody's knowledge¹⁰.

I wish to emphasise that this definition does not require that the sentences in question 'mirror', or 'picture', or be somehow 'isomorphic' to, the bit of reality they purport to represent.

Several contemporary upholders of correspondence truth are convinced that Tarski's semantic definition has succeeded in rehabilitating the traditional view of truth as correspondence (see, e.g., Popper 1963, 223 f. and Zahar 1984, 165). This is controversial, and Tarski's own remark that 'we may accept the semantic conception of truth without giving up any epistemological attitude we may have had; we may remain naive realists, critical realists or idealists, empiricists or metaphysicians - whatever we were before' (Tarski 1944, 362) is frequently taken to undermine that conviction.

If we stick to a clear distinction between the metaphysical issue of what there is and the epistemological issue of the nature of truth, Tarski's definition is indeed ontologically neutral, but nevertheless provides some form of support for correspondentism. It is true that Tarski's definition does not specify what *kinds* of entities

¹⁰ Of course there is a sense in which 'how reality is' depends on somebody's knowledge. If S knows that p, then 'p' is true, then p. But saying on this basis that how reality is depends on somebody's knowledge would be like saying that the pressure depends on the barometer reading.

Perhaps I should also specify that when I write that 'how reality is' must be independent of *anybody's* knowledge, I take that to mean that it must be independent from the knowledge of any empirical, transcendental, idealized, or God-like subject.

must exist so that the truth-value of the sentences in our formalized language may be determined. Thus Tarski's definition is compatible with different ontologies (e.g., with phenomenalism), and does not require the existence of an external, physical world. But it is equally true that satisfaction is a semantic relation obtaining (or failing to obtain) between open sentences and sequences of objects. *If* the objects involved are understood as knowledge-independent, *then* Tarski's definition provides a perfectly clear sense of how a sentence in a formalized language can *correspond* to reality¹¹. So Tarski's definition may be construed as a version of the correspondence theory of truth, even if it doesn't necessarily need to.

However, the claim that Tarski's definition provides significant support for the correspondence theory of truth ought not to be misunderstood. First, it does not entail the further claim that the facts described by the sentences on the right-hand-side of any given (T) schema are the knowledge-independent entities that make the sentences named on the left-hand-side either true or false. Hence, it does not support the thesis that a sentence is true iff there is some knowledge-independent fact to which it is isomorphic. But that thesis is not part of our definition of correspondence truth. Indeed, for someone upholding correspondence truth, facts needn't even be conceived as knowledge-independent (I believe they should be conceived as being, at least in part, linguistic constructs. In my opinion, facts supervene on reality *under a linguistic description*. For this reason, I prefer to say that a sentence is true iff it corresponds to reality, or iff it

¹¹ The fact that 'satisfaction' is *prima facie* a semantic notion is the source of some trouble for the attempt to attach philosophical significance to Tarski's definition of truth.

Can the objects in the sequences be understood as knowledge-independent? It has been pointed out that Tarski's enumerative definition of satisfaction for atomic open sentences fails to provide a truly general definition of satisfaction, as opposed, say, to different definitions for satisfaction-in- L_1 , satisfaction-in- L_2 , etc. Tarski's definition may thus appear to embody an inadequately analysed semantic primitive, whose presence may be seen as casting doubts upon Tarski's own claim of providing a philosophically unobjectionable analysis of something like 'the classical Aristotelian conception of truth' (Tarski 1944, 342). Hartry Field has argued that Tarski's definition must be complemented by a physicalist theory of reference if it is to give a satisfactory explanation of the connection between language and (extralinguistic) reality (see Field 1972).

To my mind, it is true that the lack of a detailed theory of reference showing how language may latch on to a knowledge-independent reality may encourage people with anti-realist leanings to think that reference is always internal to our conceptual schemes. But the Realist cannot be forced to give up his (commonsensical) idea that names refer to knowledge-independent objects on the sole basis that she has not yet provided a fully worked-out theory of how that may happen. So Tarski's theory of truth does help to clarify how a sentence in a formalised language can *correspond* to reality, even though it does not clarify how an object can *satisfy* a predicate.

corresponds to how the world is, and not iff it corresponds to the facts. That is not to say that a sentence cannot fail to correspond to the facts; but that failure will be just a consequence of its failure to correspond to reality).

One ought not to overestimate the *epistemological* relevance of Tarski's definition either. For indeed that definition provides a compelling analysis of what it is for a sentence in a formalized language to be true or false, but it doesn't say anything at all about the problem of the epistemic accessibility of truth. And this very problem has been the focus of many recent attacks against the correspondence notion of truth, attacks which deny the epistemological relevance of the notion, rather than its formal correctness. The development and assessment of a counter-argument aimed at the rejection of precisely this sort of epistemological attack will provide one of the major focuses of the present work.

But let's go back to our problem: What are the links between Realism and different theories of truth? First of all, note that the issues of Realism and truth appear logically independent from each other. Take the case of correspondence truth, which is the most likely candidate to a close relationship with Realism. (Note that correspondence truth is not synonymous with evidence-transcendent truth. As we remarked before, there can be versions of coherentism according to which truth is in a strong sense evidence-transcendent, but a coherentist account of truth is certainly not a correspondence one). Does Realism entail the correspondence theory of truth? According to our definition, it's clear that it doesn't. As for the converse entailment, we have already seen that Tarski's theory is ontologically neutral, even if it does lend some form of support to the correspondence theory of truth. The correspondence theory of truth, in turn, entails some form of realism, but not necessarily external world realism. A phenomenalist who rejected the incorrigibility and the self-intimation theses could subscribe to the correspondence theory of truth.

So we shall have to look for weaker links. Devitt has suggested that an abductive argument can be devised from Realism to correspondence truth. In short, his point is that 'from a Realist perspective, we need truth to explain the properties of symbols that enable them to play their wide variety of social roles' (Devitt 1991, 44; the argument is spelled out in detail in ch. 6). Devitt's argument is highly controversial, and for any sympathy I may have for it, I won't strive to defend it in these pages. On the contrary, I will dwell

on the other abductive argument to which Devitt draws our attention, the argument *from epistemic truth to anti-realism* (see Devitt 1991, 44 f.).

For explanatory purposes, an epistemic theory of truth can be contrasted with an absolutist one. An absolutist theory of truth maintains that the truth-value of a given statement is altogether independent of anybody's epistemic situation and thus a completely objective matter. On the contrary, an epistemic theory will maintain that truth is always a function of some epistemic concept, such as, for example, verifiability, rational acceptability, warranted assertability, and so on. A recent example of an epistemic definition of truth is provided by Hilary Putnam, who has defined truth as an idealization of rational acceptability, claiming that true is what would be rational to accept in an ideal epistemic situation (Putnam 1981, ch. 3¹²; but the idea goes back at least to Peirce¹³).

Devitt argues that the closest link between the epistemic doctrines of truth and Realism is displayed by an application of Tarski's material requirement. His abductive argument from epistemic truth to anti-realism goes as follows. Assuming that any acceptable definition of truth should have as a consequence all instances of the (T) schema, then, if 'T_E' stands for 'true as defined by the epistemic doctrine E', E will require that the appropriate instances of

S is T_E iff p

hold. For example:

'Caesar had five moles' is T_E iff Caesar had five moles

¹² Nobody aware of the fluidity of Putnam's thought will be surprised to hear that he has recently modified this position by suggesting that not only truth depends on rational acceptability, but rational acceptability depends on truth: 'whether an epistemic situation is any good or not typically depends on whether many different statements are *true*' (Putnam 1988, 115). Putnam claims that this is precisely what he had in mind when he wrote *Reason, Truth and History*, though the textual evidence shows that this is not what he wrote.

¹³ To be sure, Peirce also thought that scientific method is somehow *constrained by reality*, and hence cannot be read as abandoning correspondence truth altogether.

But why should such instances of our modified (T) schema hold? Devitt invites us to consider an extravagant version of epistemic theory of truth: $T_E =$ is affirmed by the Pope. Why should

‘Caesar had five moles’ is affirmed by the Pope iff Caesar had five moles

hold? An explanation can only be provided by some epistemological doctrine. We could say, for example, that the world is created by the Pope’s word, or that the Pope enjoys some infallible insight of a divine origin into how the world is. These two answers lead to different attitudes about the knowledge-independent existence of the external world, respectively an anti-realist and a realist one. The epistemic doctrine that ‘to be true is to be affirmed by the Pope’ has some bearing on the issue of realism only *via* some epistemological theory. But while in the case of our papist doctrine we can imagine that some neo-scholastic philosopher may find the epistemological theory that leads to Realism more attractive than the epistemological theory that leads to a Pope-creator, in the case of the epistemic doctrines currently advanced in the epistemological circles the situation is, for any sort of philosopher, completely reversed:

The problem which the epistemic doctrine poses for the Realist is that it is hard to find a plausible Realist epistemology to do the explanatory job for most, if not all, the likely candidates to be T_E .

(Devitt 1991, 45)

If, for example, T_E were to be construed in terms of warranted assertability, the external world realist could hardly devise an explanation of the incredibly close link between a state of the world, Caesar having five moles, and our being warranted in asserting ‘Caesar had five moles’, that the epistemic doctrine in question would commit her to. For which epistemological doctrine assuming the knowledge-independence of the external world could ever explain why those statements which are warrantably assertible are infallibly true? And surely the external world realist couldn’t do any better when confronted with other epistemic versions of truth.

In sum, even if epistemic truth does not logically imply the denial of the existence of the external world, it is nevertheless likely to lead to anti-Realism for lack of viable

epistemological explanations of why the (T) schemas that follow from a particular choice of T_E should hold on the basis of a realist view of the external world.

The strong appeal of what Devitt calls the abductive argument from epistemic truth to anti-Realism helps to explain why the issues of Realism and truth are so frequently conflated in current philosophical reflection. Acceptance of an epistemic doctrine of truth provides a strong motivation for the rejection of Realism. Given the minor popularity enjoyed by the deflationary (or redundancy, or disquotational) account of truth (but see the forceful defence of the 'minimalist' doctrine in Horwich 1990), many philosophers feel as if they are forced to pick out just one of the two following options: either *Realism + correspondence truth* or *anti-Realism + epistemic truth*. It takes just a further, little step to suppose that Realism and anti-Realism are in fact doctrines about the nature of truth. Which is false, but nevertheless contains a grain of truth. Put in a nutshell: if one wants to be a Realist but is not willing to endorse a deflationary account of truth, she had better prepare herself to defend the correspondence theory of truth from the attacks of the advocates of the epistemic doctrines.

1.5 Correspondence Truth and the God's Eye Point of View

'How can I see that my knowledge corresponds with the object? - I only can compare the object to the extent I know it. I only can compare knowledge of the object with other knowledge of the same object'. This passage from Kant (quoted in Bonsack 1989, 80) provides a nice example of the sort of convictions underlying most current epistemic attacks on the correspondence theory of truth. Kant himself did not go so far as to reject correspondence truth explicitly (see Kant 1933, 97), yet his epistemology is clearly incompatible with any notion that could satisfy our third requirement for correspondence truth, because his phenomenal world (that is, the only world epistemically accessible to the human beings) is structured by the categories of the Understanding.

Present-day heirs of Kant usually provide a motivation and an argument in support of their rejection of correspondence truth. While the argument is developed in different forms, the motivation is pretty much the same in all cases. Critics of correspondence truth are invariably prompted by the desire to bridge the epistemic gap between evidence

and truth they take as a consequence of the correspondence theory of truth. Actually, the correspondence theory of truth does not entail, by itself, any claim about the epistemic accessibility of truth. Correspondence truth supplemented by a dogmatic epistemology like that sketched in section 1.1, for example, does not imply any epistemic gap at all between evidence and truth. There can be no doubt, however, that correspondence truth makes an epistemic gap between evidence and truth at least *possible*. But while upholders of correspondence truth address the problem of that gap only after their notion of truth is already in place, upholders of epistemic doctrines build the solution to the problem into their definition of truth from the outset.

Even if the arguments purporting to rebut correspondence truth exhibit a great variety of formulations, most of them draw on a basic idea which is shared by all their proponents. Put bluntly, this is the conviction that a concept which is forever going to transcend our capacity of rational recognition is idle, and cannot play any significant role in our epistemology or in any other part of our cognitive picture. It goes without saying that the rejection of correspondence truth is normally prompted by the belief that the notion of a correspondence between linguistic entities and reality must indeed be such a concept, because of the epistemic gap it makes possible between evidence and truth.

Two well-known versions of the argument from the idleness of correspondence truth are those developed by Michael Dummett and Hilary Putnam.

Dummett's argument from the acquisition and manifestation of linguistic competence purports to show that no truth-conditional theory of meaning committed to the existence of recognition-transcendent truths will succeed in providing an adequate account of the propositional knowledge implicitly possessed by the speakers of some language L (see Dummett 1976). Anyone accepting Dummett's semantic characterization of the realism/anti-realism debate will regard his argument, if successful, as an immediate refutation of realism. But Dummett's argument doesn't require that the semantic approach be accepted to establish its point. For its thrust is that acceptance of an evidence-transcendent notion of truth prevents the formulation of an acceptable theory of meaning. So, if Dummett is right, we shall be compelled to give up correspondence truth¹⁴, endorse epistemic truth, and accept Devitt's abduction from epistemic truth to anti-

¹⁴ Unless we are prepared to subscribe to a dogmatic epistemology capable of making correspondence truth generally accessible to human cognitive subjects.

Realism. However, it should be noted that the rejection of the semantic approach has the effect of restoring the issue of the evidence-transcendence of truth to its natural place, i.e., epistemological theory. Whether the possibility of an epistemic gap between evidence and truth allowed by correspondence truth has any relevance for the possibility of a theory of meaning cannot be determined by our theory of truth alone: an answer to the question requires the decisive contribution of epistemological theory. In other terms: while acceptance of correspondence truth implies commitment to the possibility of truth transcending evidence, only acceptance of an epistemological theory will imply commitment to a definite view of the relationship between evidence and truth, which is what actually affects the construction of a theory of meaning.

Dummett claims that no account of the manifestation and acquisition of linguistic competence can possibly be developed on the basis of a recognition-transcendent notion of truth. Acquisition of linguistic competence presupposes its manifestation: nobody could learn to speak a language if what speakers know in knowing their language were not publicly observable, hence acquirable in public contexts. But

if truth-value is construed as a possibly recognition-transcendent property of sentences, as in the realist view it is, then what account is to be given of what it is to know the truth-conditions of sentences whose truth-values we are not able to establish? It is plainly impossible to associate grasp of the transcendent truth-conditions of certain sentences with the possession of an ability to recognise what their truth-values are, precisely because their truth-conditions are *transcendent*. For such sentences, then, we have no way of saying how knowledge of their truth-conditions can be manifested. But if we cannot say this, the theory does not show how sense and use connect - how, that is, sense and use determine each other. Accordingly any theory of meaning based on a transcendent concept, as in the case of a realist theory of meaning, is useless.

(Grayling 1990, 235)

Dummett's argument can be construed as an argument 'from the idleness of correspondence truth' because it claims that an evidence-transcendent notion of truth (such as correspondence truth turns out to be, if it is not complemented by a dogmatic epistemology) cannot give any contribution to our understanding of how linguistic competence is manifested and acquired. Any appeal to that notion made by truth-conditional theories of meaning is thus to be regarded as completely otiose.

If we do not readily accept this conclusion, Dummett urges, it is because we convince ourselves that understanding of the truth-conditions of undecidable sentences consists in our grasp of what the ability of using such sentences to give direct reports of observation would be for someone who had epistemic access to the God's eye point of view:

We cannot do this; but we know just what powers a superhuman observer would have to have in order to be able to do it - a hypothetical being for whom the sentences in question would *not* be undecidable.

(Dummett 1976, 99)

The idea is forcefully conveyed in the course of Dummett's attack against realism about the past (which is rendered as the claim that any arbitrary statement about the past is determinately true or false independently of our present, future, or possible knowledge of its truth-value):

What the realist would like to do is to stand in thought outside the whole temporal process and describe the world from a point which has no temporal position at all, but surveys all temporal positions in a single glance [...] The anti-realist takes more seriously the fact that we are immersed in time: being so immersed, we cannot frame any description of the world as it would appear to one who was not in time, but we can only describe it as it is, i.e., as it is now.

(Dummett 1978, 369)

Such reflexions lead Dummett to suggest that ultimately 'realism is tenable only on a theistic basis' (Dummett 1978, xxxix).

In the present work, I won't produce a detailed criticism of Dummett's anti-realist argument (for which purpose, see McDowell 1978 and Devitt 1991, ch. 14). Nor will I dwell on his proposal of an anti-realist theory of meaning developed in terms of assertability-conditions. What I am concerned with is rather his assumption that correspondence truth could play some role in our theory of meaning only if we were willing to endorse the (indefensible) thesis of the epistemic accessibility of the God's eye point of view.

A similar line of thought can be found in Putnam 1981, the main difference being that Putnam's argument deals with epistemology rather than the theory of meaning (to be

sure, Putnam does have an anti-realist argument based on the theory of reference as well; but that's a different story).

Putnam is even more explicit than Dummett in his use of the theological metaphor of the God's eye point of view. He proposes to distinguish 'two philosophical perspectives' (Putnam 1981, ch. 3), which he calls the 'externalist' and the 'internalist' perspective. The former is also characterized as 'metaphysical realism', that is, as Realism + correspondence truth + the thesis that there is exactly one true and complete description of 'the way the world is'. And Putnam calls it the *externalist* perspective exactly because 'its favorite point of view is a God's Eye point of view' (Putnam 1981, 49).

It is apparent that Putnam's definition of the externalist perspective identifies a very demanding doctrine. In fact, he conflates correspondence truth with the 'One True Theory' claim. But correspondence truth entails the 'One True Theory' claim only if the correspondence between language and world is understood as an isomorphism between sentences and facts. And we saw in the last section that isomorphism between sentences and facts is not an essential feature of an adequate notion of correspondence truth.

An *internalist* perspective is described by Putnam as dismissing any talk of reality implying an 'external' point of view:

There is no God's Eye Point of View that we can know or usefully imagine, there are only the various points of view of actual persons reflecting various interests and purposes that their descriptions and theories subserve.

(Putnam 1981, 50)

According to an internalist perspective, asking *What objects does the world consist of?* only makes sense *within* a theory or description. Truth cannot be a correspondence between our statements and a ready-made world: if truth is to be accessible, it must be understood in an epistemic fashion.

Putnam's argument against externalism can be put quite straightforwardly. To ascertain that a given statement is correspondence-true, one ought to have access both to her description of the world *and* to the world as it is in itself, in order to compare their features and determine if they fit with each other. But to assume that one could have

access to the world as it is in itself is to assume that one could place oneself in the God's eye point of view. However, it is widely agreed that our experience of the world is always theory-laden, and that its content is always shaped, to some extent, by our conceptual schemes. So, Putnam concludes, if our epistemic perspective is always determined by our conceptual choices, we will never be in a position to ascertain that a given statement is correspondence-true. As access to a God's eye view is precluded to human beings, truth will be for ever inaccessible to their cognitive efforts.

Even if Putnam presents his argument as a refutation of externalism, his point is also relevant for the minimal version of correspondence truth we are primarily concerned with. Taking truth to be some form of correspondence between linguistic entities and the knowledge-independent world does raise the epistemological problem of the accessibility of that world to human knowledge. So we may agree to construe Putnam's argument as an argument against correspondence truth, rather than as an argument against the externalist perspectives. However, Putnam's point is hardly original. Putnam himself traces its origins back to Kant. Nevertheless, his argument is representative of a widespread belief that acceptance of an evidence-transcendent notion of truth will make knowledge impossible. More or less radical versions of the argument are to be found in the writings, for example, of Richard Rorty, Arthur Fine, David Bloor and Brian Ellis. But Putnam's use of the metaphor of the God's eye point of view makes his own version particularly vivid.

To sum up: both Dummett and Putnam claim that anyone upholding correspondence truth is implicitly committed to the indefensible thesis of the epistemic accessibility of the God's eye point of view. They argue that if a God's eye view were accessible to human beings, then correspondence truth could play a role in our theory of meaning (Dummett) or in our epistemology (Putnam). But since that is not the case, correspondence truth remains a totally idle concept, and we had better settle for something more ready-to-hand.

Putnam explicitly equates his own internalist perspective to Dummett's anti-realism (see Putnam 1981, 56). Yet, when we fix our attention on the issue of the external, physical world, these two philosophers seem to have significantly different ideas about what would count as a God's eye view. While Dummett is concerned with the comprehensiveness of the evidence provided by a God's eye view, Putnam suggests that

access to a God's eye point of view would yield knowledge of the noumenal world. To my mind, this difference can be explained by reference to Dummett's residual verificationism about observational beliefs (see, e.g., Dummett 1978, 158) on the one hand, and to Putnam's open Kantianism (see Putnam 1981, 60-64) on the other. According to Putnam, the objects of our discourse are always internal to some conceptual scheme, yet they somehow point to the existence of a deeper, unknowable reality, i.e., that noumenal world that provides 'a mind-independent "ground" for our experience' (Putnam 1981, 62). On the contrary, Dummett's anti-realism seems committed to deny any deeper dimension to the world: knowledge from the God's eye point of view would not unveil a deeper reality lying beneath the appearances, but would simply provide a more extensive information about those appearances themselves.

If my interpretation is correct, Dummett and Putnam are concerned with two different species of evidence-transcendence, respectively an 'horizontal' and a 'vertical' one. In the remaining part of this work, I shall try to address the challenge arising from Putnam's 'vertical' species of evidence-transcendence, rather than that arising from Dummett's 'horizontal' species. The sort of argument I wish to reject is an *epistemological* argument from the (alleged) epistemic inaccessibility of reality in itself to the idleness of correspondence truth.

What I wish to argue for is the thesis that acceptance of correspondence-truth does not make the claim that we can *know* some features of the Realist's world unintelligible. In Kantian terms, I wish to argue that acceptance of correspondence truth needn't confine us to phenomenal knowledge: if there is anything that we know, that is reality as it is in itself. For knowledge of reality as it is in itself needn't be knowledge from a God's eye point of view.

Realism Without a God's Eye View

2.1 The Quest for Certainty

Most traditional epistemologies have taken it for granted that objectivity and certainty are closely related to each other. In fact many writers have regarded certainty as a sign of objectivity, and developed their epistemologies in accordance with that view.

In order to assess such an epistemological project, we need to answer a preliminary question. Why do people aspire after objectivity in their cognitive efforts? What is it that makes objectivity a virtue?

Claims to objectivity are usually made when the epistemic status of a certain statement is challenged by someone who does not believe in its truth. As far as I see, if S claims that p is *objectively* true, what she is actually claiming is that p is true *whatever opinion any particular subject* (and specifically her objector) *has about p*.

Suppose for a moment that Locke got his account of the primary and secondary qualities of things right. Then S cannot be warranted to claim that London's buses are *objectively* red, although she can be warranted to claim that London's buses have *objectively* such-and-such a shape. From this perspective, the search for objectivity appears to be the search for statements which are true of their object, not of the way that object happens to be perceived, known, or taken to be by some cognitive subject. This way of looking at the matter helps with the following difficulty. Someone may wish to point out that saying of a given sentence that it is objectively true does not add anything to saying of that same sentence that it is true. After all, one and the same sentence (read 'proposition', if you like) cannot be true for X and false for Y! So what is the point of all this talk of objectivity? Lacking relativistic leanings, I am entirely sympathetic with

this kind of objection. Yet, saying of a sentence that it is *objectively* true can be more than a rhetorical device. It can be a way to specify the subject-matter of that sentence. We say that the sentences about the secondary qualities of things are subjective not because they cannot be objectively true of their own subject-matter, but because their subject-matter is not restricted to the 'objective' physical world, but includes the 'subjective' perceptual apparatus of the human beings. So saying that in our cognitive efforts we aim to come up with sentences that are *objectively* true may be a way of saying that we aim to come up with sentences whose subject-matter is the world as it is in itself, rather than as it is in relation to our perceptual apparatus¹⁵.

Those philosophers who endorse an epistemic theory of truth usually claim that their theories are able to accommodate a satisfactory degree of objectivity within their framework. Such philosophers merely have to introduce some distinction between empirical and transcendental subjects, or between actual and ideal epistemic situations. That done, they can say of whatever statement they like that its truth is objective, i.e., independent of the knowledge of any *empirical* subject, or independent of the cognitive outcome of any *actual* epistemic situation.

I am not sympathetic with such carefully carapaced accounts of epistemic objectivity. Nevertheless, these accounts show how deeply entrenched is the belief that knowledge must ultimately aspire after some form of objectivity. It can hardly be denied that any knowledge *claim* involves a claim to inter-subjective validity. Roughly, if S claims to know that p, her claim will entail, among other things, that anyone presented with the relevant evidence ought rationally to believe that p. Thus Kant wrote:

Persuasion is a mere illusion, because the ground of the judgment, which lies solely in the subject, is regarded as objective. Such a judgment has only private validity, and the holding of it to be true does not allow of being communicated.

I cannot *assert* anything, that is, declare it to be a judgement necessarily valid for everyone, save as it gives rise to conviction [in any rational subject].

(Kant 1933, 645 f.)

¹⁵ Of course it is wholly legitimate to set out in search of objective truths about subjective phenomena, as various kinds of phenomenological and psychological disciplines try to do.

One may well debate how the class of the cognitive subjects to which the 'ought' applies can best be defined, but it is hardly deniable that a claim to some sort of inter-subjective validity is included in any knowledge claim¹⁶. But then, if S claims to know that p, she must be claiming that p is *objectively* true, that is, that the truth-value of p is invariant with respect to any member of the class of the relevant cognitive subjects. This suggests a possible explanation of the fact that objectivity, for all the qualifications one may wish to apply to it, has been traditionally considered an epistemic virtue.

Many philosophers have thought that the best way to gain objectively true beliefs (that is, beliefs in sentences/propositions which are objectively true) is the pursuit of certainty. If I achieve certainty about a given belief, I will have secured its truth¹⁷. But since what is true (relative to the class of the relevant epistemic subjects) must be true for everyone (in that class), I will have secured objectivity as well. So, goes the advice, look for certainty, and you will get objectivity as well.

The quest for certainty lies at the core of the most important research programme in modern epistemology - the programme of 'internalist' epistemology, as it will be called in section 2.2. Starting with Descartes, modern epistemologists have usually sought to provide incontrovertible foundations for our knowledge. The certainty of those foundations has been thought to secure their (objective) truth as well as that of their consequences, and hence to support the claim to inter-subjective validity of our putative knowledge.

This is how Descartes presented his epistemological project in a famous passage from his *Meditations on First Philosophy*:

Anything which admits of the slightest doubt I will set aside just as if I had found it to be wholly false; and I will proceed in this way until I recognize something

¹⁶ If the claim is a claim about the 'secondary' qualities of things, it will have to be construed as an inter-subjective claim about subjective phenomena. The claim that London's buses are red, for example, will have to be construed as involving the claim that any rational subject ought to believe that London's buses *look* red to normal human beings.

¹⁷ In this context, 'certainty' is not to be construed as a merely psychological feeling of conviction, but, roughly, as an objective property of the content of a certain belief. Furthermore, one must be very careful about what exactly can be warrantably described as 'certain': I may be certain that X looks red to me, but of course that X is red-relative-to-me does not entail that X is red-relative-to-a-bat!

certain, or, if nothing else, until I at least recognize for certain that there is no certainty.

(Descartes 1984-85, II, 16)

He eventually thought to have found what he was looking for in the proposition, *I am, I exist*, which he took to be self-evident and then used as the uncontroversible foundation of his epistemological construction¹⁸. As one would expect, in order to carry out his project he had to appeal to a bridge-principle saying that whatever is perceived very clearly and distinctly - i.e. with the marks of certainty - must be true (Descartes 1984-85, II, 24), that is, *objectively* true. So Descartes can be fairly regarded as a paradigmatic example of that kind of epistemological strategy that makes of certainty, conceived as an hallmark of objective truth, an essential feature of knowledge.

Of course one needs to be clear about what 'certainty' actually means. Certainty as a psychological feeling of conviction - 'I can't help believing that p' - will not secure the objectivity of p. Philosophers within the 'Cartesian' tradition have thus endeavoured to come up with non-psychologistic explications of the meaning of 'certainty' (as we noticed, Descartes' idea was to spell out certainty in terms of clarity and distinction). But for our purposes we needn't discuss all their proposals in detail. One proposal which is worth discussing, though, can be found in Husserl's *Logical Investigations*.

Husserl's position is particularly relevant to our topic as it provides an excellent example of an epistemology espousing both correspondence truth and the view that correspondence truth involves commitment to the epistemic accessibility of the God's eye point of view (Husserl 1970, II, 760-770).

Husserl lists four different meanings of 'truth', but he regards as basic the concept of truth as 'the *adaequatio rei ac intellectus*' (Husserl 1970, II, 670). According to Husserl, truth makes itself present in that kind of objectifying act in which 'the object is not merely meant, but in the strictest sense *given*, and given as it is meant' (Husserl

¹⁸ '*I am, I exist*' is the proposition Descartes claims to be necessarily true ('whenever it is put forward by me or conceived by my mind') in the Second Meditation. This is of course the *cogito* argument, but Descartes himself warns his reader against an easy misconstruction of his formula, *cogito, ergo sum*: 'When someone says "I am thinking, therefore I am, or I exist", he does not deduce existence from thought by means of a syllogism, but recognizes it as something self-evident by a simple intuition of the mind' (Descartes 1984-85, II 100).

1970, II, 765). This quotation reflects Husserl's own distinctive terminology. Note first of all that 'object' is used in a very broad sense, including states of affairs as well. Husserl conceives truth as the objective correlate of a psychological act in which a meaning-intention comes to perceptual fulfilment. Note that he does not speak of sentences, but of meaning-intentions. Meaning-intentions are temporally located psychological acts, but Husserl is well aware of the dangers of psychologism, and describes their contents, in a Fregean vein, as 'the self-identical meaning that the hearer can grasp even if he is not a percipient' (Husserl 1970, I, 290). When the cognitive subject is not facing the intended object, her meaning-intention is a *mere* meaning-intention, that is, void of intuition. But when the intended object comes before the subject, her meaning-intention may be fulfilled, if the object is *given* to her by perception (which Husserl understands in a suitably liberalized fashion) in exactly the same way as it is meant. If so much happens, Husserl says that an 'identifying act' takes place, which has as its object 'the *full agreement* of what is meant with what is *given as such*', that is, '*being in the sense of truth, or simply truth*' (Husserl 1970, II, 765).

What is particularly relevant to our inquiry is Husserl's view of the relation between truth and self-evidence:

This agreement [i.e., truth] we *experience* in self-evidence, in so far as self-evidence means the actual carrying out of an adequate identification.
(Husserl 1970, II, 765)

According to Husserl, whenever we run up against truth we also have the possibility of laying it before our consciousness:

Truth is indeed 'present'. Here we have always the *a priori* possibility of looking towards this agreement, and of laying it before our intentional consciousness in an adequate percept.
(Husserl 1970, II, 766)

When a meaning-intention is given the fulness of the object itself, 'the *adaequatio rei ac intellectus* [...] is itself given, to be directly seized and gazed upon' (Husserl 1970, II, 670).

Husserl's characterization of 'self-evidence' can be seen as an attempt to provide a non-psychologistic explication of the notion of 'certainty'. It is clear that his concept of 'self-evidence' precludes any interpretation of certainty as a mere 'feeling' contingently attached to the act of judgement (see Husserl 1970, II, 769). The relation between truth and certainty becomes so intimate that Husserl is eventually led to identify knowledge with certainty:

The synthesis of fulfilment achieved in this limiting case [the case of an objectively complete adequacy of the meaning-intention to the object itself] is *self-evidence or knowledge in the pregnant sense of the word*.

(Husserl 1970, II, 670)

For Husserl certainty is not merely desirable as a sufficient condition for objectivity, nor is it simply an essential feature of knowledge: certainty is knowledge itself, *in the pregnant sense of the word*. Truth cannot be known without the *stigmata* of certainty.

What I am most interested in emphasizing, though, is Husserl's assumption that acceptance of correspondence truth involves commitment to the epistemic accessibility of the God's eye point of view. Of course he does not use these words, but what else could he mean when he says that in the synthesis of fulfilment 'the object is not merely meant, but in the strictest sense *given*, and given as it is meant'? To be sure, Husserl allows for increasing degrees of intuitive fulfilment. Yet he says that the experience of fulfilment 'is represented by the words: "This is the thing *itself*"' (Husserl 1970, II, 720). If truth is *adaequatio rei ac intellectus* - so goes Husserl's assumption - knowledge at its best must be the self-evidence of that *adaequatio*. Husserl's contention is that knowledge in this sense is indeed possible, because in the synthesis of fulfilment 'this *adaequatio* is itself given, to be directly seized and gazed upon' by the cognitive subject. This is tantamount to saying that the cognitive subject has epistemic access, in certain circumstances, both to her description of the world *and* to the world as it is in itself. Moreover, Husserl tells us that she can compare their features, and also that, when the object is given as it is meant, she can perceive this agreement in the experience of certainty.

The textual evidence of the *Logical Investigations* allows us to regard Husserl's view as a striking illustration of exactly that kind of epistemology that Putnam strives to

reject in his attack against 'metaphysical realism'. On the other hand, Husserl's view also provides a striking illustration of that kind of epistemology that I shall argue a Realist needn't be committed to.

2.2 Three Kinds of Internalist (Externalist) Doctrines

In section 1.5 we came across Putnam's particular version of the internalism/externalism dichotomy. But there are at least two other versions of that dichotomy which are frequently referred to in the present philosophical debate.

Putnam's own version had to do with the issue of 'metaphysical realism', which he rendered as the conjunction of Realism, correspondence truth, and the thesis that there is exactly one true and complete description of the way the world is. He called the doctrines falling under this composite description 'externalist' doctrines because their favourite point of view could be described as God's eye point of view. For the sake of terminological simplicity, I shall call this version of the dichotomy the *metaphysical* version.

A different version of the internalism/externalism dichotomy is also due to Putnam (see Putnam 1975, 223-227), although he himself didn't label it that way (and neither did the other father of the doctrine, Tyler Burge, in his 1979). This second version belongs to the *philosophy of mind*. Roughly, the issue at stake is this: What determines the nature of intentional states (or, if you like, of propositional attitudes)? Internalist theories of mind claim that the nature of intentional states is entirely determined by those factors which are 'internal' to the subject. Externalist theories claim that the nature (e.g., the reference) of intentional states is affected by such 'external' factors as the nature of the subject's environment. Putnam's Twin Earth argument was originally put forward as a proof of the view that "meanings" just ain't in the *head*' (Putnam 1975, 227). So Putnam, while an internalist on the issue of realism, is to be counted as an externalist in the philosophy of mind. However, we needn't spell out this issue in detail. What really affects our topic is rather the third version of the dichotomy, to which we now turn our attention.

The third version of the internalism/externalism dichotomy, which has been the focus of extensive philosophical debate in the last decade or so, arises from *epistemological* concerns. In fact, it is closely linked to the analysis of knowledge and/or justification.

An analysis of 'S knows that p' is said to be 'internalist' when the conditions that must be satisfied if S is to know that p are 'internal' to S's awareness, that is, when they are (to some extent) epistemically accessible to S. An analysis of 'S is justified in believing that p' is said to be 'internalist' when the grounds that make S's belief justified are 'internal' to S's awareness, that is, when they are (to some extent) epistemically accessible to S.

Of course one needs an explanation of what is meant by 'internal to one's awareness' and 'epistemically accessible'. Most of all, one needs an explanation of *how* epistemically accessible the relevant conditions must be in order to justify a claim to knowledge (or to epistemic justification). Does S need to know it for certain that those conditions are satisfied? Or does she merely need to have access to some evidence that they are? Different kinds of internalist epistemologies will issue from different answers to these questions.

Take the traditional analysis of knowledge as justified true belief:

S knows that p	iff	1) p is true
		2) S believes that p
		3) S is justified in believing that p

Is this analysis internalist or externalist? That depends (1) on our concept of justification and (2) on our concept of epistemic accessibility. If our concept of 'epistemic accessibility' is not too demanding (so that satisfaction of condition 3 may be seen as providing epistemic access to the satisfaction of condition 1), and if our concept of justification is itself internalist (so that S may be thought to have epistemic access to the satisfaction of 3), this analysis can be described as 'internalist' (on the further, reasonable assumption that S has epistemic access to the satisfaction of condition 2). On the contrary, if (a) our concept of epistemic accessibility is such that nothing can count as epistemically accessible which is not known for certain, or if (b) our concept of

justification is externalist, the JTB account of knowledge will come out as externalist. If our concept of epistemic accessibility requires certainty, satisfaction of condition 3 will not provide epistemic access to the satisfaction of condition 1 (unless we commit ourselves to the rather implausible assumption that one cannot justifiedly believe a false proposition). If our concept of justification is externalist, S will not have epistemic access to the satisfaction of condition 3.

A much sharper example of an externalist analysis of knowledge is provided by the epistemological doctrine which goes under the name of 'reliabilism':

S knows that p	iff	1) p is true
		2) p believes that p
		3) S came to believe that p by means of a reliable belief-forming process

Here it is clear that conditions 1 to 3 may be satisfied without S having any epistemic access to the satisfaction of condition 3. In fact, S will be mostly unaware of the particular belief-forming process that led her to believe that p, let alone of its being reliable or not. Reliability accounts of knowledge are thus to be counted as externalist.

I will postpone to later sections a discussion of the internalist and externalist theories of *justification*. What I now want to emphasize is the deep entrenchment of the internalist approach to *knowledge* in the epistemological tradition that has dominated Western philosophy since the times of Descartes. Any epistemology which makes of certainty a pre-requisite for knowledge will be implicitly committed to an internalist analysis of knowing. Both Descartes and Husserl take epistemic accessibility to involve certain, or self-evident, knowledge of the object. So they are committed to the thesis that 'S knows that p' implies something like 'S can know for certain that the conditions which must be satisfied for her to know that p are actually met'. In other words, they are committed to the thesis that whenever S knows (in the pregnant sense of the word) that p, S can also know (in the pregnant sense of the word) that it is the case that she knows that p. While Descartes simply says that knowledge that p is incompatible with the

slightest doubt about p's truth¹⁹, Husserl refers to a particular objectifying act (that is, self-evidence) in which the fulfilment of the necessary conditions for knowledge is experienced by the cognitive subject.

In our century, Karl Popper has been one of the fiercest critics of the internalist view that knowledge requires certainty. His epistemology rests on a sharp distinction between truth and certainty:

We [...] must clearly *distinguish* between the *truth* of an expectation or a hypothesis and its *certainty* [...] There is much truth in much of our knowledge, but little certainty.

(Popper 1990, 33)

To be sure, other philosophers emphasized the fallibility of human knowledge before Popper. The distinctive feature of Popper's philosophy is rather its insistence on the practical incompatibility of the epistemological goals of certainty and objectivity. His idea is not just that, our cognitive powers being what they are, we must settle for something less than certain knowledge and make do with conjectural knowledge. The reason why Popper thinks we had better give up the cognitive aim of certainty is not merely that it is beyond our reach. While Descartes recommended that one set out in search of certainty if she wanted to find what was objectively true about the world, Popper insists that seeking certainty will never lead us to objectivity. His main contention is that looking for certainty will lead us to an over-cautious methodological attitude which will prevent us from learning anything of interest about reality. On the contrary, if what we are looking for is not assurance, but what is objectively true about the world, the most effective strategy will be to give up the quest for certainty and espouse a much bolder and critical attitude towards our own theories.

Popper's approach implies such a radical divorce of the ideals of certainty and objectivity that few philosophers have been willing to accept the unpalatable consequences of his epistemology. Popper himself has not always lived up to his own epistemological

¹⁹ 'The fact that an atheist can be "clearly aware that the three angles of a triangle are equal to two right angles" is something I do not dispute. But I maintain that this awareness of his is not true knowledge, since no act of awareness that can be rendered doubtful seems fit to be called knowledge' (Descartes 1984-85, II, 101).

standards; see, for example, his notorious “whiff” of inductivism’ footnote in the Schilpp volume about his philosophy (Schilpp 1974, 1192-1193).

Contemporary commitment to an internalist analysis of knowledge can be seen as a way to preserve the intent of traditional epistemology (which took the quest for certainty as an essential feature of any cognitive effort), while at the same time relaxing its criteria of epistemic accessibility. Nowadays it is commonplace to see certainty as an unattainable goal, and hence as an undesirably strong requirement to be set upon knowledge. Nevertheless, many people find it hard to conceive that S might possibly know that p without having at least some degree of epistemic access to her own epistemic situation. So many philosophers have given up the quest for certainty and settled for some kind of internalist epistemology according to which, if we are to have any knowledge, we need to have at least some *probable* grasp of our own epistemic situation. For how could one know that p if she had no reason to believe that she did?

To sum up, the terms ‘internalism’ and ‘externalism’ can be used to formulate three different dichotomies: (1) a metaphysical dichotomy; (2) a dichotomy in the philosophy of mind; and (3) an epistemological dichotomy. It is this last dichotomy (with its different versions in the theory of knowledge and in the theory of justification) that will provide the focus of our discussion in the remaining part of this chapter.

2.3 The Vertical Epistemic Regress Problem

Internalist analyses of knowledge face what I shall call the ‘vertical’ epistemic regress problem, in order to distinguish it from the more commonly discussed ‘horizontal’ epistemic regress problem.

The horizontal problem stems from the requirement that our beliefs be epistemically justified. Anyone will agree that beliefs can be justified inferentially, i.e., by deriving them from other beliefs; but inferential justification cannot be carried on *ad infinitum*. So the horizontal epistemic regress problem is usually developed as an argument in support of some kind of foundationalist epistemology. In other words, it is maintained that there must be some beliefs which are justified non-inferentially if there are to be inferentially justified beliefs at all.

The vertical epistemic regress problem stems from the internalist requirement that the conditions that must be satisfied for S to know that p be epistemically accessible to S (see Dancy 1985, 129 ff. and Alston 1989, 210 f.; Bonjour 1980, 54 f. describes the same problem, but confines his discussion to basic beliefs). The JTB account of knowledge will provide a useful illustration:

S knows that p	iff	1) p is true
		2) S believes that p
		3) S is justified in believing that p

Suppose that we construe 'S has epistemic access to p' as 'S is justified in believing that p'. Under this assumption, satisfaction of condition 3 entails epistemic access to the satisfaction of condition 1. Furthermore, there seems to be no doubt that S can have epistemic access to the satisfaction of condition 2. The vertical epistemic regress problem arises with condition 3: has S epistemic access to the satisfaction of this last condition? Under our construal of epistemic access, this question can be rendered as: Is S justified in believing that she is justified in believing that P? Our internalist construal of the JTB account of knowledge turns out to require the satisfaction of a new condition if S is to know that p:

4)	S is justified in believing that she is justified in believing that p
----	---

But then, S will need to have epistemic access to the satisfaction of condition 4 as well, and this will lead to an infinite regress.

Things get even worse if we replace 'S is justified in believing that p' with 'S knows that p' as our interpretation of 'S has epistemic access to p'.

The most straightforward way to stop the vertical regress is to give up the internalist analyses of knowledge altogether. In the case of the JTB account of knowledge, that means that one should stick to the three original conditions stated in the *analysans*: 'S knows that p' iff conditions 1 to 3 are met, whether S has epistemic access

to the fact that they are or not (but in fact externalist accounts of knowledge are more frequently alternatives to, rather than improvements on, the traditional JTB account of knowledge).

A different way of stopping the vertical epistemic regress might be to adopt a *weaker* form of internalism, as Jonathan Dancy has suggested. Such a weaker form of internalism does not equate epistemic access to knowledge or justified belief, but merely to belief. So what is required from S is not that she know or justifiedly believe that the relevant conditions for her to know that p are satisfied, but merely that she *believe* that they are (see Dancy 1985, 133 ff.).

For all of Dancy's assurances, I cannot see why this should count as an internalist conception of knowledge. The only way in which it is indeed internalist is that it rules out that sense of knowledge in which we sometimes ascribe knowledge to all kind of organisms which are well-adapted to their environment. Or rather, it rules out any kind of knowledge the cognitive subject is not aware of, such as behavioural dispositions and unconscious expectations. I think this to be undesirable, but I don't want to rest my argument on this point. What is relevant is rather that whatever plausibility Dancy succeeds in providing his theory with actually comes from a trivial equivocation. Dancy tries to convince us that the condition 'S believes that she is justified in believing that p' adequately captures the idea that 'all that we can ask of a man is that he retain beliefs which, so far as he can tell, meet the conditions for justification' (Dancy 1985, 133). But what 'S believes that she is justified in believing that p' actually means is just that S *happens* to believe that she is justified in believing that p, not that p, 'so far as [s]he can tell, meets the conditions for justification'! Dancy is smuggling in an assumption about the rationality of S's belief-forming processes which he leaves totally unargued for. But wishful thinking will not make of an externalist analysis of knowledge an internalist one!

2.4 Nozick's Externalist 'Refutation' of the Sceptic

As far as logic and physics are concerned, we might well be brains in a vat, properly stimulated by an evil (or benevolent) scientist to provide us with exactly those experiences we presently believe to be caused in us by (what we take to be) the actual

world. If this were our situation, we would not know what we usually think we know. But then, how is it possible that we know anything, if we cannot rule out the possibility that we are just brains in a vat being systematically deceived about our own situation²⁰?

This sceptical challenge has been recently addressed by Robert Nozick on the basis of his counterfactual account of ‘S knows that p’ (see Nozick 1981, ch. 3, esp. 172-178; 197-211). Nozick’s argument has been the focus of extensive philosophical debate in the last decade. An investigation of the merits and limits of Nozick’s argument will help to clarify what is at stake in the controversy between internalist and externalist epistemologies.

Nozick’s analysis of knowledge goes as follows:

S knows that p	iff	1)	p is true
		2)	S believes that p
		3)	$\neg p \Box \rightarrow \neg(S \text{ believes that } p)$
		4)	$p \Box \rightarrow S \text{ believes that } p$

As in section 1.2, the box-arrow stands for the subjunctive conditional ‘If it were the case that..., then it would be the case that...’. Nozick would like to claim that his definition doesn’t commit him to any particular semantics for subjunctive conditionals. But he has to refer to Robert Stalnaker’s and David Lewis’ ‘possible-worlds’ theories as to the best candidates for such a role (see Nozick 1981, 680, note 8), because there are relevant circumstances in which no informal reading of condition 4 seems to be possible.

If conditions 1 and 2 are met, condition 3 can be easily read without any ‘possible-world’ aid as ‘If it were the case that not-p, S wouldn’t believe that p’. It is more difficult to provide an intuitively satisfactory reading of condition 4. For what does ‘If it were the case that p, S would believe that p’ mean if it is actually the case that p and S believes that p? According to a suitable revision of Stalnaker’s and Lewis’ possible-worlds semantics, if one wants to decide whether the subjunctive conditional $p \Box \rightarrow q$ is true, she will have to examine those possible p-worlds that are closest (most similar) to

²⁰ Having repeatedly referred in chapter 1 to Putnam’s book, *Reason, Truth and History*, in which a very well-known ‘Brain in a Vat’ argument is devised, I should probably make clear that the argument I discuss here is definitely *not* Putnam’s argument (which is, to some extent, an argument *against* scepticism).

the actual world, and see if q holds true in all of them. If it does, the subjunctive conditional holds true in the actual world. Thus condition 4 turns out to mean something like: 'In all those worlds in which p holds true that are closest to the actual world, it is also true that S believes that p '. If, under the assumption that conditions 1 and 2 are met, one wanted to render condition 4 in a slightly more familiar language, she would probably have to go for something like this: '(p is true and S believes that p , and) S would believe that p in all those (slightly) different circumstances in which p were to hold true'.

The idea that Nozick's counterfactual analysis of knowing is meant to capture is that S knows that p only if her belief that p is 'sensitive' both to the truth and to the falsity of p . If S knows that p , she doesn't merely happen to believe truly that p : 'To know that p is to be someone who would believe it if it were true, and who wouldn't believe it if it were false' (Nozick 1981, 178). Nozick sums up this idea of a subjunctive connection between S 's belief about p and p 's truth-value by saying that S 's belief 'tracks' the truth that p :

To know is to have a belief that tracks the truth. Knowledge is a particular way of being connected to the world, having a specific real factual connection to the world: tracking it.

(Nozick 1981, 178)

Here I am not so much concerned with the adequacy of Nozick's account to our intuitions about what should be counted as knowledge (some useful remarks on this matter can be found in Goldman 1988, 57-63), as with the anti-sceptical strategy which he develops on the basis of that account (but we shall eventually see that these two questions are closely related).

The sceptical argument presented at the beginning of this section relies on the principle P that knowledge is closed under known logical implications (Nozick 1981, 204 ff.). This principle says, roughly, that if S knows that p and if S knows that ' p entails q ', then S also knows that q . Principle P enables the sceptic to argue, by *modus tollens*, that if S knows that ' p entails q ' and she doesn't know that q , then she cannot possibly know that p .

Let p be 'S is sitting in her room reading a book' and q be 'S is not a brain in a vat'. We can suppose that S knows that ' p entails q ' (i.e., that she cannot both be sitting in her room reading a book *and* be a brain in a vat). Yet S doesn't seem to know that q (for it is indeed logically and physically possible for S to be a brain in a vat). So how can S possibly know that p (i.e., that she is sitting in a room reading a book) if she doesn't know that q (i.e., that she is not a brain in a vat)? The sceptical predicament arises from the possibility of imagining a whole set of 'sceptical worlds' supposedly different from the one in which we believe we are, but which nevertheless would cause us to have the very same experiences we actually happen to have. Such 'sceptical worlds' would be *evidentially indistinguishable* from, and hence *doxically identical* to, the one which is supposedly ours (different worlds are said by Nozick to be 'doxically identical' for S iff S would have exactly the same beliefs in any of them²¹). But, so goes the sceptical argument, if we cannot tell that we are not living in one of those sceptical worlds, we cannot know anything at all about our own world. For none of our actual beliefs (about the world) would be true if we were in fact living in a sceptical world.

However, Nozick's analysis of 'S knows that p ' has the nice consequence that in general S can be said to know that p and that ' p entails q ' without being required to know that q . In other words, Nozick's analysis entails that knowledge is *not* closed under known logical implications (Nozick 1981, 204-211). Let us see how the counterfactual analysis of knowing handles our example involving the propositions p , 'S is sitting in her room reading a book', and q , 'S is not a brain in a vat'. It is immediately clear that S cannot be said to know that q , because if she were a brain in a vat ($\neg q$) she would nevertheless believe that she weren't (that is, she would believe that q), since that sceptical world would be doxically identical to what she takes to be the actual world. So Nozick will grant the sceptic that S doesn't know that she is not a brain in a vat. Yet, if S tracks the truth that p , she can be said to *know* that she is sitting in her room reading a book:

²¹ The hypothesis under discussion differs from Descartes' evil demon hypothesis (see Descartes 1984-85, II, 15) in so far as a 'sceptical world' must be understood as a world in which S has (is?) in fact a *functioning* brain, and in which S's beliefs arise by means of the ordinary belief-forming processes which occur within brains, the only difference being that the inputs transmitted by S's afferent nerve endings are not produced by the sort of facts which are usually thought to produce them.

For (3') [i.e., 'if q were false, S wouldn't believe that q'] talks of what S would believe if q were false, and this may be a very different situation from the one that would hold if p were false, even though p entails q. [...] There is no reason to assume the (closest) not-p world and the (closest) not-q world are doxically identical for you, and no reason to assume, even though p entails q, that your beliefs in one of these worlds would be a (proper) subset of your beliefs in the other.

(Nozick 1981, 206 f.)

In other terms, when we assess the truth-value of 'if p were false, S wouldn't believe that p', we need to consider those possible not-p worlds that are closest (most similar) to the actual world, and see if 'S doesn't believe that p' holds true in all of them. We don't have to consider those possible not-p worlds which, like all not-q worlds, are most distant from (most dissimilar to) the actual world. Thus, since none of the relevant not-p worlds will be a world in which S falsely believes that p, condition 3 will be satisfied, and S can be said to track the truth that p, even if she cannot be said to track the truth that q.

Can we conclude with Nozick that his account of knowing, having as a consequence that knowledge is not closed under known logical implications, successfully undermines the sceptical challenge by showing that indeed we can have some knowledge of the (external) world?

An apparently destructive criticism of Nozick's claim is contained in a brief paper by Edward Craig (see Craig 1989). The argument is very straightforward. According to Nozick's analysis, if any sceptical world is a close possible world, we don't really know what we think we know. But since the actual world *is* a close world, if we want to defeat the sceptic we need to be in a position to assert that the actual world is not a sceptical world.

But if we *are* in a position to assert that the actual world is not a sceptical world then the sceptic must somehow already have been defeated without recourse to the 'tracking' analysis; if we are *not* in a position to assert it, recourse to the analysis doesn't help.

(Craig 1989, 161)

According to Craig, Nozick may be right in his contention that principle P is false, but that doesn't show that we are warranted to claim any knowledge whatsoever about the world.

Craig's argument is valid, but what it shows is not really incompatible with Nozick's claims. Craig's argument shows that the counterfactual analysis of knowing is unable to prove that we have some actual knowledge of the world. But that has nothing to do with what Nozick's analysis is intended to prove in the first place. Nozick's analysis is intended to prove that S *may* (not: does) know that p even without knowing that she is not a brain in a vat. In other terms, the counterfactual analysis of knowing implies that *knowing that one's world is not a sceptical world is not a necessary condition for having any knowledge at all*. Craig's reply does not undermine this result.

Why should this result have any bearing on the sceptical challenge though? After all, Nozick hasn't shown us that we do know that p. Are we epistemologically any better off after Nozick has told us that we *might* be in a position to know that p, although we cannot know whether we actually are?

I think Nozick's result is indeed relevant to the sceptical challenge. If it doesn't seem so to Craig or to any other reader, it is merely because of his (her) own implicit commitment to an internalist view of knowledge. If the conditions that are to be met if S is to know that p need to be epistemically accessible to S, then no proof of the possibility of knowledge falling short of establishing the *actuality* of some instance of knowledge will ever meet the sceptical challenge. That is to say, commitment to internalism prevents any putative instance of knowledge which is not certifiable as such from within (e.g., the actual tracking relation which Nozick takes to be the referent of the phrase 'S knows that she is sitting in her room reading a book') from being a possible candidate to meet the sceptical challenge. Craig doesn't see (or doesn't want to see) what Nozick's point really is. Nozick doesn't want to meet the sceptical challenge by proving the existence of instances of knowledge which satisfy the internalist constraints. He wants to deny the very adequacy of those internalist constraints by showing that there is a sense in which we *may* have knowledge (i.e., have a 'specific real factual connection to the world') even if those constraints are not satisfied.

That Nozick's analysis of knowledge is externalist should have by now become apparent, for it does not set any requirement of justification upon S's belief that p. Furthermore, if S doesn't know that she is not a brain in a vat, surely she cannot know (in the sense required by the sceptic's challenge) that she is tracking the truth that p (if she is). Whether in any particular case S is tracking the truth that p will be a completely

objective matter, altogether independent from S's awareness of her own epistemic situation. But according to Nozick's analysis, S *needn't* know (or justifiedly believe) that she knows that p in order to know that p²²! So Nozick's analysis of knowing falls squarely in the externalist camp. As a consequence, it won't help to reject his argument merely by showing that it doesn't prove the existence of instances of knowledge which satisfy the internalist requirements on knowledge. What will have to be discussed is rather whether an externalist view of knowledge can provide a satisfactory alternative to the internalist orthodoxy, or whether externalism must be seen, on the contrary, 'as simply abandoning the traditional idea of epistemic justification or rationality and along with it anything resembling the traditional conception of knowledge' (Bonjour 1980, 70).

2.5 Externalist Knowledge and Internalist Justification

We have so far been discussing (epistemological) internalism and externalism *qua* doctrines about the nature of knowledge. We have not explored them *qua* doctrines about the nature of epistemic justification. Some writers tend to conflate these two issues. I take this to be a legacy of their past (or present) commitment to the traditional account of knowledge as justified true belief (see, e.g., Chisholm 1989, ch. 8, esp. 75 f., and Bonjour 1980, esp. 53 f.). If what turns a true belief into knowledge is its being an *epistemically justified* belief, then an externalist account of justification will necessarily yield an externalist account of knowledge. Of course the reverse entailment does not obtain. Yet, the influence of the JTB account of knowledge has been so strong that many internalist epistemologists seem unable to understand their colleagues' attempt to develop an externalist view of knowledge other than as an attempt to provide a different, externalist version of the notion of epistemic justification. When confronted with an externalist definition of *knowledge*, they will interpret it straightaway as the proposal of

²² 'To know that p is to actually be related to the world in a certain way, namely, to track it. But the nature of the tracking relation is such that you can track the fact that p without also tracking the fact that you are tracking p. [...] If knowledge is a real relationship in the world, such as tracking, then it will be a fact that you stand in that relationship to p; so room will be left for failing to stand in that very (tracking) relationship to the fact that you stand in it to p. If the knowledge relationship is a stringent one, not easily satisfied, there will be many cases of knowing without knowing that one knows' (Nozick 1981, 246).

an externalist account of *epistemic justification*. They will then proceed to dismiss it as totally irrelevant to 'the analysis of any ordinary concept of knowledge or of epistemic justification' on the ground that it cannot provide any answer to such questions as 'What can I know?', 'How can I be sure that my beliefs are justified?' and 'How can I improve my present stock of beliefs?' (Chisholm 1989, 76).

This I maintain is a mistake. An externalist view of knowledge will indeed look irrelevant to any traditional epistemological issue if it is taken to imply commitment to an externalist view of epistemic justification. But such an implication cannot be taken for granted.

I take it that for many people the main obstacle to accepting an externalist view of knowledge is the fear of ending up with the impossibility of ascribing any kind of rationality to human beliefs. When phrased in terms of epistemic justification, this objection can be summarized by the following question: If one cannot tell whether a given belief is epistemically justified, why should its acceptance be counted as rational merely because that belief happens to track the truth? For it seems that acceptance of a belief can be counted as rational only if the cognitive subject has some epistemic access to whatever makes her belief justified. Essentially the same objection can be phrased in terms of knowledge, without recourse to the notion of justification: If one cannot tell whether a given belief provides any true knowledge of the world, why should its acceptance be counted as rational merely because that belief happens to track the truth?

These objections arise from a legitimate concern. If one of the aims of epistemology is to advise cognitive subjects about the most effective strategies for knowledge acquisition ('How can I improve my present stock of beliefs?'), we shall need what has been called a 'regulative' notion of justification (Goldman 1980, 28 f.; see also Nagel 1986, 69). I take it for granted that any prescriptive epistemology will advise us to believe (accept) exactly those statements which come out as (comparatively) best justified on the basis of its own criteria of justification. But if one is to follow this suggestion, she will need to have epistemic access to the *justifiedness* of her own beliefs. A suggestion like 'Retain justified beliefs and reject (or suspend judgement on) unjustified beliefs' will be completely idle if one is not in a position to tell which of one's own beliefs should be counted as epistemically justified and which should not. It seems clear

to me that the notion of justification can be assigned a 'regulative' role in our epistemology only if justifiedness is taken to be epistemically accessible.

That an externalist conception of knowledge will not provide, by itself, any relevant answer to the traditional questions of prescriptive epistemology is obviously true. But that does not mean that an externalist conception of knowledge will prevent the development of any relevant answer to those questions.

I argued in the first chapter that Realism does not entail the claim of the epistemic accessibility of the God's eye point of view. My working hypothesis was that Realism could best be defended by severing its links with any form of epistemological theory involving commitment to that claim. In order to flesh out that hypothesis, I now wish to examine how an epistemology consisting of an accurate blend of externalism in the theory of knowledge and internalism in the theory of justification (plus an externalist meta-theory of justification) could allow the Realist to renounce the God's eye view without being compelled to endorse a sceptical attitude towards knowledge and human rationality.

2.6 Is Epistemic Justification Essential to Knowledge?

In the limits of this work, I cannot dwell on the details of the various versions of knowledge-externalism which have been proposed in recent years. Fred Dretske, for example, defines '*K* knows that *s* is *F*' as '*K*'s belief that *s* is *F* is caused (or causally sustained) by the information [in the technical sense of communication theory] that *s* is *F*' (Dretske 1981, 86), while David Armstrong's 'reliabilist' account of knowing requires that there be 'a *law-like connection* between the state of affairs *Bap* and the state of affairs which makes "*p*" true, such that, given *Bap*, it must be the case that *p*' (Armstrong 1973, 166). Other externalist analyses of 'knowing' can be found in the writings of Alvin Goldman, Alan Goldman and William Alston. Each of these accounts has its own pros and cons, and I cannot embark on a discussion of how faithfully each of them analyses our ordinary concept of 'knowledge'. What I am most interested in is the fact that all these analyses understand knowledge as a real relationship in the world, which can obtain (or fail to obtain) independently of the subject's having epistemic access to its own epistemic situation. This relationship can be variously characterized as a

counterfactual (Nozick), nomic (Armstrong), or causal-plus-information-theoretic (Dretske) relationship. However, it is not essential to my purposes to decide which of these different characterizations gets closer to the truth. In the remaining part of this work I shall focus my attention on the general features shared by all versions of knowledge-externalism, leaving aside the question of *which* specific relationship obtains (if knowledge-externalism is right) between S and the fact that p when S *knows* that p.

Can the conjunction of an internalist account of justification with an externalist account of knowledge constitute a legitimate epistemological option? And can it suffice to dispel the worry that the mere acceptance of the latter is enough to make an account of the rationality of human beliefs impossible?

As far as I am concerned, the main point to be stressed in this context is that *epistemic justification is not essential to knowledge, although it is essential to knowledge claims*²³. I do not claim any particular originality for this point (see Dretske 1981, 123-128), but I do have to say that part of the recent debate about internalism and externalism has been prejudiced by the opposite assumption that epistemic justification is indeed essential to knowledge.

If epistemic justification has to do with knowledge *claims*, rather than knowledge as such, there will be no difficulty in accepting an externalist conception of knowledge *and* an internalist conception of justification. Furthermore, acceptance of an externalist notion of knowledge will not preclude the understanding of such activities as discussing, evaluating, criticizing, and claiming knowledge as *rational* activities. The traditional questions of prescriptive epistemology will then be discussed from the perspective of an internalist theory of justification, which will be taken to be relevant not to our definition of knowledge, but to the critical assessment of our claims to knowledge. As a consequence, the rationality of our cognitive efforts will be understood not as the rationality of our beliefs, but as the rationality of our critical evaluation of their claims about reality.

²³ The notion of a 'knowledge claim' can be understood in a narrow sense and in a broad sense. In the narrow sense, a knowledge claim is the act performed by uttering the sentence, 'I know that p'. In the broad sense, a knowledge claim is the act performed by asserting that p. One cannot assert that p without committing oneself to provide some justification for the claim that p.

In the present work I shall generally understand the notion of a knowledge claim in the latter, broader sense.

In order to clarify this point, let us have a look at a short story devised by Laurence Bonjour as a counter-example to the externalist theories of knowledge:

Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable.

(Bonjour 1980, 62)

Bonjour originally devised his story having in mind Armstrong's particular version of externalism, which he took to be a typical example of what is usually called a 'reliabilist' theory of knowledge. A reliabilist theory of knowledge explicates 'S knows that p' as (roughly) 'S came to believe that p by applying a reliable cognitive method, that is a method which is likely to produce mostly true beliefs' (Armstrong's own version further explicates the reliability of cognitive methods in terms of a law-like connection between S's believing that p and the state of affairs which makes p true). However, we may safely disregard the circumstance that Bonjour's example was originally devised in terms of Armstrong's reliability theory and take it as a challenge to externalist theories of knowing in general.

Bonjour's contention is that Norman cannot intuitively be described as being *epistemically justified* in believing that the President is in New York City: 'why should the mere fact that such an external relation obtains mean that Norman's belief is epistemically justified, when the relation in question is entirely outside his ken?' (Bonjour 1980, 63). So far, so good. But Bonjour takes this to imply that Norman cannot be described as *knowing* that the President is in New York City either: 'From his standpoint, there is apparently no way in which he *could* know the President's whereabouts' (Bonjour 1980, 62). This suggestion I believe should be resisted.

I think Norman's epistemic situation can be better described on the basis of a separate treatment of knowledge and justification by saying that he does know that the President is in New York City, although he is not justified in believing that and thus cannot claim that he does.

Saying that Norman *knows* that the President is in New York City under the circumstances described in Bonjour's story may sound peculiar only if the possession of knowledge is taken to imply the possession of justification, and thus to warrant the statement of a knowledge *claim*. But if knowledge that *p* is not taken to licence the claim that one knows that *p*, the peculiarity of our own way of describing Norman's epistemic situation rapidly disappears. Should Norman wish to *claim* to have knowledge that the President is in New York City, he would clearly need to assess the reliability of the cognitive method which led him to endorse his belief. As Bonjour points out, he would need to assess the possibility of reliable clairvoyance in general and to decide whether he himself possesses such a cognitive faculty. If he could succeed in carrying out this task and if the outcome of his inquiry were to be favourable, he would then be able to provide a justification for his belief and could be inclined to claim that he not only believes that the President is in New York City, but that he knows that.

This way of looking at Norman's epistemic situation turns out to be much more natural than that provided by the Cartesian approach of traditional epistemology. After all, a 'fundamental facet of animate life, both human and infra-human, is telling things apart, distinguishing predator from prey, for example, or a protective habitat from a threatening one. The concept of knowledge has its roots in this kind of cognitive activity' (Goldman 1976, 791).

Nowadays we believe that natural selection has programmed in our nervous systems (and in those of most living creatures) a whole set of (quasi-)beliefs and (quasi-)belief-forming processes which enable us to cope with our environment and to acquire new, individual knowledge. After Chomsky it has been widely accepted that some kind of highly structured innate disposition to language-acquisition must be ascribed to human beings if we are to account for their competence to understand a virtually infinite set of sentences. But if Chomsky's hypothesis cannot be regarded as completely uncontroversial, different evidence of a more unobjectionable sort can be provided. It is widely agreed, for example, that our central nervous system is programmed to interpret certain visual inputs according to some rigid rules which in most cases provide the 'right' perceptual output, but which in a very limited range of circumstances don't. But even in those circumstances in which one is aware, on the basis of some independent evidence, that a given interpretation of one's visual inputs is actually incorrect, one is not in a

position to modify that interpretation on the *perceptual* level, but only to disbelieve it on the *cognitive* level.

There would be no point here in multiplying the examples of such preprogrammed beliefs and belief-forming processes we constantly rely on in our cognitive practice without even being aware of their being operating. What I wish to emphasize is rather that the overall picture suggested by currently accepted theories in the fields of psychology, neuroscience and evolution theory can hardly be reconciled with anything resembling Descartes' attempt to question the validity of the whole corpus of our beliefs and then reestablish it from scratch. 'Internalism encourages the idea that [...] we must first select a *criterion* of truth - a principle for deciding which propositions are true - before we form any beliefs' (Goldman 1980, 47). But in fact biological evolution has endowed us with a whole set of beliefs and belief-forming processes which form the basis of our most sophisticated cognitive activities but which are themselves unconscious and at any rate had long been accepted and relied on before any philosopher even began to think of them as in need of epistemic justification. Of course this doesn't mean that Descartes' epistemological project is meaningless or *a priori* doomed to failure. But it does suggest that the combination of an externalist view of knowledge with an internalist view of justification may provide, after all, a more natural picture of the epistemic situation of human beings than Cartesian internalism.

An externalist view of knowledge may provide a concept of knowledge general enough to cover such instances of cognitive situations as an electric-eye door *knowing* that someone (something) is coming, or an ant *knowing* that forage is available at the end of the tracks of the foragers who have come back heavy laden, while an internalist view of justification may provide the 'regulative' notion of epistemic justification required for an account of the rationality of such activities as discussing, evaluating, criticizing, and *claiming* knowledge. The resulting picture will be able to accommodate all the phenomena epistemology is traditionally taken to be committed to account for.

2.7 The Divorce of Certainty and Truth

The price to be paid for the adoption of such a picture is of course that one must be prepared to give up the cognitive goal of certainty and make do with objective truth. We saw at the beginning of this chapter why the pursuit of certainty has been frequently taken to be the best way to acquire objectively true beliefs. However, correspondence truth provides a perfectly clear sense in which a sentence can be said to be objectively true (or false) independently of the epistemic warrant one may have for accepting (or rejecting) it. If what makes a sentence either true or false is its structure, the referential relations between its parts and reality, and how reality is independently of anybody's knowledge (see 1.4 above), then the issue of the truth-value of a sentence will be totally distinct from any epistemic issue concerning the degree of confidence any particular (or idealized) cognitive subject is willing (or justified) to grant to that sentence. I may believe that a given sentence is true, I may be justified in believing that it is true, I may even be certain that it is true, and yet, according to the correspondence theory of truth, the truth-value of that sentence will not depend on my belief or on the grounds of my belief. That sentence's being true will never be identical with, or a function of, its being *known* (or *justifiedly believed*) to be true.

The correspondence theory of truth provides a perfectly clear sense in which we can say that our cognitive efforts, though aimed at the achievement of objective truth (i.e., of objectively true statements about the world), do not aspire after cognitive certainty (or high cognitive probability). Cognitive certainty (or probability) may provide a symptom of objective truth, but we needn't set out in search of the former if what we want to obtain is the latter. Certainty and truth are different things, and we can search for the latter without worrying about the former.

Earlier in this chapter we referred to Popper's contention that truth and certainty may even turn out to be incompatible goals for our cognitive efforts: an over-cautious methodology is unlikely to lead us to new and deeper insights into the structure of reality; it is more likely to yield minor (and usually *ad hoc*) revisions of our old theories whenever they appear to require adjustment to restore their accordance with the available experimental evidence. However, there are circumstances - especially practical circumstances - in which the reliability of our beliefs is of critical importance. When it

comes to the technological applications of science, it is clear that we are strongly concerned about the reliability, not only the objective truth, of our theories. So the problem of the reliability of the technological applications of science will have to be addressed, at some stage, by any satisfactory theory of knowledge. I suspect that the concept of epistemic certainty will play no role in the explanation of that reliability, but, however this may be, the cognitive goal of objective truth remains conceptually distinct from that of epistemic certainty, probability or reliability.

We may wish to ascribe to an electric-eye door objectively true quasi-beliefs about its surroundings, but should we ascribe it any access to its own epistemic situation? Can we ascribe to the electric-eye door any awareness of its (generally reliable) quasi-beliefs? And when it opens, is it certain that someone (something) is coming, or does it merely believe that someone (something) is coming with degree of confidence 0.51? The same problem arises with the ascription of beliefs to animal agents. The instinctive behaviour that natural selection has endowed an ant with may be interpreted as the outcome of objectively true quasi-beliefs, but has the ant any access to its own epistemic situation? Talk of cognitive certainty (or probability) in the description of such elementary cognitive situations seems superfluous or definitely out of place. I can see no compelling reason why epistemic certainty (or probability) should become an essential feature of those more complex epistemic situations in which human agents are involved.

The correspondence theory of truth plays an important role in the shift from epistemological dogmatism to epistemological fallibilism. If truth is no longer seen as a function of some epistemic concept, there can be knowledge without certainty. This fact makes it possible to understand knowledge as a naturalistic relationship between an agent and its environment. The obtaining of such a relationship becomes a completely objective matter, which has nothing to do with the agent's capability of providing an uncontroversial ground for her own beliefs. The lack of any conclusive justification of the agent's beliefs is wholly compatible with her being in possession of some real knowledge. Such an externalist view may be seen as the most consequential form of an anti-psychologistic conception of knowledge.

In its classical formulation in the writings of Frege and Husserl, anti-psychologism represents a reaction to the empiricist identification of the laws of logic with the laws of human thought. Both philosophers took great pains to persuade their contemporaries that

the truths of logic must not be understood as empirical generalizations about what human beings can (or cannot) believe, and as a whole they succeeded in their effort.

Both Frege and Husserl were prompted to oppose psychologism in the philosophy of logic by their objectivism about truth: they could not accept the idea that (logical) truth could be in any way dependent upon the judging subject. Making logical truth dependent upon the judging subject would have meant relativizing it to the contingent (and possibly variable) constitution of the human species. Being reluctant to accept such a relativization, Frege and Husserl strenuously argued that logical truths, far from being empirical generalizations about mental processes, must be regarded as 'boundary stones set in an eternal foundation, which our thought can overflow, but never displace' (Frege 1967, 13).

Even if I think that we still have good reasons to subscribe to an anti-psychologistic view of logic, today we are unlikely to be so dogmatic about the status of logical truths. However, such an externalist view of knowledge as I am trying to defend seems to me to provide a natural complement for an objectivist view of truth, and represents the most thoroughgoing outcome of an anti-psychologistic attitude towards rationality. The view that knowledge does not necessarily entail justification prompts the elimination of a further psychologistic element from our epistemology. On this view, not only truth is seen as independent of the knowing subject, but knowledge itself comes to be understood without reference to the knowing subject's being justified in her own (objectively true) belief. Since the justifiedness of a belief is in ordinary circumstances subject-relative (my being justified in believing that *I* have an headache does not entail that everybody is justified in believing that *I* have an headache) as well as species-relative (my being justified in believing that *I* see a red spot does not entail that any alien being confronted with the same situation would be justified in believing that she saw a red spot), this amounts to a de-psychologization of our notion of knowledge. In other terms, in the same way as truth is thought by Frege and Husserl to be independent of what any particular subject may believe, knowledge is now thought to be independent of what any particular subject may be justified in believing. In *this* sense I describe my favourite view of knowledge as naturalistic.

2.8 Is There any Internal Link Between Justification and Truth?

Knowledge as such does not require epistemic justification. But the activity of justifying (or criticizing) any sort of knowledge claims is central to human rationality. This fact has to do with the *growth* of human knowledge and with the claim to inter-subjective validity characteristic of its most developed forms, especially of its scientific forms.

Epistemic justification may be thought to have made its first appearance in the cognitive activities of those human beings who originally began to reflect upon the reliability of their own beliefs about themselves and their environment. However, people are usually less critical of their own opinions than they are of the opinions of their neighbours²⁴. So epistemic justification is much more likely to have made its first appearance in the social context of some primitive form of linguistic communication (further reasons in support of this hypothesis will be given in section 3.3). Epistemic justification is likely to have appeared in connection with some form of social cooperation between individuals, when language began to be used to describe facts, and not merely to control the behaviour of other members of the group (of the various philosophical speculations about the origin of language, I favour Karl Bühler's hypothesis according to which language originally emerged as a signalling, or triggering, device, and only later developed its expressive and descriptive functions; see Bühler 1927, ch. 2). Some individuals will have uttered claims about the right way to perform a certain task or achieve a certain end, and they will have supported their claims with rudimentary forms of epistemic justification, particularly of an analogical or metaphorical kind. If I am right, epistemic justification may have emerged as a way of backing particular knowledge claims uttered with the intent of securing the success of some form of cooperative behaviour.

In such a situation, the *personal* knowledge of the individual is given a linguistic objectification, so that it may become the object of a claim to *inter-subjective validity*. The individual discloses what she believes and presents it as knowledge, claiming everybody's assent to the content of her belief. In this way, she commits herself to defend her knowledge claim and to engage in the activity of *justifying* it.

²⁴ If we believe that p, normally we do not ask ourselves if our belief is justified. But if someone else claims that not-p, we ask her to justify her claim, because it contradicts our own belief.

The idea that epistemic justification is to be regarded as necessary to knowledge *claims* rather than knowledge does not prevent an adequate consideration of its role in those public cognitive activities which are best represented by modern science. For it is precisely the claim to inter-subjective validity made on behalf of the results of those public activities that accounts for the relevance of the problem of the epistemic justification of those claims. In other terms, the requirement that the claims of science be epistemically justified does not arise from science's being a form of knowledge, but from science's own claim (or rather, from the scientists' claim) to the inter-subjective validity of its cognitive results. The requirement of epistemic justification is built into the inter-subjective nature of science, which is not the sum total of the beliefs privately held by the scientists, but a multifarious (and changing) corpus of theories and hypotheses claiming the assent of every rational subject. In order to emphasize this idea, Karl Popper uses the phrase 'objective knowledge' as a technical term referring to the corpus of the cognitive claims of science as opposed to the 'subjective knowledge' represented by the personal beliefs of the individual scientists (see Popper 1972). But of course his 'objective knowledge' is rarely *true* knowledge, unless we decide to restrict the label (contrary to Popper's own intention) to that limited part of contemporary science which consists in fact of true claims about the world. This is why I prefer to keep the word 'knowledge' for what I have called 'personal' knowledge, i.e. for a particular relationship a given individual can bear to her environment, and treat the corpus of the theories and hypotheses of science as a corpus of *claims* to rational assent.

There would be more things to say about epistemic justification and how the ideal of inter-subjectivity characteristic of modern science affects its 'internality' (can scientific justification be epistemically accessible merely to one particular scientist or does it need to be accessible to any scientist - or rational being - in general?). However, I want to address the general problem of the relationship between epistemic justification and truth first, because a discussion of this problem will help to clarify some confusions in the recent debate between externalist and internalist approaches to epistemology.

The problem can be phrased as follows: Is there any *internal* link between epistemic justification and truth? Does epistemic justification *guarantee* any specified degree of epistemic access to truth? Or should we rather consider the link between

epistemic justification and truth as merely contingent? In what circumstances can a justified belief fail to be true?

That a justified belief may fail to be true is widely accepted today (or the huge literature on the Gettier's counter-examples to the JTB account of knowledge could never have flourished). Foundationalist epistemologists who believe in the existence of incorrigible beliefs of some sort deny that every kind of epistemic justification should be regarded as fallible. From the perspective of (strong) foundationalism, there is indeed some kind of epistemic justification (usually the justification provided by the self-evident character of those beliefs which express the subject's own psychological states) which secures an infallible access to truth. (Strong) foundationalism can thus be seen as positing an *internal* link between some non-inferential form of epistemic justification and truth. But what can be said of those epistemologies which do not contemplate any form of infallible justification?

Carnap's inductive logic, interpreted (as Carnap himself originally interpreted it) as providing an *analytical* appraisal of the degree of confirmation a certain body of evidence confers upon a given hypothesis, surely assumes the existence of an *internal* link between justification and truth. But we needn't require that an internal link between epistemic justification and truth be necessarily *analytical* or *a priori*. What we need to require is only that its approximate strength be known to the knowing agent prior to any further investigation of the subject. That is to say, the knowing agent must be able to tell, on the basis of her evidence and of her background knowledge, approximately how likely (that is, how *objectively* likely) her belief is to be true. So much epistemic access to truth is needed if we are to say that an *internal* link obtains between epistemic justification and truth.

It seems apparent to me that any internalist view of *knowledge* must be committed to the existence of an internal link between epistemic justification and truth. For if no such internal link obtains, the justifiedness of a belief will provide no epistemic access to the belief's truth, and the satisfaction of the components of the analysans of 'S knows that p' will not be epistemically accessible to the cognitive subject. In other terms, if no internal link obtains between justification and truth, the cognitive subject will be deprived of any access to her own epistemic situation and will be unable to tell whether she actually *knows* what she is justified in believing.

Personally, I believe that epistemic justification bears no internal relationship to truth, and I am perfectly happy with the suggestion that we can never tell whether we know those things we justifiedly believe. I see the relationship between justification and truth as a contingent one. Taking justification to be an internal matter, I allow for the possibility that S could be equally justified in believing p in two evidentially indistinguishable (to S) worlds W and W', p being true in W and false in W'. But then, our criteria of justifiedness may be more or less conducive to truth depending on the world we happen to live in (and we may not know whether we actually live in W or in W!). This way of looking at the matter of the relationship between justification and truth fits in very well with an externalist view of knowledge and an evidence-transcendent notion of truth such as correspondence truth (although, as the reader may have begun to suspect and as we shall see in greater detail in the next chapter, it appears to hamper what Feigl called the *vindication* of the proposed methodological rules). But sometimes it appears to creep also into some theories of knowledge which are officially presented as internalist. So I believe that the upholders of internalist theories of knowledge should be challenged to be very clear about their view of the relationship between justification and truth. I suspect that in many cases they may turn out to be objectively less committed to internalism than they explicitly profess.

Laurence Bonjour is definitely consistent with his own profession of internalism:

What knowledge requires is *epistemic* justification. And the distinguishing characteristic of this particular species of justification is, I submit, its internal relationship to the cognitive goal of *truth*. A cognitive act is epistemically justified, on this conception, only if and to the extent that it is aimed at this goal - which means at a minimum that one accepts only beliefs that there is adequate reason to think are true.

(Bonjour 1980, 54)

Bonjour also maintains that epistemic justification must be capable of being shown *a priori* to be adequately truth-conducive (Bonjour 1985, 10). He quotes with approval Chisholm's suggestion that 'one's purely intellectual duty is to accept beliefs that are true, or likely to be true, and reject beliefs that are false, or likely to be false' (Bonjour 1980, 55), which leads him to endorse the thesis that the epistemic justification of a belief must

confer high probability on the belief's truth. But is Chisholm equally consistent with his own internalist creed?

Chisholm's conception of justification is surely internalist, as the ten material principles which constitute his theory of justification are such that 'the proper use of them at any time will enable us to ascertain the epistemic status of our own beliefs at that time' (Chisholm 1989, 62). But then, it turns out that according to his theory of justification (which he takes to express the 'traditional conception of "internal" epistemic justification') 'there is no *logical* connection between epistemic justification and truth. A belief may be internally justified and yet be *false*' (Chisholm 1989, 76). But is there at least a *non-logical* connection between justification and truth, such that epistemic justification may be thought to confer a suitably high degree of probability upon a belief? For all my efforts, I haven't been able to locate in Chisholm's *Theory of Knowledge* any explicit admission of the existence of such a connection. To be sure, Chisholm appears to endorse the claim that 'autopsychological' statements are certain, and thus, one may suppose, certainly *true* (see Chisholm 1989, 22-25). Yet his treatment of epistemic justification is officially quite separate from any question of truth: certainty is treated as an epistemic notion, and no mention is made of its bearing upon truth. His theory of knowledge is really a theory of justification, and no effort is made to explain why justification should represent an epistemic virtue with respect to the cognitive goal of the search for truth. Chisholm's epistemology may well give expression to the 'traditional conception of "internal" epistemic justification', but it can hardly be said to espouse anything like the 'internalist' conception of knowledge. Indeed, no epistemology allowing a substantive cognitive gap to separate justification from truth (or, in a slightly different context, corroboration from verisimilitude) can legitimately claim to solve the problem addressed by the traditional, internalist theories of knowledge. (That Chisholm's epistemology cannot be seen as solving that problem is also argued, from a different perspective, by Goldman 1980, 41-42).

That a substantive cognitive gap may separate justification from truth and yet knowledge of the Realist's world may be possible to human beings is the claim an externalist view of knowledge of the sort I am trying to defend is committed to make intelligible over against the internalist protests that such a claim cannot be supported without giving up any traditional understanding of words like 'knowledge' and

'rationality'. In the present chapter I have argued that justification is not necessary to knowledge. But that is not enough. In the next chapter I shall have to argue that the overall picture emerging from an externalist view of knowledge plus an internalist view of justification, for any substantive epistemic gap it may open between justification and truth, does not preclude an understanding of how it is possible that human beings not only have some *knowledge* of their environment, but also act *rationally* in order to improve that knowledge.

Do We Really Need to Know That We Know?

3.1 The KK and Objectivity Principles

Having argued that justification is not necessary to knowledge, I have hinted at the possibility of developing a naturalistic notion of knowledge according to which S can be said to know that p even if she has no access to her own epistemic situation (recall Norman's example in section 2.5). Endorsement of correspondence truth needn't make knowledge impossible, because knowledge is a real factual connection to the world which can obtain (or fail to obtain) independently of the subject's being justified in her beliefs. Correspondence truth may be evidence-transcendent, but that doesn't affect its epistemic accessibility, let alone the epistemic accessibility of the Realist's world. For epistemic justification is no longer seen as a requirement for knowledge, and a cognitive subject needn't gain access to the God's eye point of view in order to lie in a correct factual relationship to the world. While an epistemic notion of truth makes knowledge depend upon the availability of adequate grounds for belief, an absolutist notion of truth delivers knowledge from the domination of evidence.

A consequence of an externalist view of human knowledge is that saying that 'S knows that p' comes to be regarded as more similar to saying that 'S is six feet tall' than to saying 'S has a toothache'. Suppose that S is, as a matter of fact, six feet tall. That makes the sentence 'S is six feet tall' true. But does that warrant S's utterance of the claim, 'I am six feet tall'? Of course it doesn't. S can assert that she is six feet tall only if she has measured her own height, or had it measured by someone relative. Her sentence is made true by the fact that she *is* six feet tall, but her statement is unwarranted unless she can provide some justification for making it. The same is true if we substitute 'S knows that p' for 'S is six feet tall'. According to an externalist conception of knowledge, the sentence 'S knows that p' is made true by S's bearing a specific factual

connection to the world, but the obtaining of that connection doesn't warrant S's claim that she knows that p. Norman cannot claim that he knows that the President is in New York City, unless he can provide some justification for his statement.

This picture contrasts sharply with the internalist view that possession of knowledge entails access to one's own epistemic situation. There is a widespread belief that people cannot be wrong about sentences like 'I have a toothache'. The mere fact that I have a toothache will be taken to warrant the utterance of the claim, 'I have a toothache'. Epistemologists in the Cartesian tradition have often thought that the behaviour of sentences like 'I know that p' resembles more closely that of sentences like 'I have a toothache' than that of sentences like 'I am six feet tall'. They think, roughly, that if S knows that p, S must be in a position to *know* that she knows that p. That is to say, if S knows that p, she must *ipso facto* be warranted to claim that she does.

This is how H.A. Prichard formulates the thesis that knowing that p includes knowing that one knows that p:

We must recognize that whenever we know something we either do, or at least can, by reflecting, directly know that we are knowing it.

(Prichard 1950, 86)

Following Hintikka, Roderick Chisholm refers to this thesis as to 'the KK principle' (Chisholm 1989, 99 f.). However, the KK principle cannot be true, because on any internalist or externalist account of knowledge a person can know that p and yet lack the very concept of knowledge. Such a person will not understand the meaning of propositions like 'S knows that p', or 'S knows that S knows that p'. Assuming, reasonably enough, that a proposition cannot be known by a person unless that person understands it, someone lacking the concept of knowledge will not know that she knows that p (Danto 1967). This is why Chisholm proposes to replace the KK principle with 'the objectivity principle':

The objectivity principle tells us that, if a person knows a given proposition to be true, and if he also *believes* that he knows that propositions to be true, then he *knows that he knows* that proposition to be true.

(Chisholm 1989, 100)

It is easy to see how acceptance of an internalist analysis of knowledge may lead to a commitment to the objectivity principle: if S must have epistemic access to the satisfaction of the conditions which must obtain if she is to know that p, then S must be able to tell the difference between knowing that p and merely apparently knowing that p.

However, neither the KK principle nor the objectivity principle say that a person must be able to tell, of any proposition she happens to believe, whether that proposition is known by her or not. Both principles are conditionals: 'If S knows that p, then...'. So 'sceptical worlds' and cases of merely apparent knowledge which are introspectively indistinguishable from cases of actual knowledge do not count against them. For this reason, the KK and objectivity principles have in fact little epistemological relevance. We cannot answer the sceptic's challenge - 'How can you say that you know?' - by replying, 'I know that I know!', while pointing to the KK or to the objectivity principle!

On a counterfactual analysis of knowing like that proposed by Nozick, it is a factual question whether a person that knows that p also knows that she knows that p. However, the real problem is what one *does* when uttering a claim like 'I know that I know that p'. Even if one can, in certain circumstances, know that one knows that p, what does making the claim that one knows that one knows that p amount to?

If S makes the claim that p, she is implicitly committing herself to providing some justification for her claim that p (unless she is joking, acting, or involved in other kinds of activities in which one can utter a descriptive sentence without making a statement). If somebody challenges the truth of S's claim, S may reply, 'I *know* that p'. Saying, 'I *know* that p', S may be simply laying open the claim that she has sufficient evidence for stating that p, or she may be meaning, say, that she is tracking the truth that p. In both cases, she is issuing a promissory note about her ability to justify her claim that p. But one can hardly devise any ordinary situation in which S might think it appropriate to say, 'I *know* that I know that p' - unless, that is, S belongs to the epistemologists' circle. In this latter case, S might be a cognitive optimist trying to convince her sceptical colleague of the existence of some actual instances of knowledge: 'Dear sceptic, you correctly say that I merely *think* that I know that x, y, and z; but look: I do *know* that I know that p'. This is the only situation I can think of in which there is care to utter the sentence 'I know that I know that p'. However, the sceptic won't be satisfied by S's claim that she

knows that she knows that *p* unless *S* can *prove* that she *knows* that she knows that *p*. In other words, the sceptic's challenge arises (as we saw in section 2.4) from adopting the strongest internalist notion of knowledge, which has as its consequence that 'if *S* cannot tell the difference between knowing that *p* and merely apparently knowing that *p*, then *S* does not know that *S* knows that *p*' (Feldman 1981, 269). In this sense of knowing, the possibility of sceptical worlds, as well as the possibility of less exotic cases of apparent knowledge, show that there is no true proposition that we can ever know that we know²⁵. But if we retreat to some sense of knowing in which it might be true that there are some propositions about the external world that we can know that we know (e.g. Nozick's or other externalist senses of knowing), then I can see no situation in which the sentence, 'I know that I know that *p*', might be put to any practical use²⁶.

In sum: the only thing one may wish to *do* with the sentence, 'I know that I know that *p*', could be to provide a refutation of the sceptic's denial of the possibility of (external world) knowledge. But one cannot *do* that by truly uttering the sentence, 'I know that I know that *p*', because there is no sense in which that sentence can be truly uttered that will quite do the job.

3.2 Why Should We Value Internalist Justification?

If the argument developed in the last section is correct, it appears that we can truly say that there are some contingent truths about the world that we can *know* that we know only by construing 'knowing' in an externalist fashion. But if our notion of knowing is such that, Q: *if S cannot tell the difference between knowing that p and merely apparently*

²⁵ Perhaps one could argue that indeed there are some sets of propositions that we can know that we know, i.e., necessary and autopsychological truths. However, that won't contribute to the solution of our problem, which has to do with the epistemic accessibility of the Realist's world.

²⁶ But consider the following conversation: 'I am six feet tall'. 'How do you know?'. 'My daughter measured me. So I do know'. 'But she is extremely unreliable'. 'Not in this case. I watched her and she was very careful, so *I do know that I know*'. This is a case in which the sentence, 'I know that I know that *p*', is uttered to reject an objection to a knowledge claim. But note that the objection can be successfully rejected only because it is of a *local* character. No *global* objection of a sceptical character could be rejected by merely uttering the sentence, 'I know that I know that *p*'. One can *deny* the assertion that we live in a sceptical world and therefore cannot know that *p* by saying, 'I know that I know that *p*'; but denying is not the same as *refuting*.

knowing that p, then S does not know that S knows that p, then there is no contingent truth about the world that we can ever know that we know. For the sake of brevity, I shall write ‘know’ for any sense of ‘knowing’ that satisfies Q, while I shall write simply ‘know’ in any other case. Assuming this linguistic convention, I shall say that *we can never know* that we know some contingent truth about the world* (call that the Fallibility Principle, or FP). But I take it that FP does not entail that there is no contingent truth about the world that we can ever know.

FP may appear to have disruptive consequences for epistemology. If the task of a *prescriptive* epistemology is that of providing cognitive subjects with appropriate methodological rules for enhancing their cognitive success, how shall those rules be appraised with respect to their cognitive effectiveness if there is no (relevant) p we can ever know* that we know? In other terms, how shall we assess whether a given rule is, say, truth-conducive, if we can never know* of any singular instance of putative knowledge that it is an instance of actual knowledge? This sort of predicament arises from the conjunction of an internalist theory of justification with an absolutist notion of truth:

The point here is a simple one [...] aren't some of the propositions you believe epistemically rational for you to believe? And wouldn't whatever it is that makes those propositions epistemically rational for you also be present in a world where these propositions are regularly false, but where a demon hid this from you by making the world from your viewpoint indistinguishable from this world (so that what you believed, and what you would believe on reflection, and what you seemed to remember, and what you experienced were identical to this world)?
(Foley 1985, 190)

Foley's definition of reliabilism leads him to believe that this can count as a refutation of that doctrine. But from the perspective of the present work, his passage can be taken as providing an explanation of why acceptance of an internalist notion of justification will give hard time to any upholder of a non-epistemic notion of truth engaging in the activity of vindicating methodological rules. For Foley's passage makes the point that internal justification can bear only a contingent relationship to non-epistemic truth.

By adopting an internalist view of epistemic justification, we assume that we can tell whether a belief is more or less justified with respect to the concept of justification

provided by our methodology, but we are forced to conclude that we cannot tell how truth-conducive our criteria of justification are, since we cannot know* if the world we happen to live in is more similar to what we take to be the actual world or to some kind of sceptical world (and note that there may be many stages between those extremes).

Accordingly, for meta-methodological externalism a methodological rule will be *vindicated* not when it is known* to be truth-conducive, but when it is in fact truth-conducive (compare Goldman 1980, 33; the similarity of this doctrine to the externalist doctrines of justification cannot go unnoticed).

This gives rise to the kind of difficulty I had in mind when I hinted, in the previous chapter, at the challenge of combining an externalist view of knowledge with an internalist view of justification without giving up the idea that human beings not only can have some knowledge of their environment, but also behave *rationally* in order to increase that knowledge. The fact that there is no (relevant) proposition about the world that we can ever know* that we know seems to preclude any possibility of comparing competing concepts of epistemic justification with respect to their truth-conduciveness, and hence to preclude any possibility of judging any cognitive methodology as a *rational* cognitive methodology. As Larry Laudan puts it:

if we cannot ascertain when a proposed goal state has been achieved and when it has not, then we cannot possibly embark on a rationally grounded set of actions to achieve or promote that goal. In the absence of a criterion for detecting when a goal has been realized, or is coming closer to realization, the goal cannot be rationally propounded even if the goal itself is both clearly defined and otherwise highly desirable.

(Laudan 1984, 53)

We may tell whether a given belief is justified with respect to the notion of justification spelled out by a set S of methodological rules, but it seems as if we are unable to explain why S-justifiedness should be regarded as an epistemic virtue a rational belief ought to possess. One may be tempted to think that adopting an internalist concept of justification which cannot be known* to be truth-conducive is not really different from adopting an externalist concept of justification.

(By the way, is it really true that no internalist concept of epistemic justification can be combined with an internalist view of the vindication of methodological rules?

Surely a principle like, 'We should adopt those methodological rules that are *known*' by us to be truth-conducive', will not work. But what happens if we substitute justified belief for knowledge? What happens, that is, if we define the internalist character of the vindication of methodological rules by saying that 'We should adopt those methodological rules that we justifiedly believe to be truth-conducive'? I am afraid that will not do either, because 'what we justifiedly believe' is exactly what is being defined by adopting a given set of methodological rules. And if one attempts to avoid circularity by claiming that there can be a hierarchy of different (internalist) concepts of justification (e.g., justification₁ for beliefs and justification₂ for methodological rules), she will merely succeed in shifting the problem. For then she will have to explain why justifiedness₂ should be regarded as an epistemic virtue a good methodological rule ought to possess, and so on *ad infinitum*).

My claim is that adopting an internalist concept of justification which cannot be known* to be truth-conducive is in some relevant respects different from, and preferable to, adopting an externalist concept of justification. In section 3.3 I will argue that the conjunction of an internalist view of justification with an externalist view of the vindication of methodological rules may provide a better account of successful argumentative interaction than straightforward appeal to an externalist view of justification. And in the subsequent sections I will argue that acceptance of meta-methodological externalism does not bar the possibility of ascribing rationality to the cognitive efforts of human beings.

3.3 Justification and Successful Argumentative Interaction

In section 2.6 we argued that epistemic justification is not essential to knowledge, but to knowledge *claims*. Epistemic justification has its proper place in the activities of discussing, evaluating, criticizing, and claiming knowledge. If that is true, and if there exist anything like successful²⁷ argumentative interaction between human beings, then we need an internal concept of epistemic justification, even if it is conceivable that

²⁷ I take argumentative interaction between human beings to be successful when it leads to agreement, not necessarily to truth.

nobody should ever have cared to develop a prescriptive epistemology and a 'regulative' notion of justification. In other words, my point is that the existence of successful argumentative interaction between human beings cannot be accounted for without the introduction of an internalist concept of epistemic justification.

My argument is as follows. Both Mary and John truly believe (indeed, they know) that q , that p entails q , that $\neg q$ entails $\neg p$, and that p provides the best available explanation of q . Neither of them thinks that explanations must be true. Yet Mary believes that 'inference to the best explanation' preserves truth, while John doesn't. So Mary believes that p , while John withdraws his judgement. We can further suppose that, unknown to Mary (and John), p is not only the best available explanation of q , but the only possible explanation of q . So it may be the case that Mary came to believe that p by means of a reliable belief-forming process (because in this particular case the best available explanation of q happens to be in fact the only possible explanation of q , so that the truth-value of q may turn out to be a reliable indicator of the truth-value of p). Under such circumstances, we may agree to say that Mary's belief that p , being the result of a reliable belief-forming process, is *externally* justified. Will that provide Mary with any argument to convince John that p ? Surely it will not, because Mary and John don't share the same concept of *internal* justifiedness - Mary believes that inference to the best explanation provides justification, while John believes that it doesn't. But if John did believe in the reliability of inference to the best explanation, Mary could indeed convince him that p by appealing to an instantiation of that argument.

This example shows that whether a cognitive subject can engage in a successful argumentative interaction with another depends (among other things) on whether the two subjects happen to have the same concept (or overlapping concepts) of internal justifiedness, that is, on whether they happen to have the same ideas (or overlapping ideas) about what can confer justification upon what. The better the match between their concepts of internal justifiedness, the higher the chances for a successful argumentative interaction to take place. Furthermore, two cognitive subjects who lacked any concept of internal justifiedness wouldn't be able to engage in any form of argumentative interaction, because they would lack any *argumentative* (as opposed to physical and rhetorical) means of persuasion.

Thus it seems to me that the existence of successful argumentative interaction between human beings calls for the introduction of an internalist concept of epistemic justification, even if that concept is prevented by its very nature from bearing anything more than a contingent relationship to (non-epistemic) truth. The fact is that internal justification appears to play a relevant explanatory role in a theory of human argumentation.

The reason why I prefer to appeal to successful argumentative interaction, rather than to the private activity of an individual subject's evaluating her own beliefs, in order to show the explanatory relevance of internal justification is that I take the latter situation to be a limiting case of the former. If we had no evidence of the existence of successful argumentative interaction between human beings, I think that Wittgensteinian worries might prove fatal for an argument purporting to establish the explanatory relevance of internal justification on the basis of the invariance of a lonely subject's criteria of cognitive appraisal. Nor might we hope to legitimate our preference for an internalist view of justification merely by appealing to its significance for a prescriptive epistemology (whose importance and feasibility is far from being generally recognized), if the prescriptive function of epistemic justification were not already built into many established social practices.

From a slightly different perspective, we may sum up our argument by saying that there is a basic sense in which epistemic justification must be internal because its aim is not to 'epistemize' beliefs (i.e., to change them into knowledge), but to *affect* argumentatively other people's beliefs in order to bring them into a closer agreement with what we take to be reality. And one cannot affect *argumentatively* other people's beliefs without inducing them to *see* why a change is being required.

3.4 Styles of Reasoning

Mary and John did not agree, in our example, on whether inference to the best explanation should be regarded as truth-preserving. So we can say that they had different concepts of epistemic justifiedness (in the sense that some beliefs would count as justified on the basis of Mary's criteria of justification but not on the basis of John's criteria), or

that the sets of methodological rules accepted by Mary and John defined different *styles of reasoning*. Indeed, we could have called our characters Richard Boyd and John Stuart Mill, and we would have provided an historical example of the different views philosophers can take of the style of scientific reasoning. W.H. Newton-Smith refers to the absence of inference to the best explanation in Mill's list of the methods of science as a symptom of the change undergone by scientific methodology with the postulation of ever more theoretical items and properties for the explanations of those correlations between observables which constituted the main focus of Mill's methodology (see Newton-Smith 1981, 211 f.). I am not particularly keen on inference to the best explanation, which I do not see as playing more crucial a role in the methodology of contemporary science than it did in the methodology of 17th-century physics (in fact I think that inference to the best explanation is simply invalid). But it is true that scientific methodology changes through time, giving rise to what I have called different *styles of reasoning*, and that provides some *factual* support for our picture of justification and of the vindication of methodological principles.

A methodological rule which came to be adopted only relatively recently in medical science is that clinical trials are to be performed 'double blind'. Its adoption was the consequence of the recognition that patients are subject to the placebo effect and can be affected by the therapeutic expectations of people administering drug tests. The adoption of this rule clearly makes some difference to what medical beliefs will be regarded as epistemically justified by present-day physicians.

The use of sophisticated statistical techniques to test scientific hypotheses is another relatively recent development in the methodology of science which has strongly affected the style of reasoning of the scientific community.

Non-deterministic theories failing to ascribe sharp values to some of the quantities involved in the description of a physical system would have been methodologically unacceptable for 18th and 19th century physics. Yet they are now regarded as fully acceptable by the style of reasoning of present-day physics, since they play a major role in contemporary quantum theory.

The history of science (of post-Galilean science!) is a mine of examples of methodological changes affecting the criteria of epistemic justifiedness accepted by the scientific community. Saying this, I am not denying the possibility that there might be

a few very general methodological criteria endorsed by post-Galilean science as a whole, or even a few very general principles constitutive of rationality. But it is hardly deniable that there have been in fact several *local* methodological shifts which have progressively modified our view of what can be counted as an epistemically justified scientific belief. Again, I claim that this fact is consistent with, and to some extent predictable from, our epistemological picture of epistemic justification and of the vindication of methodological principles. While the existence of several distinguishable styles of reasoning within modern science couldn't even be described without appealing to the notion of internal justification, the idea that the vindication of our methodological rules is an external matter provides a most appropriate framework for a dynamic view of scientific methodology.

With the possible exception of strictly formal rules such as, 'One should eliminate logical contradictions from one's system of beliefs', it appears that most methodological principles make in fact some substantive assumptions about the world we happen to live in. That means, as Larry Laudan puts it, that

the cogency of any methodological principle is, at least in part, hostage to the vicissitudes of our future interactions with the natural world. But that is just another way of saying that methodologies and theories of knowledge are precisely that, theories. Specifically, our methodologically rules represent our best guesses about how to put questions to nature and about how to evaluate nature's responses. Like any theory, they are in principle defeasible. And like most theories, they get modified through the course of time.

(Laudan 1989, 374)

Most of our methodological rules are such that they cannot be vindicated in all possible worlds, but only in a small subset of them, the actual world hopefully included. In section 2.8 we argued that the relationship between justification and truth can only be a contingent one: S could be equally justified (with respect to a set R of methodological rules) in believing p in two evidentially indistinguishable (to S) but structurally different worlds W and W', p being true in W and false in W'. Our styles of reasoning may thus turn out to be more or less truth-conducive depending on the world we happen to live in. By defining her own style of reasoning with reference to R-justifiedness, S is implicitly committing herself to the claim that the world she lives in is more similar to W than to

W'. I believe that all this sits very well with the fact that scientific methodology, rather than having being established *from within* once and for all, changes through time along with our beliefs about the nature of our world. We wouldn't require double blind clinical trials if we had not discovered the existence of the placebo effect! Nor would we have abandoned the requirement that our physical theories be deterministic had we not come across something like Heisenberg's uncertainty principle.

This view lays more emphasis on the 'local' character of most of our methodological rules than Alvin Goldman's own externalist view of the vindication of what he calls the 'total' optimal DDP [Doxastic Decision Principle] (see Goldman 1980, esp. 45 f.). I doubt if there is anything like *the* total optimal DDP, and furthermore I believe that what counts as a good methodological rule is relative not only to the world we live in, but to the place we occupy within that world. As Popper wrote,

All we can do is to conjecture that we live in a part of the cosmos where conditions for life, and for succeeding with our knowledge enterprise, seem to be favourable at the moment. But if we know anything then we also know that almost anywhere else in this cosmos conditions for life and for knowledge are highly unfavourable, because our cosmology tells us that the world is almost everywhere completely empty, and where it is not empty it is almost everywhere too hot.

(Popper 1972, 98)

Popper concludes that we should not hasten to provide an explanation of the apparent success of the cognitive enterprise of science, for 'this strange fact cannot [...] be explained without proving too much' (Popper 1972, 204). If the apparent success of our style of reasoning is such a contingent achievement, setting out in search of an internalist vindication of the methodological rules of science will indeed represent an attempt at explaining too much²⁸!

²⁸ As a matter of fact, Popper used the argument from the unlikelihood of knowledge to support his criticism of induction, by pointing out that present-day science 'tells us that only under very special and improbable conditions can situations arise in which regularities, or instances of regularities, can be observed' (Popper 1972, 29). As an argument against Carnap's project of a purely analytic theory of inductive reasoning that will do the job. But Popper's argument needn't be taken as meaning that there cannot be locally reliable methodological rules, although it supports the view that, if our scientific picture of the world is approximately correct, we must settle for something less than an *internalist* vindication of our style of reasoning.

3.5 Internal and External Rationality

On the basis of the epistemological picture I am trying to defend, the cognitive behaviour of an individual may be said to be internally or externally rational. I propose to say of the cognitive behaviour of an individual that it is *internally* rational when it displays a consistent agreement with a given style of reasoning (e.g., that of the social group the individual belongs to). And I propose to say of the cognitive behaviour of an individual that it is *externally* rational when it follows a style of reasoning which is 'vindicated' as truth-conducive in the individual's actual environment.

It seems to me that the cognitive behaviour of an individual may be externally rational even if the style of reasoning it conforms to wouldn't be vindicated as truth-conducive in all conceivable environments²⁹. Moreover, the manifest (diachronic and synchronic) variety of styles of reasoning characteristic of the human species does not undermine the possibility that the beliefs of some cognitive subjects adopting a very imperfect style of reasoning may nonetheless stand in the correct factual relationship to the world and hence constitute true instances of knowledge. For we saw in section 2.7 that on an externalist understanding of 'S knows that p', knowledge may be thought to be independent of what any particular subject may be justified in believing. Hence knowledge will be *a fortiori* independent of how generally effective the style of reasoning employed by any particular subject may actually be.

Can we know* that our cognitive behaviour is externally rational? Unfortunately we cannot. But of course that does not mean that our cognitive behaviour cannot *be* externally rational, if our style of reasoning is adequately truth-conducive in our present environment.

But even if it is granted that we cannot know* whether our cognitive behaviour is externally rational, no account has yet been given of the fact that people do discuss and evaluate methodological rules and styles of reasoning. How can that be if the truth-conduciveness of those rules is an external matter? Do those people merely pretend, or delude themselves, that they are evaluating the external truth-conduciveness of those rules?

²⁹ External rationality does not imply infallibility (i.e., faultless truth-conduciveness in all possible environments) any more than internal rationality does.

In the next section I will examine Alvin Goldman's answer to these questions, and then I shall try to sketch a picture of how the activities of epistemic justification and methodological appraisal can contribute to the growth of our knowledge of the world.

3.6 Epistemological Contextualism

According to Alvin Goldman, we do assess the truth-conduciveness of our methodological rules (or, as he calls them, DDPs), and while the external character of their vindication fails to *guarantee* epistemic access to the optimal DDP, 'this fact should not be confused with the claim that the optimal DDP is necessarily *inaccessible*' (Goldman 1980, 45). But how do we make our choices among competing DDPs? Goldman's answer is that if we are not to be caught in an infinite regress of DDP choices, the selection of DDPs must rest, ultimately, on antecedent doxastic habits:

there are native, or constitutional, doxastic processes that generate beliefs independently of our will and independently of the deliberate selection of a DDP. Perceptual processes automatically produce representations that, unless inhibited by other cognitions, serve as beliefs. Similarly, we are all ground-level inductivists. Expectation based on past experience is part of our animal heritage. Thus we do have means of forming doxastic attitudes before choosing doxastic *principles*.

(Goldman 1980, 46 f.)

So first comes a set of doxastic habits which generate beliefs through automatic, preprogrammed processes. Later on the cognitive subject begins to reflect upon the reliability of those processes (which implies the development of a rudimentary notion of truth and falsity) and comes to believe that some of them are more reliable than others. As a consequence, the creature sets out in search of *regulative* principles that may enable it to maximize the effectiveness of its own cognitive behaviour. But the creature will have no other starting point for its search than its actual beliefs. Goldman cites Quine, Popper and Peirce (and forgets Neurath) as earlier proponents of the idea that in epistemology there is no point of cosmic exile: 'We have to start, epistemologically speaking, from the beliefs we have at a given time' (Goldman 1980, 48). This idea is described by Goldman

as epistemological '*contextualism*'. So this is how we should assess the truth-conduciveness of methodological rules:

At first we should (and must) use our constitutional doxastic habits. Once these habits generate the choice of a DDP, that DDP should be used (together with the habits which it does not wholly displace) to form any new views about the optimal DDP.

(Goldman 1980, 48)

I think that Goldman is right and that our methodological judgements should correctly be described as 'contextual', that is, as originating, ultimately, in the context of an empirical cognitive situation involving a whole set of uncritically accepted beliefs and belief-forming processes. We have no alternative to developing our concept of epistemic justifiedness from the starting point of our present beliefs about what counts as a true instance of knowledge and what does not. Of course we have no guarantee that those beliefs are in fact true, but this does not make the development and assessment of methodological rules a mere game. *For it can be rational to develop criteria of epistemic justifiedness (and to follow them in one's cognitive behaviour) even if they cannot be proved to be truth-conducive.*

If our constitutional doxastic habits are not wholly unreliable, our initial epistemic situation may involve a number of true beliefs about our environment. Reflection about past performances of different belief-forming processes may thus lead to comparative judgements about their reliability (e.g.: visual perception is more reliable than divination for the purpose of knowing physical truths about material things of ordinary size) and to the development of rudimentary methodological rules for successful cognitive behaviour. At a more sophisticated level, the appraisal of competing methodologies will rest upon the availability of a number of intuitive judgements about some clear-cut cases of 'good' and 'bad' science, as well as upon some particular beliefs about the actual structure of the world³⁰. Scientific methodologies are built out of particular beliefs about the

³⁰ A detailed version of this view of the assessment of (scientific) methodologies is developed in Lakatos 1978, ch.2, esp. 121-136. A summary of Lakatos' view and a comparison with the 'transcendental' view of epistemological criticism put forward by Popper in his first important epistemological work, *Die beiden Grundprobleme der Erkenntnistheorie*, can be found in Zahar 1984, 150-154. But the 'theoretical' character of our methodological rules is emphasized, as we have seen, by Laudan 1989.

structure of the world (see section 3.4 above) and appraised with regard to their ability to account for our judgements about cases of 'good' and 'bad' science. If a scientific methodology turns out to conform to these judgements, it will provide some guidance for the evaluation of further, less clear-cut cases, thus providing useful advice for our cognitive practice. The possibility of developing an effective methodology rests on the correctness of our ideas of 'good' and 'bad' science, that is, ultimately, on our being in possession of some actual knowledge of the world.

Now, it is true that we don't know* whether our initial epistemic situation (the set of beliefs produced by our constitutive doxastic habits, or the set of our intuitive judgements about singular instances of 'good' and 'bad' science) involves a number of significantly true beliefs large enough to set in motion an effective mechanism of methodological appraisal. But if it does, engaging in the epistemological appraisal of competing methodological rules will increase the external rationality of our cognitive behaviour. In other words, if our initial doxastic habits provide us with a suitable set of true beliefs about our environment, from which we can derive a suitable set of true beliefs about what counts as a true instance of knowledge and what does not, then we have everything we need to develop a concept of epistemic justifiedness (i.e., a style of reasoning) which is likely to increase the *external rationality* of our cognitive behaviour. We do not know* if that supposition is the case, but we do know that at any rate no set of methodological rules can be vindicated from within³¹. So I claim that it is rational to develop criteria of epistemic justifiedness for the conduct of our cognitive behaviour *on the basis of our initial doxastic habits and beliefs* even though, failing to know* whether those habits and beliefs are respectively cognitively effective and true, we know that we will never succeed in giving a proof of the truth-conduciveness of those criteria of justifiedness we shall eventually come up with. I claim that such course of action is rational because it represents in fact the only strategy which will possibly (i.e., under the

³¹ Endorsement of an epistemic conception of what it is for a sentence to be true or false may seem to increase the accessibility of truth, but it cannot give a rationale for our methodological choices. It seems to me that the very possibility of appraising the cognitive merits of competing methodologies requires that truth be understood independently of any specific set of methodological rules. For if truth is defined epistemically in terms of a given methodology, in what sense can one methodology be said to be more truth-conducive than any other? Epistemic truth makes the choice between competing methodologies, and hence between competing notions of epistemic truth, cognitively quite arbitrary.

right initial circumstances) succeed in making our cognitive behaviour an *externally rational* cognitive behaviour.

Let R be [increasingly] externally rational cognitive behaviour; S the development of criteria of epistemic justifiedness on the basis of one's own present epistemic situation; and T favourable initial epistemic circumstances. Then the form of my argument is: You cannot get R unless you do S under circumstances T; you do not know* whether circumstances T actually obtain (so you cannot know* whether you will ever get R); but if you want to get R (since it is not impossible that you will), then it is rational for you to do S - hoping that circumstances T, unknown* to you, do obtain³².

3.7 The Context-Relativity of Epistemic Justification

In brief, I have argued that acceptance of an internalist view of epistemic justification (1) is required to provide an account of the fact of successful argumentative interaction, (2) entails, if combined with a non-epistemic notion of truth, an externalist view of the vindication of methodological rules. I have further argued that the conjunction of an internalist view of epistemic justification with an externalist view of the vindication of methodological rules (3) provides a suitable framework for a dynamic view of scientific methodology, such as is required by the fact that the style of reasoning of post-Galilean science is not static, but evolves along with the change of our scientific picture of the world, (4) does not prevent our cognitive behaviour from being *externally* rational, and (5) provides a rationale for the activity of justifying beliefs and evaluating methodological rules without assuming that we can get to know* whether our criteria of epistemic justifiedness are in fact truth-conducive.

(1) and (2) show that accepting an internalist view of justification *and* an externalist view of the vindication of methodological rules is, under certain circumstances, the only epistemological option we can make; (3) shows that the resulting picture has independent explanatory virtues; (4) and (5) show that, although preventing

³² Note that 'rational' here refers to means-end rationality, not to the internal rationality of a belief which is justified according to the methodological criteria of a given style of reasoning, nor to the external rationality of a cognitive behaviour which is in fact truth-conducive.

us from knowing* whether our cognitive strategies are in fact externally rational, that picture enables us to regard the development and assessment of criteria of epistemic justification as a rational means to the end of increasing the external rationality of our cognitive behaviour.

Having said that epistemic justification is not essential to knowledge but to knowledge claims, or in other terms that the aim of epistemic justification is not to 'epistemize' but to affect beliefs, I ought to provide an alternative account of *how* the activity of justifying knowledge claims is carried out and of *what* exactly it is meant to achieve. Unfortunately, I have no well-worked out account of these matters to offer. Here are a few very tentative proposals.

I think that any viable theory of the activity of epistemic justification should reject the assumption that people justify *beliefs*, as if internal justification could establish the truth of a belief. Strictly speaking, what people can (and do) justify are not beliefs, but knowledge *claims*. Knowledge claims can be justified because they are not sentences (or propositions) but linguistic *acts* performed in real, interpersonal contexts.

When I produce to my Italian friend Marco, who is happily ignorant of British politics, a recent copy of 'The Times' in order to justify my claim that Mr. Lamont, rather than Mrs. Thatcher, is the present Chancellor of the Exchequer, I am not trying to establish the truth of the sentence 'Mr. Lamont is the present Chancellor of the Exchequer'; I am really trying to appeal to Marco's belief that 'The Times' is a reliable source on British affairs to convince him of the truth of my claim. Alternatively, I could attempt to persuade him by producing a copy of 'The Guardian', but since my friend Marco has never heard of any reliable British newspaper but 'The Times', and furthermore suspects that 'The Guardian' could be a satirical paper, I would fail to justify my claim that Mr. Lamont is the present Chancellor of the Exchequer. 'The Guardian' might succeed in convincing Ilaria, who read it when she was in Britain last summer, but surely will not convince Marco.

According to our example, the epistemic justification of knowledge claims appears to be context-relative. The beliefs of the person to whom the justification is addressed appear to determine what can count as a justification of a knowledge claim in a given context. If the aim of epistemic justification is not to epistemize but to affect beliefs, I surmise that the kind of epistemic justification one commits oneself to provide to one's

audience when uttering the claim that *p* is precisely this kind of context-relative epistemic justification. Asserting that *p* does not commit one to establishing the truth that *p* (which is, generally speaking, impossible), but to providing arguments that may persuade one's audience to believe that *p*. According to this notion of justification, a knowledge claim (unlike a belief or a sentence) can indeed be justified, but only *relative to a certain context* (i.e., relative to the beliefs and style of reasoning of a certain audience).

I shall say that *S* is justified in making the claim that *p* in context *C* iff *S* can persuade argumentatively every subject in *C* to believe that *p*. Whether *S* can induce *argumentatively* every subject in *C* to believe that *p* crucially depends on the style of reasoning of the subjects in *C* and on the evidence they have access to. If both *S* and *Q* believe that private revelations provide epistemic justification, but *Q* has not received the private revelation that justifies *S*'s belief that *p* and is not willing to accept *S*'s word for the fact that she has, surely *S* will not be able to persuade *Q* that *p*³³!

The point I believe should be emphasized is that within the present framework the recognition of the context-relativity of epistemic justification need not support any relativistic argument against the possibility that human beings may have developed a context-independent knowledge of the world.

Widespread agreement among members of the scientific community and remarkably successful technological breakthroughs are often taken (by 'scientific realists')

³³ The methodology of scientific knowledge is probably our best shot at the de-contextualization of epistemic justification. We recognized in section 2.1 that any knowledge claim involves a claim to inter-subjective validity. The larger the context in which *S* is capable of justifying her assertion that *p*, the closer *S* can come to vindicating the claim to inter-subjective validity implicit in her assertion. Scientific methodology sets severe constraints upon the accessibility of scientific evidence, so that the context of validity of scientific justification may coincide with the context in which the standards of scientific methodology are actually accepted. Scientific evidence must be public and open to inter-subjective testing. Revelations, feelings of conviction, intuitions, and other private psychological states do not qualify as scientific evidence, because only a few privileged individuals can have access to them. Access to scientific evidence must be granted, in principle, to any member of the scientific community. In this way, the only constraint on the possibility of carrying out scientific justification is connected with the acceptance of the scientific style of reasoning, because the universal accessibility of the relevant evidence is built into the very definition of 'scientific' evidence.

Karl Popper laid much emphasis on this point (see Popper 1959, section 8), ending up endorsing a conventionalist view of the acceptance of the basic statements of science (see Popper 1959, 106). If the goal of epistemic justification is to *epistemize* beliefs, Popper's conventionalism on basic statements may seem to have disruptive consequences for his falsificationist methodology (how can a decision refute a theory?). But if the goal of epistemic justification is to *affect* beliefs, it may be perfectly rational to rest the outcome of a discussion on a previous agreement about a few shared beliefs. Furthermore, if those shared beliefs happen to be true, the outcome of the discussion may be not merely consensus, but truth.

to provide evidence for the effectiveness of the methods and the approximate truth of the theories of 'mature' sciences. I have no settled opinion about this matter, but I do think that we should not require that our theory of epistemic justification *explain* the fact (if indeed it is a fact!) that 'mature' sciences provide an (approximately) correct picture of the world and that scientific inquiry represents an instance of (externally) rational cognitive behaviour.

The character of our epistemic access to our own epistemic situation is such that we can only *hope* that the epistemic engine of justification may have geared with the actual structure of our environment and set in motion an effective mechanism of knowledge acquisition, rather than be merely running idle while producing groundless consensus.

Is Fallibilism Merely Scepticism in Disguise?

I am under no illusions that the arguments developed in chapters 2 and 3 will necessarily have succeeded in convincing the critical reader of the attractions of a Realism without a God's eye view.

After all, renouncing the epistemic accessibility of the God's eye point of view while keeping the correspondence theory of truth has a few quite unattractive consequences as far as our knowledge of the 'external' world is concerned. First, it forces one to give up any hope of cognitive certainty. Second, it forces one to admit that one can never know* that one knows that p. Third, it precludes any possibility of knowing*, of any given set of methodological rules, that it is in fact truth-conducive. If this is what a Realism without a God's eye view amounts to, no wonder that someone will feel inclined to dismiss it as nothing more than a clumsily disguised form of scepticism.

Granted, this Realism without a God's eye view does not provide a refutation of the sceptic's arguments. The sceptic is right when she claims that no belief can ever be known for sure to represent a true instance of knowledge. This is why our Realism without a God's eye view is a form of fallibilism. But commitment to fallibilism is not commitment to scepticism. For our view does prevent the sceptic from drawing two of her favourite conclusions, namely (1) that there is nothing one can ever know, and (2) that no cognitive strategy can be more rational than any other.

The first conclusion is avoided by naturalizing the concept of knowledge. The second by adopting a contextual view of the vindication of methodological rules. Following one style of reasoning may be more truth-conducive than following another one. But even if we grant to the sceptic that we shall never be able to know* that one

style of reasoning is in fact more truth-conducive than another, we can nevertheless show her that it is *rational* to adopt the style of reasoning which fits in best with (what we assume to be) our present knowledge of the world. If methodologies were purely prescriptive, the sceptic could legitimately maintain that no methodological choice could be more rational than another. But we saw that methodologies incorporate substantive assumptions about the world and our knowledge of it. This provides a basis to, or at least sets constraints on, their prescriptions. If we got those assumptions right in the first place, the ensuing methodologies will provide some effective guidance for our future cognitive activity. So it is rational to adopt that style of reasoning which fits in best with our present beliefs about the world, hoping that these beliefs incorporate some actual knowledge of the world. If we do know that people are subject to the placebo effect, performing double-blind clinical trials will be an effective cognitive strategy. The sceptic points out that we cannot know* that we know that people are subject to the placebo effect. So what? Does that remark lead to a different methodology? Should it turn out that the placebo effect is an experimental artifact, would have it been (means-end) irrational to perform double-blind clinical trials? Of course not. For performing double-blind clinical trials would still have been the only cognitive strategy which could have enabled us to adopt (if the right circumstances had obtained) an externally rational behaviour. In our situation (which means in our *objective* situation, and not from our *subjective* perspective), no alternative course of action could have been more rational than performing double-blind clinical trials.

It is wrong to maintain that correspondence truth cannot play any role in our epistemology because we can never know* whether we have achieved it or not. Correspondence truth acts as an ineliminable *regulative idea* in our cognitive efforts and especially in the development of our styles of reasoning.

I believe that any adequate epistemology should recognize the fact that the contents of our beliefs about the world always transcend the evidence that supports them, thus making it impossible to provide a conclusive justification of the truth of those beliefs. Of course one may try to get round this fact by reinterpreting the contents of our beliefs as making claims merely about the world as it appears to us, and not about the world as it is in itself. One shall then have a very hard time in explaining how ordinary and scientific knowledge of the world can *progress*. But if the fact that the contents of

our beliefs typically transcend the evidence that supports them is accepted, as I believe it should be accepted, at its face value, I think that the possibility pointed out by the sceptic - that the world may be different from how we believe it to be in ways that perhaps we cannot even imagine - must be taken seriously, and that any attempt at refuting it *a priori* is doomed to failure. As Thomas Nagel puts it,

sceptical problems arise not from a misunderstanding of the meaning of standard knowledge claims, but from their actual content and the attempt to transcend ourselves that is involved in the formation of beliefs about the world.

(Nagel 1986, 69)

After all, if the primitive cognitive processes natural selection originally endowed human beings with were aimed at increasing the chances or survival of the species, it is hardly surprising that no certification procedure was built into those processes. For the elimination of sceptical problems surely would have failed to affect positively the fitness of the species. And if later on human beings came to apply their cognitive faculties to more abstract and speculative tasks, there is no reason why natural selection should have endowed them with the capacity of assuring themselves of the good outcome of their undertaking.

If the evidence-transcendence of the contents of our beliefs represents the necessary outcome of the biological development of our cognitive faculties, I suppose that our Realism without a God's eye view could be seen as ascribing to human beings as much epistemic access to truth, knowledge and rationality as one can reasonably hope for. This is less than is promised by anti-Realism with an epistemic theory of truth³⁴, but definitely more than pure and simple scepticism.

³⁴ In the limits of this work, I cannot discuss whether anti-Realism with an epistemic theory of truth can be trusted to carry out its promises, which I seriously doubt.

References

- Alston, W.A. 1989. *Epistemic Justification. Essays in the Theory of Knowledge*, Ithaca-London: Cornell University Press.
- Armstrong, D.M. 1973. *Belief, Truth and Knowledge*, London: Cambridge University Press.
- Austin, J.L. 1962, *Sense and Sensibilia*, Oxford: Oxford University Press.
- Bonjour, L. 1980, 'Externalist Theories of Empirical Knowledge', *Midwest Studies in Philosophy* 5, 53-73.
- Bonjour, L. 1985, *The Structure of Empirical Knowledge*, Cambridge Mass.-London: Harvard University Press.
- Bonsack, F. 1989. 'Prolegomena to a Realist Epistemology', *Dialectica* 43, 67-81.
- Boyd, R. 1973. 'Realism, Underdetermination and a Causal Theory of Evidence', *Nous* 7, 1-12.
- Bühler, K. 1927. *Die Krise der Psychologie*, Jena: Gustav Fischer.
- Burge, T. 1979. 'Individualism and the Mental', in P. French, T. Ühling and H. Wettstein eds. *Midwest Studies in Philosophy 4: Studies in Metaphysics*, Minneapolis: University of Minnesota Press, 73-122.
- Chisholm, R. 1989. *Theory of Knowledge*, 3rd edition, Englewood Cliffs N.J.: Prentice-Hall.
- Craig, E. 1989. 'Nozick and the Sceptic: the Thumbnail Version', *Analysis* 49, 161-162.
- Dancy, J. 1985. *An Introduction to Contemporary Epistemology*, Oxford: Blackwell.
- Danto, A.C. 1967. 'On Knowing That We Know', in A. Stroll ed. *Epistemology: New Essays in the Theory of Knowledge*, New York: Harper & Row, 32-53.

- Descartes, R. 1984-85. *The Philosophical Writings*, trans. J. Cottingham, R. Stoothoff and D. Murdoch, 2 vols., Cambridge: Cambridge University Press.
- Devitt, M. 1991. *Realism and Truth*, 2nd edition, Oxford: Blackwell.
- Dretske, F.I. 1981. *Knowledge and the Flow of Information*, Oxford: Blackwell.
- Dummett, M. 1976. 'What is a Theory of Meaning? (II)', in G. Evans and J. McDowell eds. *Truth and Meaning*, Oxford: Clarendon Press, 67-137.
- Dummett, M. 1978. *Truth and Other Enigmas*, London: Duckworth.
- Dummett, M. 1982. 'Realism', *Synthese* 52, 55-112.
- Ellis, B. 1990. *Truth and Objectivity*, Oxford: Blackwell.
- Feldman, R. 1981. 'Fallibilism and Knowing that One Knows', *The Philosophical Review* 90, 266-282.
- Field, H. 1972. 'Tarski's Theory of Truth', *The Journal of Philosophy* 69, 347-375.
- Foley, R. 1985. 'What's Wrong with Reliabilism?', *The Monist* 68, 188-202.
- Frege, G. 1967. *The Basic Laws of Arithmetic*, Los Angeles: University of California Press.
- Goldman, Alan H. 1988, *Empirical Knowledge*, Berkeley-Los Angeles-London: University of California Press.
- Goldman, Alvin I. 1976. 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy* 73, 771-791.
- Goldman, Alvin I. 1980. 'The Internalist Conception of Justification', *Midwest Studies in Philosophy* 5, 27-51.
- Grayling, A.C. 1990. *An Introduction to Philosophical Logic*, 2nd edition, London: Duckworth.
- Grayling, A.C. 1992. 'Epistemology and Realism', *Proceedings of the Aristotelian Society* 92, 47-65.
- Haack, S. 1987. "'Realism'", *Synthese* 73, 275-299.
- Hacking, I. 1983. *Representing and Intervening*, Cambridge: Cambridge University Press.
- Horwich, P. 1990. *Truth*, Oxford: Blackwell.
- Husserl, E. 1970. *Logical Investigations*, trans. J.N. Findlay, 2 vols., London-Henley: Routledge & Kegan Paul.

- Kant, I. 1933. *Critique of Pure Reason*, trans. N. Kemp Smith, 2nd edition, London: Macmillan.
- Lakatos, I. 1978. *The Methodology of Scientific Research Programmes. Philosophical Papers, vol. 2*, eds. J. Worrall and G. Currie, Cambridge: Cambridge University Press.
- Laudan, L. 1984. *Science and Values*, Berkeley-Los Angeles-London: University of California Press.
- Laudan, L. 1989. 'If It Ain't Broke, Don't Fix It', *The British Journal for the Philosophy of Science* 40, 369-375.
- McDowell, J. 1978. 'On "The Reality of the Past"', in C. Hookway and P. Pettit eds. *Action and Interpretation: Studies in the Philosophy of the Social Sciences*, Cambridge: Cambridge University Press, 127-144.
- Nagel, T. 1986. *The View From Nowhere*, New York-Oxford: Oxford University Press.
- Newton-Smith, W.H. 1981. *The Rationality of Science*, London-New York: Routledge & Kegan Paul.
- Nozick, R. 1981. *Philosophical Explanations*, Oxford: Clarendon Press.
- Popper, K.R. 1959. *The Logic of Scientific Discovery*, London: Hutchinson.
- Popper, K.R. 1963. *Conjectures and Refutations: The Growth of Scientific Knowledge*, London: Routledge and Kegan Paul.
- Popper, K.R. 1972. *Objective Knowledge*, Oxford: Clarendon Press.
- Popper, K.R. 1990. *A World of Propensities*, Bristol: Thoemmes.
- Prichard, H.A. 1950. *Knowledge and Perception*, Oxford: Clarendon Press.
- Putnam, H. 1975. *Mind, Language and Reality. Philosophical Papers, vol. 2*, Cambridge: Cambridge University Press.
- Putnam, H. 1981. *Reason, Truth and History*, Cambridge: Cambridge University Press.
- Putnam, H. 1988. *Representation and Reality*, Cambridge Mass.-London: The MIT Press.
- Schilpp, P.A. ed. 1974. *The Philosophy of Karl Popper*, 2 vols., La Salle Ill.: Open Court.
- Tarski, A. 1944. 'The Semantic Conception of Truth', *Philosophy and Phenomenological Research* 4, 341-375.
- Van Fraassen, B. 1980. *The Scientific Image*, Oxford: Clarendon Press.

- Worrall, J. 1989. 'Structural Realism: The Best of Both Worlds?', *Dialectica* 43, 99-124.
- Zahar, E. 1984. 'The Popper-Lakatos Controversy in the Light of *Die beiden Grundprobleme der Erkenntnistheorie*', *The British Journal for the Philosophy of Science* 35, 149-171.