



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■

Jarvie, Alexander Max (1999) *Justification and acceptance*. [Masters thesis].
London School of Economics and Political Science.

<https://researchonline.lse.ac.uk/id/eprint/132555/>

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without the author's written consent.

[LSE Research Online](#) is the repository for research produced by the London School of Economics and Political Science. For more information, please refer to our [Policies](#) page or contact lseresearchonline@lse.ac.uk

Justification and Acceptance

Submitted for the degree of M.Phil. in Philosophy

London School of Economics and Political Science

By Alexander Max Jarvie

UMI Number: U126579

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 U126579

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

7730

715215

Abstract

What status does a justification confer on a claim? I try to improve upon my ignorance in this matter by examining how we regard the acceptance of claims. Justification is a sufficient condition for acceptance – that is, once a claim is justified, it is thereby accepted. Taking a claim to be true is also a sufficient condition for acceptance. Does that mean that justification leads to taking to be true? I argue not; because, although taking to be true is a sufficient condition for acceptance, it is not a necessary condition. In particular, it is not a necessary condition in those cases where the acceptance issued from a justification. There can be cases in which an individual takes an accepted claim both to be justified and to be true; but the former does not induce the latter. If justification does not lead to taking to be true, then can we at least say that it leads to taking to be probable? No; because, even if it would be descriptively accurate to view the operations of the brain as performing probability calculations, claiming that something is probable is not a sufficient condition for acceptance of that something. So it can't be that a justification issues in *regarding* something as probable, though one's acceptance of it may indeed vary in a way that conforms to the transformations of the probability calculus. If these arguments hold, then what is it that justification does for a claim? If it doesn't license regarding it as true, and it doesn't license regarding it as probable, then what else might acceptance indicate? I suggest – and I only suggest, I do not argue – that a claim is acceptable only in so far as it has not been refuted, and insofar as it appears *unavoidable*.

Acknowledgements

I would like to thank my supervisor, David Ruben, my father, Ian Jarvie, my colleague, Craig Walmsley, and Mario Bunge, for their helpful comments on the many transformations of this paper, and the illuminating discussions I had with each.

CONTENTS

Introduction.....	5
The Search for Secure Foundations.....	12
Foundationalism and its Alternatives.....	23
Our Attitude Towards Acceptance.....	36
Interlude: Van Fraassen, Belief and Acceptance.....	59
Justification and Acceptance.....	68
Conclusion.....	95

Introduction

What is the goal of a theory of knowledge? It would seem that, at minimum, it should provide answers to the following questions. What is knowledge? What is the relationship between knowledge and opinion?

There is a received answer to these questions. Knowledge is characterized, in this tradition, as justified true belief. Opinion is characterized as belief alone. If we accept the received view, justification will be a focus of concern. That something *known* should be both true and believed seems quite intuitive; what exactly is added, however, by justification? One might hope that it provides some sort of support for the truth of the belief. No matter how a justification accomplishes this task, it should at least do something else – it should somehow lead to the acceptance of the belief. If the belief was accepted without being justified – perhaps by way of some sort of intuition or revelation – then we should be reluctant to call that claim *known*, even though it may in fact be true.

So a justification should somehow lead to, or result in, belief. A substantial literature in epistemology is given over to worries about how the one leads to the other: in particular, how one can *rationally* lead to the other. One form of the worry is expressed in Hume's problem of induction. How is it that one can rationally believe that the sun will rise tomorrow, given that the evidence for such a proposition is inadequate to justify its truth? The problem arises because to believe something *just is* to take it to be true. I have no quarrel with such a characterization of belief; but I want to argue that room should be found for another epistemic attitude, which, like belief,

can serve as a basis for action and for assertion, yet does not carry the burden of a truth-attribution. Now Hume solved his problem by locating the connection between non-deductive justifications and beliefs in our psychology: it is a fact about us that, given certain kinds of evidence or, more generally speaking, reasons, our minds tend to move towards belief in such propositions. Hume's solution is not thought to be very satisfactory, for it suggests that the mind makes inferences which are not all that they could be. Our assent to propositions justified inductively puts us in an uncomfortable position: we allow ourselves to commit to something without adequate support, without guarantees. The riskiness of relying on shaky beliefs is itself disheartening, since we cannot allow ourselves to breathe easily, as it were. The demoralizing thought that we could always be wrong compels us to remain vigilant: whenever possible, we should be resubmitting those beliefs to tests (whether conceptual, looking for internal inconsistency, or empirical, looking for agreement with facts). However, most of us (if not all) overcome the depressing implications of Hume's problem and his solution, and move on. After Hume, we find an increased interest in ways of making the relationship between beliefs and justifications more precise. Bayes' application of the probability calculus to this problem has borne some of the ripest fruit; yet even the work done there has produced some unsatisfactory results. One difficulty is that while it seems sometimes correct that our attitude towards a proposition is adequately captured by thinking it plausible or probable, this does not always do justice to the phenomenology of accepting a proposition. There is a cup on the table before me; however, I am not certain of this fact. Nor does it seem right that I should say that it is probable. There is also a picture on the wall above this table; yet I could not say that I think that I accept this with a greater degree of probability than the presence of the cup. I regard each proposition as having an equal

epistemic value, so far as I am able. That is, it appears to me to be acceptable that there is a cup on the table to the same degree that it is acceptable that there is a picture on the wall. Yet given that I have looked at, handled, and otherwise observed the cup much more frequently than the picture, should I not in some sense *feel* that the actuality of the one is much more probable than the other, if my mental processes genuinely reflected something like Bayes' application of the probability calculus? I think that I should, yet I do not. In such cases, I seem to just accept the one as much as the other; and the collection of further instantial evidence towards the truth of the one or the other does not change my attitude towards them at all. Nonetheless, as I said, I do not regard them as certain. If I do not regard them as certainties, can my acceptance of these propositions be said to include an attribution of truth? That is, should I find no difficulty in passing from 'there is a cup on the table' to "'there is a cup on the table' is true'? For some reason this particular inference, which seems so natural to many, sticks in my throat. There seems to me to be a genuine difference between my acceptance of there being a cup on the table, and regarding it as true that there is a cup on the table. In any case, I find it difficult to make the move from assent to an attribution of truth. In the third chapter I explore this difficulty and try as best I can to explain why it troubles me. The upshot of the argument is that attributions of truth presuppose an unrevisability of opinion even when one acknowledges in the same breath that "one could be wrong" about such an attribution; in other words, the argument applies to fallibilist conceptions of knowledge as well as infallibilist ones. If the argument is accepted, then it follows that since we do as a matter of fact frequently revise our opinions, then we cannot ever have regarded those opinions as true. It further follows that in accepting propositions we do not thereby regard such propositions as true, nor do we regard them as either false or as neutral. They can't be

regarded neutrally, because we are after all inclined to act upon and assert these propositions we accept. Whatever the phenomenology of the status of these positively-regarded propositions, I do not want to be taken to be offering anything as ambitious as an alternative account of *belief*. As I said above, beliefs are rightly characterized as propositions which are taken to be true. What my paper intends to suggest is that we very rarely find ourselves entertaining such an attitude towards propositions we accept; in other words, we accept far more than we believe. Generally speaking, *all* we do is accept. To regard a proposition as true is something in addition.

Accepting what has been said so far, a further question arises. We said that knowledge was justified true belief. Justification leads to acceptance. Does justification also lead to belief?

In addition to arguing that we generally only accept propositions rather than believe them, I shall try to argue that while a justification leads to acceptance of a claim, it does not lead to taking that claim to be true. At this point in my argument, it will likely occur to the reader that all I have done is given one more reason for switching from an absolutist conception of claim acceptance to a multivalent one – one which, perhaps, is best captured by the probabilistic approaches developed in this century. This would make my conclusions much more palatable, since they would feed nicely into what is now a well-established tradition. However, I have not found that a probabilistic approach serves as a good characterization of claim acceptance; at least, it does not serve well as a *general* characterization. There is a place for probabilism which harmonizes with my own conclusions; I discuss this in chapter four.

I have undertaken exploration of these matters in the hope of being able to answer the question of what role a justification plays in acceptance of a claim. What I

want is some sort of account of why justification leads to acceptance. Though psychology can hardly be entirely excluded from this account, the preferred result would be a formal one.¹ I have not solved the problem in this paper, merely clarified and moved towards that goal. What I try to do in this paper is take a close look at the character of acceptance itself, and canvass some suggested criteria. By doing this, I can at least identify what it is, exactly, that justification is supposed to lead *to*. That is, when one has accepted a claim, what does it amount to? As I asked above, does it mean that one takes the claim to be true? To be probable? Or something else? I have said that I do not find either of the first two characterizations satisfactory. I venture instead that we should make, and that our practices in fact reflect, a distinction between taking a claim to be true and accepting it. A distinction between attributions of truth, and acceptance is not new. Van Fraassen, for example, tries to distinguish these from one another in **The Scientific Image**; but his target is science, not justification and acceptance in general, and his motivation for the distinction arises almost directly from skeptical considerations (in combination with his antirealism). My own reasons for wanting the distinction are less directly tied to skeptical considerations, although it will be clear (if it is not already clear) that I have a great deal of sympathy for the skeptic. In any case, the argument I make for this distinction is not one I have seen in Van Fraassen's work.

I proceed in the following way. I begin with a brief rehearsal of some well-known epistemological themes: a rational reconstruction of the passage from infallibilism to fallibilism, and a discussion of the difficulties historically associated with foundationalism. I want to discuss the former, so as to recount the normal motivation for abandoning infallibilism; I want to discuss the latter, because I want to

¹ 'Formal' in the sense of logical. Although the Fregean ideal of a total separation of the two may not

include certain theories of knowledge within the scope of the difficulties raised in this paper: in particular, coherentist theories of knowledge.

In the chapter following, I examine some possible interpretations of claim acceptance and related notions like assertion. I try to motivate the distinction by way of a thought experiment, which focusses on the problem of changing belief relative to new evidence or argument. The thought experiment is intended to show that truth-attribution prevents change of belief. In case there are worries that I have conflated a truth-taking attitude with infallibilism, I have made especial effort to show that attributions of truth are indefeasible, *even if the agent attributing truth acknowledges that the attribution is defeasible*.

Afterwards, I try to characterize some features of acceptance, in order to highlight its difference from belief; and I discuss Van Fraassen's distinction, to show that my conception differs from his own. In the next chapter, I discuss acceptance and belief in relation to justification, and show how acceptance cannot be conceived probabilistically. If these latter arguments are accepted, then it seems that we should abandon some traditional conceptions of fallibilism. Peirce, for example, is generally referred to as the first philosopher who recognized the importance of statistical or probabilistic reasoning;² and as a result is considered a fallibilist. On this view, accepted hypotheses cannot be taken to be true, but only probable to some degree. If my arguments have any weight, however, then accepted hypotheses should not even be taken as probable. Once again, I should qualify this before any misunderstanding arises: it isn't that *no* accepted hypotheses turn out, on my view, to be regarded as probable; it is only that this cannot be taken as a *general* characteristic of acceptable

turn out to be plausible, such a separation is, for me, a working hypothesis.

² At least in science. See, for example, the entry on Peirce by C. Hookway, *Oxford Companion to Philosophy*, p. 648-651.

hypotheses. And it is such a general characteristic that I am after here.

Other forms of fallibilism are not probabilistic, but merely acknowledge that attributions of truth can be revised. I think that the arguments I offer against the infallibilist count against these fallibilists as well, since, as I just mentioned, they demonstrate that attributions of truth are indefeasible, even if the agent attributing truth acknowledges that the attribution is in principle defeasible.

I shall let the text speak for itself. I have little doubt that it is wrong in detail, yet I think the general form of the argument is basically correct. There is much work to be done before the explorations undertaken here can yield a positive thesis about the nature of acceptance and justification; the most that I can claim is to have shown some ways in which their nature should *not* be understood.

The Search for Secure Foundations

The following remarks constitute a sketchy rational reconstruction of the history of epistemology. Their historicity is not germane to my argument. The views I sketch, I have gleaned from my philosophical education; they are, in most cases, what I have been told, or what I have extracted from passages selected for my perusal by various instructors. Given that my education in the history of philosophy has not been *exceptionally* poor, any errors I have absorbed ought to be common ones. The story told here is intended to remind, not instruct, the reader. What it should remind him of is the usual quasi-historical argumentation given for the general shift among philosophers away from an infallibilist conception of knowledge and towards a fallibilist conception.

Knowledge, as said above, is traditionally characterized as justified true belief. The impetus for such a characterization derives from the following problem, recorded by Plato:

Socrates. Let me explain. If someone knows the way to Larissa, or anywhere else you like, then when he goes there and takes others with him he will be a good and capable guide, you would agree?

Meno. Of course.

Socrates. But if a man judges correctly which is the road, though he has never been there and doesn't know it, will he not also guide others aright?

Meno. Yes, he will.

Socrates. And as long as he has a correct opinion on the point about which the other has knowledge, he will be just as good a guide, believing the truth but not knowing it.

Meno. Just as good.

Socrates. Therefore true opinion is as good a guide as knowledge for the purpose of acting rightly.

Meno. It seems so.

Socrates. So right opinion is something no less useful than knowledge.

Meno. Except that the man with knowledge will always be successful, and the man with right opinion only sometimes.

Socrates. What? Will he not always be successful so long as he has the right opinion?

Meno. That must be so, I suppose. In that case, I wonder why knowledge should be so much more prized than right opinion.

Meno's wonder is well-placed. What distinction obtains between knowledge and true opinion? Socrates goes on to say that knowledge is 'fastened' while true opinions are evanescent and fleeting. In what, however, does the fastening consist? The bulk of the dialogue from which the above text is extracted is taken up with a lesson in geometry for a slave boy. He is shown a procedure for doubling the area of a square. At the end of the lesson, the slave boy has gained some knowledge: that the second square is the double of the first. The procedure he followed in order to obtain that knowledge can and does serve as his *reason* for claiming to know that the second is the double of the first. He presumably could have simply opined (by way of an arbitrary guess or intuition) that the second was twice as large; in that case he would not have knowledge. The fastening, then, is accomplished by the reasons provided. An obvious benefit of having such reasons for one's claims at one's disposal is that they can be passed on: if I have knowledge that the second square is twice the first, then I can show you why, and you shall have knowledge, too.

But that is not the primary advantage for Plato. True opinions are fleeting; without the provision of reasons, one has (literally) no reason to make the claim at all. Having made it once, the slave boy might the next minute claim something else; for instance, that the second is triple the area of the first. His claims can be blown to and

from by whatever tides of fancy motivate him. But having hold of reasons provides, it seems, a normative function. They transform true opinions into truths away from which one cannot and ought not be persuaded. The slave boy, having accepted the procedure that Socrates had taught him,³ could be accused of irrationality were he then to affirm that the second square is the triple of the first; for the procedure he accepted compels him to affirm only that the second is the double of the first, if he is to affirm anything at all.

For Plato, this is the essence of the power of justification. True opinions are just as good a guide where our actions are concerned; but they lack reasonable impetus for their affirmation. Some sort of explanation of them can be given, perhaps causal, perhaps psychological, but not logical. If logical reasons are given, then whatever claim is being affirmed is *ipso facto* justified.⁴

This conception of knowledge was, for a long period of western history, undisputed in its form. That these three elements – truth, belief, and justification – were the individually necessary and jointly sufficient conditions for knowledge was not seriously attacked until the twentieth century. There were and still are, however, disputes over what justification consists in. From the Sixteenth century onwards, various doctrines were put forward describing possible sources of authority for knowledge claims. Although it is usually difficult to decide whether a philosopher falls solely into one or another of these camps, the division between rationalists and empiricists remains an important one. Justification can have one or another source of authority: reason or sense experience. The adoption of either side leads to difficulties.

If we take the *Meno* as a paradigm example, then we can see that Plato's

³ In keeping with the Platonic doctrine of recollection, it may be better to say 'helped him to recall'.

conception of what can count as a reason falls squarely on the rationalist side of the divide. Although he may not have put it this way himself, what grounds claims to knowledge are deductive arguments. The attractions of such a position are obvious. Logic, for whatever reason, has a powerful intuitive appeal: arguments conducted by the rules of deduction seem to have an iron-clad quality, and yield conclusions which are inescapable. Someone who justifies their claim to know by way of deductive argument can assert it with a confidence not given by other forms of argument, such as inductive argument. That distinction, between deductive and inductive forms of argument, is not quite so cut and dried as it may appear;⁵ but in any case, the rationalist approach suffers from another difficulty worth mentioning here.

The inescapable conclusions of deductive arguments remain inescapable only so long as their premises are undisputed. To the extent that the conclusions of such arguments are acceptable, that acceptance must extend to the premises from which they were derived. The acceptability of the conclusions rests on their being deduced from the premises; so we are invited to speculate on the source of the acceptability of those premises. Once such speculations are made, the threat of a regress looms. Given that acceptability consists in being properly deduced from acceptable premises, the regress looks vicious; if no end can be found, then it seems that no acceptance of claims is genuinely possible.

Two solutions to the regress problem are quite standard. The first is simply to

⁴ Discussion of Plato's conception of knowledge as justified true belief in the *Meno* can be found in C. Janaway's article, 'The Pre-Socratics and Plato', in **Philosophy: A Guide Through the Subject**, pp.366-368.

⁵ I allude here to two broad conceptions of how inductive reasoning works. It is possible that the impacts of particular experiences on human consciousness, by means of some underlying and inaccessible process, themselves *suggest* the generalizations or predictions that constitute the conclusions of inductive arguments. However, the acceptability of such conclusions may issue not from such occult processes but from premisses tacitly assented to, which when added to the sensory reports, constitute a deductive argument. I am undecided on this matter, and given its contentiousness, I will not pursue it further here.

declare or to stipulate that certain premisses are acceptable. This solution has found favour with some twentieth century philosophers, among them A. J. Ayer.⁶ The other solution has been to identify a way in which certain premisses can be made acceptable by being given (some sort of) undeniable backing. For some faithfully minded folk, the introduction of God was an attractive manoeuvre. Descartes used God's benevolence as his guarantee for the connexion between clarity and distinctness, and truth. Others sought earthly assistance: trying to find ways of linking sensations – which could not be doubted – and sentences describing them. Such sentences could be treated as acceptable premisses.

This last ploy is what we know as empiricism. Empiricists do not deny the importance of reason in the process of extracting knowledge from opinion. They do, however, stress that to rely on reason alone will produce nothing.⁷ Because the justificatory power of premisses in a deduction is questionable, deductive reasoning cannot escape the regress involved in using it as the ground for accepting claims. As a matter of fact, a similar problem turns out to affect the empiricist route as well.

Empiricists built on the intuitive idea that our knowledge of the world comes to us through the senses. Sensations do not seem to admit of any sort of doubt; while one can doubt that one has just seen a red ball, it is considerably more difficult to deny that one has seen a red round bulgy patch in one's visual field. This illustration somewhat misdescribes the situation, however. It seems right to say that some elements of sensible experience cannot be doubted, and some can. We can doubt that the red bulgy patch is a ball; we may be in the grip of an illusion of some sort. We can

⁶ See *The Problem of Knowledge*, p.74; see also discussion of this point in connection with Ayer and Morton White in W.W. Bartley, *The Retreat to Commitment*, pp. 95-99. Keith Lehrer also mentioned to me in conversation that he takes it for granted that some premisses will remain undefended. See also his *Theory of Knowledge*, chapters four and five.

also, however, doubt the red bulginess of the sensation, even though we cannot doubt the sensation itself. The predicates ‘red’ and ‘bulgy’ seem conceptual in the same way that the predicate ‘ball’ does; they are all characteristics that our sensations can have. But the very fact that these characteristics can be thought of in this way suggests that they are separable – at least in principle – from the sensations themselves. *Something* is given to us in sensation that cannot be doubted. Our only access to this given, however, is mediated by our descriptions of it; and since they are *our* descriptions, we cannot be sure of them. That is to say, that describing the world as red, round and bulgy – not to say three (or four) dimensional, material, and so on – may not cut nature at its joints. Our descriptions of bits of the world may be grossly erroneous. Given that our claims to knowledge are couched in terms of the way in which we describe our sensations of the world, then even the indubitability of those sensations cannot give us firm ground to support inferences to knowledge. The reason this problem resembles that of the rationalist, is that one can characterize the doubtfulness of our descriptive terms as engendering a vicious regress; one can ask *why* such terms as ‘red’, ‘round’, and ‘bulgy’ are appropriate. Once one asks for *some* reasons, one can go on to question the reasons given, too.

This latter point did not, to my knowledge, trouble the earlier empiricists such as Bacon, Locke, and Berkeley; the first detailed exploration of such concerns is found in Kant. There was yet another difficulty with sensationalism, however, which was emphasized by Kant’s predecessor, Hume; and, while the corrosive power of Hume’s point threw the empirical approach into disarray, it also opened up the possibility of severing the so far largely undisputed link between knowledge and

⁷ This can be found in Bacon, but also in Descartes, curiously enough. Syllogistic reasoning cannot in and of itself produce knowledge; but in the *Meditations*, it seems clear that Descartes’ attempt to sift out knowledge from opinion is an exercise in deduction.

deductive logic.⁸

The difficulties just examined have had to do with *particular* knowledge, such as ‘here is a red ball’. This, however, is generally speaking a rather uninteresting sort of knowledge; what we crave more often is general knowledge. The value of collecting particular observations of, for instance, green emeralds, is that one hopes to be able to discover the truth or falsity of generalizations such as ‘All emeralds are green’. Even if we can assume that the problem of a skeptical gap between sensations and their descriptions is solved, Hume showed that a skeptical gap also arises between general knowledge and the evidence that supports it. In other words, his problem had not to do with the validity of our descriptive framework, but rather with the validity of a certain class of inference.

Hume observed that sensory experiences cannot deductively justify the generalizations that they are usually taken to support. The collection of empirical reports of the sun rising each morning can be used to reasonably support neither the generalization that the sun always rises, nor the singular statement that it ^{will} ~~shall~~ rise tomorrow. There is no logical conflict to be found between asserting that the sun ^{will} ~~shall~~ fail to rise tomorrow, and the collection of undisputed reports of its having risen every day since records began. Because there is no logical conflict, Hume found himself puzzled as to why the inference from the sun’s having risen in the past to its rising in the future should seem so compelling. He hoped (I presume) to find a rational basis for the intuitive power of this and other *inductive* arguments. Hume was motivated by his fondness for empirical science. For Hume and other empiricists, our senses give us structured information about the world, upon the back of which we can make predictions and generalizations about that world. However he recognized that it was

⁸ I say *largely* undisputed, since at least both Bacon and Descartes were not convinced of such a link.

not a matter of deducing generalities from the singular statements delivered by sense experience; the relationship between empirical information and generalizations was not iron-clad.

Hume solved his problem by recourse to psychology; it was by means of the repetition of particular instances that such inferences became so compelling. Perhaps unsurprisingly, this answer did not satisfy anyone but Hume (assuming that it satisfied the man himself). The status of science, its purported laws and generalizations, seemed, after Hume's enquiries, not nearly as firm and abiding as it ought to have been. The methods of the sciences had been taken to yield knowledge, where knowledge was understood as the Platonic ideal of justified true belief. Although that view of science could no longer endure in the wake of Hume's critique, he did pave the way for a more sophisticated understanding. Instead of treating the evidence of the senses as making the claims it supported *true*, it was taken to make them more probable or more likely to be true than claims which were less well supported. This meant that, although the collection of empirical evidence for some claim did not logically exclude any actual or potential competing claims, it did count towards the truth or accuracy of that claim in some way. A popular way of metricating the relationship between claims and evidence, so understood, is the use of Bayes' theorem from the calculus of probability.

This approach was seriously challenged by Nelson Goodman's new problem of induction. Goodman argued that the collection of empirical evidence cannot be used as support for some particular claim, since it could equally well support a different claim. His notorious example involves the collecting of reports about the greenness of emeralds. On the basis of many such reports, we conclude that (it is very likely that) all emeralds are green. But suppose we introduce the predicate *grue*,

defined as

Something is grue iff it is examined before time t and is green, otherwise blue.

If all of the emeralds we observed were checked before time t , then that body of reports supports both the claim that all emeralds are green and that all emeralds are grue. Goodman says that it is obvious that no one will want to claim that all emeralds are grue; however, it is difficult to see why the preference we have for the claim that all emeralds are green is not somehow arbitrary. Goodman himself finds the solution in practice: certain predicates become entrenched and others do not; and the original choices are indeed arbitrary.

It is to Goodman's credit, I think, that he recognizes that such originally chosen predicates must be chosen arbitrarily, or at least conditioned only by human psychology: by our motives and presumably our imaginative capacities. He argues that these choices are legitimized by their origins. If one were to ask, 'Why *green* instead of *grue*?', Goodman would respond with a description of the origin of our use of *green*; this description, which could involve reference to culture, indoctrination, and to the constraints of human psychology, serves as the justification of our preference for *green* over *grue*. Goodman, in fact, generally speaking equates justification with description;⁹ if that description involves reference to human psychology, so be it.

Goodman's understanding of what a justification consists of is a radical departure from earlier understandings. We began with Plato's analysis of knowledge into justified true belief. His conception of the function of justifying led to the

⁹ Goodman, Nelson, *Fact, Fiction and Forecast*, p.64.

conclusion that a justification must entail the claim it supports; the relationship between justifier and justified is a deductive one. We saw also that Plato's approach invites a vicious regress with respect to the rational acceptance of claims to know; it also invites speculation as to the origin of the premisses that form the foundation of all subsequent deduction. The empirical philosophers of the modern period sought to stem the regress by declaring a certain class of statements acceptable without having deduced them from earlier premisses. This class included all empirical observations. It was noted, however, that such singular statements could not provide an ironclad backing for the general statements 'derived' from them; this worry was most sharply formulated by Hume. His objections stimulated significant interest in the calculus of probabilities, which seemed to offer a way of formalizing the relationship between general claims and inconclusive empirical evidence. If successful, such a formalization could conceivably rationalize the procedure of decision-making under uncertainty. This would be true, incidentally, even if it could not be said of us that we actually do perform (or our brains perform) probability calculations on all of our decisions.¹⁰ It would be enough to treat such a formalization as normative: that is, we could judge someone rational or irrational by attending to the deviation between their own decision patterns relative to evidence and the patterns that the probability calculus would dictate.

A new understanding of justification had emerged. A justified claim was no longer to be considered true, as in the Platonic tradition; it would be considered only likely to be true, or probable. In other words, a successful justification for Plato and Descartes warranted the use of the predicate 'true' in connexion with the claim justified; for those of an empirical frame of mind, a successful justification allowed

¹⁰ It has been noted by many that such calculations would quickly result in a combinatorial explosion of

only the application of the weaker predicates 'probable', 'likely', and their cognates.

Without pretending to have a clear understanding of what justification *does* in fact amount to in both our common sense and our scientific uses of the term, I shall argue that none of the above accounts will do.

This concludes the rational reconstruction. In what follows, I shall discuss three familiar approaches to knowledge and their relationship to one another.

Although this will require some reference to this rational reconstruction, the chapter is about different approaches to knowledge. The approaches are foundationalism, coherentism, and externalism; I discuss them with a view to explaining which approaches fall within the scope of my investigation.

such proportions as to render it quite implausible that an individual could actually be computing them.

Foundationalism and its Alternatives

All the approaches discussed in the preceding chapter can be characterized as foundationalist. Once seen that way, one can argue that the problems of sensationalism, regress, and induction, which look so formidable, need not afflict approaches that eschew foundationalism. Both coherentism and externalism endeavour to avoid these problems. I shall discuss foundationalism and its relationship to the coherence theory of knowledge, and show that the coherentist approach is just a species of foundationalism. This is important because the coherence theory was developed to evade the difficulties of the foundationalist approach, and because I want it to be clear that my critique of justification and acceptance is meant to apply to coherentism too.

Foundationalism traditionally comes in two forms: rationalism and empiricism. As mentioned above, any given philosopher will generally turn out to defend aspects of each in their work; but the division makes discussion somewhat clearer. We have already discussed a number of aspects of foundationalism in the foregoing chapter. Foundationalist theories treat claims as acceptable only when justified: that is, backed by reasons. But foundationalists face the problem of an infinite regress. If every claim must be backed by reasons, then those reasons must themselves be backed – by other reasons. If the infinite regress is allowed, then given that we *do* accept some claims, our acceptance will not be rational. Under such circumstances, our acceptance of claims could only be rational if we had also accepted the entire infinite chain of reasons; it seems highly counterintuitive, to say

the least, that we do in fact accept an infinite chain of argument every time we accept any particular claim.

It is far more likely, then, that the regress is stemmed in some way; different approaches have been tried. For rationalists like Plato, such a regress was supposed to be halted by the Forms. All knowledge, for Plato, is knowledge of Forms; it was the slave-boy's unconscious knowledge of the Forms that made it possible for him to recollect the procedure for doubling a square. Plato's doctrine of recollection explained why it was that the slave boy was not born with an articulate understanding of all the Forms. Through careful study, however, and by attempting to articulate why he would accept certain claims, he would eventually be able to articulate what Forms his acceptance depended upon.

A host of difficulties plagued Plato's theory, but one should be singled out, from our perspective, as particularly troublesome. If a chain of reasoning began with premises involving the Forms, how was one to know when one had in fact reached those Forms? That is, how were Forms to be distinguished from intermediate premisses? They had, apparently, a special character which Plato described in his *Seventh Letter*; 'at last in a flash understanding of each blazes up, and the mind, as it exerts all its powers to the limits of human capacity, is flooded with light'.¹¹ It can be seen that the criterion for distinguishing Forms from intermediate premises was one of intuition. The manifest unsatisfactoriness of this criterion encouraged alternative accounts.

Empiricists located the end of the regress in sensation. This approach at least satisfied the intuition that our knowledge comes from contact with the world around us. However it seems unlikely that such contact could, in and of itself, yield the sorts

¹¹ Plato, *Seventh Letter* (344b), appearing in *Collected Dialogues*.

of general claims which we make. Our contact with the world is parcelled into singular events. If reports of these particulars are to form the foundation for our inferences to generalizations, it must be noted that they cannot be as intuitively compelling as Plato's deductions are.

If we could secure the foundations of a rationalist approach, then what knowledge our reasoning produced could be claimed as true; securing such foundations, however, seems impossible. If, on the other hand, we could secure the foundations of an empiricist approach, then at most we could claim that such knowledge as our reasoning produced was probable or in some way supported by said foundations.

In either case, however, these strains of thinking about the problems of knowledge and justification are properly referred to as foundationalism. They are characterized by their appeals to an authority; whether that authority has its origin in rational intuition or in sense experience is immaterial. The objections which may be levelled against a foundationalist approach do, however, vary according to the nature of the authority proposed. In the *Meno*, Plato treated the Forms as the foundations of all knowledge; yet, the ultimate authority for his system derives from the intuitive recognition of Forms which each of us is capable. This rational intuition is itself left unsupported. Descartes endorsed a form of this rational intuition explicitly; he describes it at the opening of the Third Meditation:

I am certain that I am a thinking thing. But do I not therefore know what is required for me to be certain of anything? Surely in this first instance of knowledge, there is nothing but a certain clear and distinct perception of what I affirm. Yet this would hardly be enough to render me certain of the truth of a thing, if it could ever happen that something I perceived so clearly and distinctly were false. And thus I now seem able to posit as a general rule that everything I very clearly and distinctly perceive is true.¹²

When some claim seemed clear and distinct, Descartes decided that he was licensed to call that claim true. The authority of such intuition is, however, itself doubtful. Descartes recognized this, and in the arguments that follow the passage quoted he attempted to bolster such intuition with arguments for the existence of God. Whatever the other merits of Descartes' criteria of clarity and distinctness, the problems encountered in the attempt to give authority to those criteria made his system less firm than he had originally wished it to be.¹³

But just as the rationalists have their headaches, so do the empiricists have theirs. Let me reiterate them. Although empirical evidence always comes in the form of singular *statements*, it seems natural to take the *sensations* to which those statements correspond as themselves indubitable. That very distinction, however, causes serious problems. If someone makes the claim 'I see a red patch', one can ask what, precisely, was *given* in the sensation that prompted the claim. The sensation cannot itself be characterized as consisting in *a red patch* or *red-patchiness*; such a characterization already attributes far more to the sensation than it could possibly yield. *Red* is a concept and a predicate, and so applies to any and all instances of redness. To suppose that a single impression of a red patch could give such a rich understanding of itself to the individual sensing is to suppose a great deal. Whatever *is* given in sensation cannot, it would appear, be characterized in conceptual terms; this makes it well nigh impossible to characterize at all.¹⁴ If sensation must be divorced from conceptualizing, it opens a skeptical gap of a very worrying sort. It may be that the way in which we conceive of the world is wholly wrong, even if our sensations are not themselves open to doubt. Conclusions arrived at by means of such

¹² Descartes, *Meditation* Three, p. 24 in Hackett (marginal ref. 35).

¹³ Descartes wanted to establish 'firm and lasting' results in the sciences. Med 1, (marg ref 17), Hackett 13.

doubtful descriptions of the world will themselves be doubtful.

Both sorts of foundationalists, then, encounter difficulties from which escape is difficult if not impossible. While foundationalism can be called a theory of knowledge, it is also - or at the very least includes - a theory of justification. It explains the acceptance of claims by reference to reasons provided, and says that some of those reasons must have an authority that prevents the vicious regress discussed above. Although I have pointed out the most widespread objections to the various forms of authority proposed, I don't want to suggest that the foundational approach is now obsolete. There is still a lively debate on many of these issues; but the problems are of such a depth that they have led to the development of a different approach, known as the coherence theory of knowledge. I wish to discuss that approach now in order to show that, despite hopes to the contrary, in all the relevant respects, the coherence theory is just as much a foundationalist theory as the varieties described above, and is challenged by the same sorts of problems.

A coherence theory of knowledge is succinctly if coarsely described as saying the following of a particular knowledge claim: that it is justified if it coheres with an already accepted body of beliefs. The justification of a claim rests on the connexions between that claim and the other claims which go into making up a corpus of knowledge. It should be noted that to claim that a statement coheres with other statements in a corpus is to claim more than mere consistency. The claim that there is a red book in front of me is consistent with my current body of opinions about the world; so too is the claim that I am suffering from a hallucination of some sort. If coherence were consistency only, there would be no way of deciding which of these claims was justified. Consistency is a necessary condition for coherence, but it is not

¹⁴ Not for want of effort, however. There is a healthy tradition from Kant onwards which attempts to

sufficient.

In order to clarify what would be sufficient, some coherence theorists speak of different claims enjoying different grades of coherence with respect to some system of opinions. One or the other of the above statements, it is claimed, could be more or less coherent with respect to some body of opinions. It may be that one of my background opinions consists in accepting that, at any given time, my senses *could* be deceiving me. Both the claim that I see a red book and the claim that I am hallucinating are *consistent* with this background opinion. One of the claims may, however, be more *coherent* than the other with respect to my background opinions. It may be that the claim that I see a red book is more coherent, because it simplifies my understanding of experience to assent to it, or because of certain ideas I have about the trustworthiness of evidence, especially empirical evidence. Both *simplicity* and *trusting evidence* can be objected to as criteria of acceptance. If we allow, however, that such background opinions do not themselves stand in need of justification, then we can see that they serve as justifications for accepting the claim that I see a red book rather than the claim that I am subject to an illusion or deliberate deception designed to look like a red book. These background considerations provide *reasons* for the acceptance or rejection of such claims.

A counterargument can be made to the effect that such a scheme in fact shows how coherence is nothing more than consistency. The individual who holds that the trustworthiness of evidence is a criterion for accepting the one claim and denying the other could be understood as accepting the claim that there is a red book rather than an illusion present on the basis of the consistency of that claim with the criterion of evidence. The claim that one is merely hallucinating is just not consistent with that

characterize sensation without reference to concepts.

criterion. Thus coherence is just consistency; which means that coherentists are open to the objection that, since there are many consistent systems of opinion available, coherentists need to show how the position they hold can be justified, relative to another consistent system. This they cannot do except by rendering their system foundational; since whatever criterion is chosen to decide between consistent systems will itself serve as the fundamental justification of that system.

Such an objection may not be a fair representation of the coherentist position, however. The claim that coherence is mere consistency can be scotched by refining our understanding of the background opinions being used as criteria of acceptance: in this case, trustworthiness of evidence and/or simplicity. For one need not opine that these criteria are categorical; other considerations may intrude, and the criteria themselves may be characterized with generous escape clauses. The trustworthiness of evidence, for example, may not be so wholehearted as it at first appears. Rather than saying that evidence is always trustworthy, one could say that it is generally trustworthy and deserves serious consideration; but depending on other factors, it may not be the sole judge of the status of claims contending for acceptance.

It could be further argued that such a refinement would introduce such factors as assessment in terms of probability or plausibility. For example, it might be claimed that the likelihood of the simplest explanation being the best in this case, is low; or that the trustworthiness of the evidence is low, and so on. The introduction of such factors is not, however, problematic, since they would be just a few among a host of factors internal to the system in question. Thus they would be just part of the coherent system and not a foundation for it.

The next question to ask is, however, the one just put aside above. What justifies the acceptance of the background opinions? How did they come to have the

status of authority by way of which these new claims could be judged? The answer is that they derived their justification purely from their coherence with other background opinions. Some opponents of the coherence theory say this defense falls into a series of circular arguments. Claim A is justified by background opinions C, D, and E, which are in turn justified by B, F, G, and so on, until at some point we discover that one or more of B, F, or G are justified by one of A, C, D, or E. Now this may be a correct way of interpreting the coherentist position; but it isn't clear to me that, even if that were the case, this kind of circular argument would be in any way vicious. For it is part and parcel of the coherentist position that justification is a phenomenon which can only be understood as internal to the system under which one is operating. To ask for justifications of the opinions one holds is to ask for more opinions held, and perhaps an explanation of *how* they cohere, since it may not be immediately obvious. To demand reasons which have their own basis *outside* the coherent system of opinions of the agent, is to ask for something which doesn't make sense. Besides which, such a demand exposes a foundationalist bias which coherentists need not answer to, since they deny that knowledge has any foundations. It might be objected that coherentism, so understood, seems to imply a sort of idealism; there may be a world external to the coherent system of an individual, but that world may not be accessible for such an agent. Their opinions may or may not be true of it, but their justifications for accepting their own opinions will not reflect the truth or falsity of those opinions. This objection seems correct to me, in that it is unanswerable. If it turns out that the coherence theory is in some way the best route to take with respect to understanding the nature of knowledge, then perhaps that unfortunate consequence must be accepted along with it. I shall return to that point later in this paper.

What isn't clear to me, however, is that the coherentists can genuinely escape

the charge that their position is covertly foundationalist. I can identify two worries. Coherentists generally claim that their advantage over foundational theories issues from their rejection of the idea that certain opinions are intrinsically basic and others are intrinsically mediated, where mediated opinions are those which require reasons (specified in the form of other opinions) for their support, and where basic opinions are those that do not. However, proponents of the coherence theory generally agree that in a given context, certain opinions will play a foundational role with respect to others. This could be taken as too much of a concession to the foundationalist.

Suppose I were attempting to justify an opinion on the basis of other opinions I hold. By allowing my justification to take this form (i.e., giving reasons for my acceptance of the opinion), I am open to the infinite regress argument. Suppose that, in my attempt to satisfy my interlocutors, I continue to pile reason upon reason until I have exhausted my whole stock. It may appear that a foundationally minded interrogator could now accuse me of irrational acquiescence in those final opinions which I scraped from the bottom of my barrel. The rationale for this criticism would be that my agreeing to play the game of offering reasons to support my opinions constrains me to take such a game to its logical conclusion - which is an infinite regress.

Coherentists, however, need not accept this rationale. For those opinions given as the ultimate reasons can, without circularity, be understood as being themselves justified by their coherence with one another. Because these 'contingently foundational' opinions do not pretend to be the last word on the matter, the regress argument does not affect their acceptance.

Such a response, however, invites contemplation of the second worry. How does the criterion of coherence derive its authority to act as a justifier? I have no answer to that question. I don't mean to suggest that coherence is a bad criterion of

justification; but it seems as though it does itself play that role in a coherence theory. That is why, I think, such theories are given the name *coherence* theories: the justification of claims within those systems rests on the coherence of those systems. I think also that this is clearly the case, no matter how ‘coherence’ is to be construed: whether in one of the ways discussed above, or in any others that I have not encountered.

One more objection must be scotched, however. One could argue that, having reached those final reasons in the chain, it is not their coherence with the rest which justifies them. Simply by adjusting one’s internal perspective, so to speak, those final links can be treated now as claims that themselves require justification; such justification can be provided by other opinions within the system. Depending on whether or not the entire stock of opinions was exhausted in the defense of the claim first made, such further justification of the ‘foundational’ reasons for that claim may turn out to be question-begging. As I said above, this does not constitute a difficulty for coherence theorists, but is rather a natural consequence of their view.

We may now ask: why is it that this incipient circularity does not constitute a difficulty for coherentists? Because the argument for accepting a claim does not rest on the chain of reasons given *per se*; it rests on the coherence of that claim with opinions already held. The process of giving reasons performs a sort of ostensive function in coherent systems: it is by *showing* the coherence obtaining among the opinions held and the candidate claim, that the claim becomes acceptable. Rather than it being a matter of “If P and Q and R, then S”, the inference is better understood as “P and Q and R, and P coheres with Q and R, and P, Q, and R cohere with S, therefore S”. The second way of describing the influence of P, Q, and R on the acceptability of S does not exclude the first, but it does allow for the variation in

interpretation on the word 'cohere'. Those coherentists who interpret coherence as more than consistency but less than entailment would find the second description less troublesome than the first. I want to keep the description at this general level, because I want it to be clear that my interpretation of coherence theories as foundationalist applies to all of them.

If the second description of inferences within a coherentist framework is unobjectionable, then I would like to suggest that it reveals a foundational form. It is the coherence of S with P, Q, and R that constitutes the criterion by which S can be accepted; there is no other court of appeal. If the claim coheres (however that term gets spelled out), then it is acceptable; if it does not cohere, then it is not acceptable. One cannot go on to ask why one should accept coherence as a criterion; coherence is just taken for granted as a criterion. As I said above, I believe that this is why these theories are *called* coherence theories.

None of this should be taken as a criticism of coherence theories, however. The foundationalists face the dilemma of the infinite regress or the unwarranted terminus; that the coherentists should turn out to face this dilemma too, does not put them in any a worse position. My only concern was to show here how these theories do, after all, share a common form in at least one respect: that the acceptance of claims is based on reasons given for those claims, and in both cases, these chains of reasoning have a definite terminus.

There is another objection to be met. The way in which I have construed coherence theories above is, I hope, correct; but it fails to capture all of the coherence theories on offer. Some coherentists do not consider the justification for acceptance of claims to issue from a conscious awareness of the coherence of those claims with the system of opinions of some individual. The opinions held and the claim may, as a

matter of fact, cohere; but the individual may not be aware of that coherence. The individual knows which claims are acceptable and which are not, but the (internal) *reason* for their judgements may be no more than intuition. These are coherence theories that take an externalist approach to the coherence of a system of opinions.

I have little to say about the nature of justification in theories which take this kind of approach. There are more than just coherentists in this camp; reliabilism, for example, might be called a sort of externalist foundationalism. The proponents of these theories of knowledge do not generally speak of providing justification for accepted claims, but rather of rationalizing acceptance of those claims. Reliabilists, for example, declare a claim to be *known* when it has been produced by a reliable process. It is rational for an individual to hold a particular claim, if that claim was somehow produced by a process that has had a high success rate; this is true, whether or not the individual accepting the claim is at all aware of the operation of this process.

Externalists generally do not speak of justifying claims because they recognize a difference between the sort of account they offer and the internalist accounts discussed above. The significance of the difference can be brought out by a (purported) counterexample to reliabilist theories of knowledge. Suppose you have just been in a serious car accident. This accident resulted in severe bruising and laceration of the face; you find that your eyes are in a state of excruciating pain. Moreover, you believe that they have been damaged as a result of the accident. Under the circumstances, this seems a reasonable conclusion. When you try to use them, everything appears murky. You see a person in front of you, but cannot make out who it is. As a matter of fact your eyes have not been damaged in any way; they are just as reliable as they ever were. The murk is a result of blood in your eyes. Now according

to the reliabilist, you should trust what your eyes tell you; that is, that you know there is person in front of you, because it is true that there is a person in front of you and this information is coming to you via your reliable eyes. Yet this seems wrong; there is an intuition pulling the other way, suggesting that it would be incorrect to say that one *knows* that there is a person in front of one in such circumstances. That tension results from the intuition that knowledge claims require justification, in the sense that the individual concerned must be aware of the reasons for her acceptance of those claims. Externalists either lack such an intuition, or treat it as misguided; in any case, their approach to knowledge does not make any use of internal reasons, and so they avoid the problems of the justificatory regress.

There seems to me something to be said for externalist approaches, but they are not the main target of this paper. I think that externalist theories should not exclude internalist treatments; we may need an account of knowledge which incorporates elements of both. While externalism may be an optional component of a theory of knowledge, however, it does not seem appropriate to say that internalism is similarly optional. I want to understand what knowledge claims amount to; I think that a primary way of understanding them is via the relationship between internal reasons and claims, and I think a first step in that direction involves examining what knowledge claims *can't* be. That is the focus of the next chapter.

Our Attitude Towards Accepted Claims

Generally speaking, we accept propositions on the basis of reasons that can be mustered in their defence. Sometimes, when these propositions are offered by others, we demand those reasons in the form of an argument; at other times, we accept certain propositions as the outcome of a train of thought. In both cases of acceptance reasons play a role; yet, while it is usually because of reasons that we find certain propositions acceptable, it does not seem to me that we treat certain propositions as *true* because of those reasons. That we do treat *some* of them as true, I don't deny; whether such treatment varies proportionally with the successfulness of the justifications offered in their defence is another matter.

In what follows, I want to assume that justification issues in acceptance; that is, that justification is a sufficient condition for acceptance. It is by no means a necessary condition. I can accept something, because I take it to be true, without having any justification for it whatever. Taking something to be true, then, is likewise a sufficient condition for acceptance. If it were also a necessary condition for acceptance, we could presumably infer that justification issued in taking something to be true; that it was the purpose of a successful justification that it should result in the accepting individual's taking something to be true. But I do not think that taking something to be true is a necessary condition for accepting it; in what follows, I'll try to explain why. If I'm successful, then it would be proper to conclude that it is not a general characteristic of justifications that they should result in taking the claim justified to be true. Whatever a justification does, it cannot be so categorical.

It seems natural to think that if one has accepted a claim, then one is also inclined to assert it under some circumstances. It is hard for me to conceive of a way to understand 'acceptance' which would exclude this inclination. I accept, for instance, that I am sitting at a computer, typing on a keyboard. If anyone were to ask me what I would now be willing to assert, one of the things I would assert would be the proposition that I am sitting at a computer, typing on a keyboard. The connexion between acceptance and assertion, or at any rate a suitably qualified disposition to assert, seems so intuitive and natural that I cannot think of circumstances under which they could be properly separated. Certainly I can think of situations in which I would accept a proposition but never assert it; yet even in such situations, I cannot imagine being *unwilling* to assert something if the circumstances were different. If I were living in the former Soviet Union, there are certain things I might accept but wish never to assert; yet I could be in such a situation and still say that if things were different, I would assert those things I accept. It is this sort of connexion between acceptance and assertion which seems unbreakable.

The acceptance of this very natural connexion leads to some difficulties, however. When someone asserts something, it is generally taken as an indication that the person takes that something to be true.¹⁵ If we can assume that someone's assertion of a proposition implies acceptance of that proposition,¹⁶ it then seems that acceptance itself indicates that the proposition is taken to be true. If the proposition in question only came to be acceptable because of the reasons offered in its defence, then the reasons offered must bear some sort of truth-making relationship to the proposition. By 'truth-making', I do not mean that the reasons offered can actually

¹⁵ See *Oxford Companion to Philosophy*, entry on *assertion*.

¹⁶ I mean to exclude such cases as assertions made by automatons, or assertions made without being conscious of uttering them; I am loath to call these cases of 'assertion' at all.

make the proposition true; I mean only to indicate that these reasons somehow contribute to the attitude of the individual towards the proposition: that attitude being *taking* it to be true, rather than simply entertaining it, for example. Since in what follows I hope to show that acceptance does not issue in attributions of truth, and because of the connexion between acceptance and assertion just mentioned, I conclude that assertions do not indicate attributions of truth any more than acceptance does.

People are apt to change their opinions from time to time. This includes those opinions which they found acceptable on the basis of reasons offered for them. We might call these justified opinions, which may or may not overlap with the class of propositions picked out by the usual definition of knowledge: justified *true* opinions. Is it sensible to treat accepted opinions as propositions taken to be true? Or ought they to be treated in some other way? In other words, should someone's acceptance of, and willingness to assert, a proposition be treated as an indication of that individual's affirmation of its truth? I think that it should not; I shall now explain why.

Rather than speaking of opinions, I shall speak of theories; but this alteration is only intended to help make what follows tidier. Let us imagine that an opinion possesses the same epistemic status as a hypothesis; one can opine that 'the force of gravity varies inversely with the square of the distance', and one can hypothesize the same (if there is any distinction to be made here at all). Accepting this proposition is not quite the same as accepting the theory of universal gravitation. Accepting the *theory* involves accepting all sorts of consequences, both theoretical and empirical, which follow from the hypothesis just mentioned (along with auxiliary hypotheses needed to yield empirical consequences). Generally speaking, when people accept a

proposition, they will also assent to logical consequences of that proposition.¹⁷ In any case where they do so, it seems reasonable to treat them as accepting not only a hypothesis but a theory, or a portion of a theory (the portion which terminates in the particular consequence to which they will assent). People may work out the various consequences of theses they consciously accepted, in unconscious ways. It is because of this that people can recognize when an observation conflicts with a proposition they hold, even though that proposition itself is not empirical; they have made a comparison of their observation with the expected empirical consequences of the hypothesis(-es) they accepted. When I speak of someone accepting a theory, I mean that they have accepted some proposition or body of propositions and some, if not all, of the consequences of those propositions.

Insofar as individuals generally do change their opinions, I will present this characterization as descriptive. Insofar as such flexibility is desirable (for which one might argue on skeptical grounds), I venture to suggest that this characterization should be taken as normative as well. It is the descriptiveness with which I am primarily concerned.

Let us imagine a scientist. The scientist accepts a body of theories, which shapes her understanding of her work and guides her inquiries. Let us suppose that, in the course of research, this scientist encounters a datum the description of which contradicts one of the theories that she accepts. As is well known, the scientist has two options in this situation. She can deny the datum; alternately, she can reject the theory. The rejection of the theory would be a change in opinion. A denial of the datum could take many forms: the scientist could say that the equipment was faulty,

¹⁷ At least, we *expect* people to assent to logical consequences of propositions which they explicitly assert. When they do not, we tend to call them irrational (or sometimes just bad reasoners), since logic has a normative role to play here.

that the lab assistant was inept (much the same complaint as the last), or that some exceptional circumstance obtained and so the datum is not representative. In any case, there are many ways (of varying degrees of plausibility) of denying that some datum counts as a refutation. Sometimes such arguments are worth making; sometimes they are not. Sometimes it is better to let a theory go, sometimes it is better to cast aspersion on the data. In any case, there are certain circumstances under which the scientist will be constrained to do the one or the other.

Suppose the scientist takes the theories that she accepts as *true* theories. Suppose, as we did before, that said scientist happens upon some datum which appears to refute one of these theories. May the scientist exercise either of the two options available to her above? No: she may exercise only her right to deny the datum, by means of whatever explanatory mechanism (ad hoc or otherwise) she deems best. For, if the scientist has genuinely taken the theories she holds to be *true*, then she has reason to suspect that there could not be any evidence against them. Thus, when a counterexample is presented, the scientist is bound by reason, if nothing else, to judge the counterexample to be only an *apparent* refutation. For how could anything that is true be refuted? It cannot; sometimes, however, it may appear to be. It may be the case that the theory that the scientist takes to be true, is in fact false; but the scientist's taking the theory to be true is sufficient for the impossibility of her being able to change her opinion on the matter.¹⁸

Similar arguments can show that the scientist will be bound to change her opinion, if she takes the *datum* in question to be true. For in this case, there would be no question of doubting the reliability of the equipment, the circumstances of the

¹⁸ At least, the option to change one's mind *rationally* is not open to the scientist; for if no evidence could possibly make any difference, then any change in opinion effected would have to be arbitrary. Here we would have to include a sudden awareness of a theory's inconsistency as a sort of evidence,

recording of the datum, or any other concerns that might cast doubt on the datum itself. If the datum was being taken as *true*, then doubt would be an inappropriate (if even possible) reaction. Thus if the scientist only *accepts* her theories, but takes the data that might refute them as *true*, then she is constrained to reject the theories she previously accepted.

This may not be a wholly intuitive move, even though I do not think that the essence of the argument is new. It is a formulation of the ancient critique of dogmatism; the individual who holds to their assertion come what may is taken to be a dogmatist. Perhaps, however, it is unfair to lump together those who take the theories they accept as true, and these intractable folk. Does it not make sense that one could say 'I take p to be true, but I could be wrong'? I think it does make sense to say such a thing. What must be recognized is that assenting to such a proposition can have no impact on the reasoning process. Suppose one takes p to be true; that means that for that individual, p *is* true. Now one could in all modesty simultaneously assent to the proposition 'p could be false', which I take to be the specific content of the proposition 'I could be wrong' relative to p. Interestingly, this specification effectively tells us that p is a contingent proposition; but more importantly, recognition of this contingency cannot put us into any better position with respect to impartially judging evidence that appears to conflict with p. I may be able to assent to the possibility of my own fallibility while maintaining the truth of p; but it is hard to see how I could ever come to assent to the proposition 'p is false'. If one honestly takes P to be true, then conflicting evidence which threatens that ascription can only be rejected, if we proceed by reason alone. What sort of scenario would result, if we tried to make the fallibility of the assenting individual relevant to her treating

and say that that too could not affect the status of the scientist's acceptance (it would have to be

evidence with the impartiality it deserves? Suppose that, as the troublesome datum registers with the scientist, she says to herself, 'Well, this datum appears to conflict; now since I could be wrong about the truth of p, I suppose I had better suspend judgment on p for now'. Then we could see the bearing that acknowledgment of one's own fallibility could have on the reasoning process. That acknowledgment serves as a reason for doubting the truth of p; it is a reason for suspending judgment on p. If that is correct, then why was it that the individual ever took p to be true at all? That is, given that she is responsive to reasons such as her own fallibility at all, why did she not apply that reason to her assessment of p when p first came to be a candidate for her acceptance? Perhaps because it is *not* the recognition of her own fallibility which is playing the role of a reason for suspending her judgment; perhaps it is the datum itself which plays this role. Thus before encountering the datum, she took p to be true; afterwards, she became an agnostic with respect to p.

Then one wants to ask: How did the datum acquire this power? If she took p to be true, how could something which is merely an acknowledged piece of evidence overthrow her attitude towards p? The answer, it seems to me, is that it has no such power. Remember that we were considering the situation in which the theory p was taken as true, but the status of the datum was left unspecified. If our scientist's attitude towards this datum is agnostic, then the datum seems to lack any kind of force as a reason. If we do not want to say that the scientist's acknowledgment of a conflicting datum consists merely in her treating it as a possibility (i.e. remaining agnostic as to its truth value), then we must presume that the scientist ascribes a truth value to that datum.¹⁹ Obviously we will not want to say that she takes it to be false.

explained away, just as any empirical evidence would).

¹⁹ If we consider human cognition to operate on a trivalent logic, then being agnostic about a datum would correspond to a truth-value (which we could call 'middle'). Such a system would indeed have

The only alternative is to take the datum that produces the contradiction to be true, as well as p ; but there are insuperable difficulties with such a suggestion.

Treating a theory as true has dangerous consequences for the possibility of changing one's opinion, as we have noted; perhaps, however, we could counterbalance this effect by according the same status to data received as to the theories accepted. Suppose the datum in question is perceptual; just as theory-acceptance may imply taking-to-be-true, so perception may imply judging the percept as true. For perception involves a sort of acceptance; when one perceives something, one takes on board the informational content of the perception.

Now there are serious difficulties associated with treating perception as a form of judgment. For example, if my perceiving something involves judging it to be true, it becomes difficult to imagine what sort of account one could give of contemplating the veridicality or illusoriness of one's perceptions. This problem is difficult but it may not be intractable; if the current proposal -- to treat perception as involving a judgment of the truth of the content of the perception -- will solve our current worry, we should be willing to take it seriously. I have discussed this particular point with reference to perception, only because it affords a clear view of what is at stake here; but I mean these arguments to apply to all sorts of data, however specified. Becoming aware of an inconsistency in p , for example, could be considered a datum.

Under these conditions, when we encounter a conflict between theory and data, the symmetry is restored. Change of opinion and data rejection are equally possible, since neither theory nor data possesses an epistemic or evaluative advantage over the other.

Unfortunately this proposal fails; to accept it would entail grotesque

more descriptive power in many respects; but I take it that the tradition with which I am arguing

consequences for the character of reason. Suppose our scientist takes her theories to be true; then suppose also that she encounters some conflicting datum, which she also takes to be true. Our scientist would then be in the position of believing a manifest contradiction. If the datum conflicts with the theory held, it can do so only because one of the empirical consequences of the theory in question turns out to be the negation of the datum. Thus, if it is really a *theory* which the scientist holds true, and not just some hypothesis or set or hypotheses which *imply* the denial of the datum, then the scientist can be said to hold the denial of the datum, and *ex hypothesi* take it to be true. The scientist could then be faithfully described as taking $(A \wedge \sim A)$ as true. This is clearly an undesirable consequence of the proposal just adumbrated. It might be suggested that it is precisely because of the manifest contradiction that the scientist is plunged into a quasi-agnostic state, in which she tries to evaluate the theory and the data as objectively as possible. But the proposal did not just imply that the scientist would recognize a conflict between the data, if she were to *attempt* to take both as true; it suggested that the scientist *actually* accept a contradiction -- take both A and $\sim A$ as *true*, at the same time.

So it seems that there is no positive solution to the problem engendered by interpreting acceptance of a theory as taking it to be true. The problem is not just that it seems to introduce a possibly undesirable inflexibility into our reasoning processes; it's that *as a matter of fact* we change our opinions quite frequently, opinions which we previously accepted. So, if it were true that accepting a theory or just some stray opinion amounted to taking it to be true, then it is difficult to see what sort of rational account could be given of that tendency to change our minds. It could, after all, be the case that we change our opinions at random; this would be compatible with the

conceives of matters on the traditional bivalent model.

interpretation of theory-acceptance as taking-to-be-true, but it wouldn't be a satisfactory account. We do not want to say, for each and every one of the opinions we accepted and then changed our minds about, that we revoked our assent to them at random.²⁰ We want to say that at least in some cases we were being responsive to reasons, and that is why our changes of opinion were rational. If that is the case, then it is not possible to accept the interpretation of acceptance as taking-to-be-true.²¹

Which is not to suggest that we should have two categories of acceptance: a sort which varies according to the whims of reason, and a sort which varies according to changes in some non-rational source. Acceptance is acceptance, whether it issues from reasons or not. In a case where one has accepted an opinion on the basis of reasons, and has also decided to take that opinion as true, one has done at least two things: the first was to accept; the second was to take it as true. One need not have accorded the accepted claim this special status of being true, although one is not prevented from doing so. The important point to note is that the two are wholly distinct from one another; we shall take up these and other points in the next chapter. Similar arguments apply to opinions which were not accepted on the basis of reasons, that is, opinions that were perhaps arbitrarily assented to. One can assent to them without taking them to be true.

It is not as if, however, a scientist generally decides without reason to treat the theories she accepts as true. If one were to ask such a scientist how she arrived at her judgment, one would likely receive a lengthy explanation, involving reference to various kinds of evidence, connexions with other theories, and so on. Evidence and evidence of coherence with previously accepted theory, however, is not enough to

²⁰ Randomness is only one of a large class of non-rational ways of changing opinion. Emotional problems, mental deficiencies, and psychosis would also belong to this class.

²¹ At least, not without doing some violence to the sense of the word 'true'.

establish truth.²² So we might say instead that our scientist, being reasonable, and aware of the concerns discussed above, would never admit that she was taking the theories she accepts as *true*; just that she was taking them as more likely or more probable than others which might be available. This would seem to solve the problem of how to treat accepted theories and accepted evidence; both theory and evidence may be accepted, because when theories are merely probable, it means that their expected empirical consequences are probabilized too. Strictly speaking, in such cases, there can be no conflict between the probable consequences and the improbable.

Suppose a certain theory that she holds highly probable entails that, of two empirical phenomena A and B, A has a probability of .75 of occurring, and B a probability of .25.²³ Without wanting to suggest that we ought to always interpret probabilities as indications of frequency-distributions, it may be helpful to so think of them here. Let us suppose, then, that repeated investigation of the theory yields the following distribution: A occurs twenty-five percent of the time, while the occurrence of B is a result of seventy-five percent of the trials. Here it would seem that the results conflict with the original probability distribution. Now it is possible, even here, to say that there is no real conflict occurring; one can always reply, in these cases, by saying that while it is very unlikely that well-conducted trials should yield distributions so much at variance with what was expected, it is still possible under the circumstances and so does not constitute a real problem. It is hard to imagine, however, that anyone who took their projects seriously would say such things. Such an attitude would make it impossible to refute, or even provide reasonable cause for adjustment of, theories

²² As I said earlier, it cannot establish truth unless we have accepted a coherence theory of truth over a correspondence theory.

held to be probable to some degree. Of course, this attitude *is* exercised some of the time; when unexpected distributions occur, it can be reasonable to respond with, say, a demand for further trials. If one tosses twenty times a coin deemed to be fair, and observes that it comes up heads fifteen of those times, one does not necessarily conclude that the coin was *not* fair, after all. But one might want to toss it a few more times.

Then again, one might not; and this is the heart of the problem. Our scientist, faced with the unexpected distribution, has no reason-based motivation for accepting the unexpected distribution as an indication that the theory held requires adjustment. Since she has already accepted, albeit as merely highly probable rather than true, the theory which indicated that $p(A)=.75$ and $p(B)=.25$, she can *only* be *reasonably* motivated to *explain away* the actual recorded distribution of $A=25\%$ and $B=75\%$.

One might object that the scientist can *rationally* entertain acceptance of the less likely outcome, because that outcome is taken as merely less probable rather than falsifying. If probability valuations are supposed to govern our decision making process, however, then it would seem that even the option of *entertaining* the truth of the unlikely result would be irrational. For, to entertain the truth of the ill-favoured outcome would be to say that, despite the low probability of that outcome, it could still be appropriate to decide to adjust the theory rather than explain away the unexpected outcome – that it is entirely possible and indeed (under appropriate circumstances) reasonable to so choose it. In that case, one's decision making process is no longer being governed by probability valuations; moreover, it entails accepting that there is no correlation between high probability and truth. This would be an odd consequence, as the most intuitive motivation for treating probabilistic valuation as a

²³ I am assuming here that A and B are mutually exclusive and exhaustive of the possible outcomes for

good substitute for truth is that we presume a correlation between increasing probability and truth.

The scientist, then, seems to be barred from saying of a theory that she accepts, either that it is true, or that it is probable. In either case, the scientist will be unable to modify her opinions; and this is an undesirable result, even if the current opinions of the scientist are true.

However, one might object that this argument appears to cheat. For the above argument seems to rely on the assumption that an individual who accepts some p on the basis of its probability value, is tacitly assenting to the proposition “‘ p is probable’ is true”. So perhaps these problems can be solved by a more *thorough* application of probability. Assuming that an individual is willing to assent to p , then if their attitude towards p can be characterized as one of high probability, they should be willing to assent to ‘ p is probable’. Let us call this assertion q . What status should we accord to the acceptance of q ? We seem to face the same options as before: ‘ q is true’; or ‘ q is probable’.

The first options looks like it will produce the same inflexibility of opinion we sought to avoid; the second response does not, on the face of it, have any particular difficulties. The solution, then, may be to treat assertions such as p as nested in a series of probability statements, such as q .²⁴ Now two methods of assigning probabilities to such statements are possible: arbitrary and conditioned. If the assignments made are arbitrary, the regress of nested probability statements can be easily stopped. Each statement serves as a reason for the assertion of the statement nested within it. If our scientist asserts p , and is asked why she does so, she may reply ‘Because q ’. The fact that the chain is potentially infinite does not in and of itself

the situation in question.

make this a vicious regress. For, if the assignments are made arbitrarily, then our scientist may stop the regress at a point of her choosing. ‘Why p?’ – ‘Because q’. ‘Why q?’ – ‘Because I *stipulated* that q’. Clearly such a strategy has serious problems; although it prevents a vicious regress, it cannot be taken seriously as a good description of scientific practice. One would hope that our scientist, if she deserved that title, would have better reason for asserting p than ‘Because I said so’.

Most probabilistically minded philosophers would likely agree; thus the other option, that assignments are conditioned by various factors, must be examined. Bayesians, for example, while allowing that some probability assignments must be made arbitrarily (in particular, the initial probabilities), take it that these assignments can be adjusted by the assimilation of evidence for or against the statements under discussion.²⁵ Evidence can be said to *condition* initial probability assignments. I would like to explore probabilism and acceptance in greater detail, but to do so involves a discussion of justification as well, insofar as we will want to discuss Bayes’ theorem and its applications. Thus I reserve discussion on this point until the next chapter.

I want to turn to some issues now which might have occurred to the reader regarding this notion of acceptance. The first has to do with whether acceptance is extensional or non-extensional. Consider the proposition

p iff ‘p’ is true.

²⁴ Each of which will be given its own probability assignment

²⁵ See, for example, the discussion of assignments for prior probabilities in Howson and Urbach’s **Scientific Reasoning: The Bayesian Approach**.

If acceptance as I intend it was an extensional concept, then one would expect that p could be substituted for 'p' is true, salva veritate, since p and "'p' is true" are logical equivalents. So

I accept that p

should be logically equivalent to

I accept that 'p' is true.

Clearly this is something which I have tried to deny. This means that I must also deny that logical equivalents can be substituted salva veritate in all contexts. There is a set of contexts which we may call propositional attitude contexts. This set includes such contexts as believing, thinking, wishing, entertaining, and accepting. Now in some of these contexts, the equivalents seem to be substitutable without difficulty. Consider

I believe that p.

We can substitute 'p' is true for p, yielding

I believe that 'p' is true

without difficulty. The only thing that may seem funny about such a substitution is its redundancy, but this only shows why we can do it without difficulty. It seems implicit

in believing something that one takes it to be true, so it is natural to say that belief in p is belief that 'p' is true.

Now substitutivity does not necessarily indicate sameness of meaning.

Consider

I entertain that p

And

I entertain that 'p' is true.

Strictly speaking, "I entertain that..." followed by a proposition is ungrammatical. One always (as far as I can see) speaks of "entertaining the thought that...". So let us take

I entertain the thought that there are Martians

and

I entertain the thought that 'there are Martians' is true.

What happens when one entertains a thought? I think that what one usually does is imagine the sorts of consequences the thought in question can have. So it might be a consequence of there being Martians that if we were to observe the surface of Mars we would be able to detect the activity of life on that planet. So one might unpack the notion of entertaining the thought that there are Martians in some such way as the

following: “If there were Martians, then we would observe their activity by looking at that planet”. The question then becomes, do we need to be provisionally treating the antecedent as true in order to infer the consequent? To put it another way: we assume that there are Martians, and infer that there will be observable activity on that planet. Do we need, in assuming the former, to attribute truth? Certainly attributing truth does no harm. But is the truth attribution *necessary*? If we were constructing a truth-table for the conditional, we would assign T’s and F’s to the components in assessing the overall truth-value of the whole. It may a mistake, however, to interpret the Ts and Fs as meaning ‘true’ and ‘false’ in the sense with which this essay has been concerned. T and F are often substituted, in truth-tables, with 1 and 0, for example.

To clarify, we might look at Frege’s distinction between his content stroke and his judgement stroke in the *Begriffsschrift*. The content stroke (a horizontal line) precedes a sign which indicates what Frege calls a ‘content’. Adding a small vertical line (the judgement stroke) to the left end of the content stroke yields the judgement symbol, commonly called the turnstile. Omission of this small stroke transforms a judgment into “a mere combination of ideas..., of which the writer does not state whether he acknowledges it to be true or not...it is to produce in the reader merely the idea..., say *in order to derive consequences from it* and to test by means of these whether the thought is correct [my italics]”.²⁶ Frege clearly thought that deriving consequences from a speculation did not require attributing any kind of truth-value to the speculation.

I am skirting perilously close to a debate into which I do not wish to get entangled right now. One problem is that, in the case of so called logical truths, it seems that we do want to acknowledge the truth of such statements in something like

²⁶ Frege, *Begriffsschrift*, p.11.

the robust sense which is the focus of this essay. I am unsure of whether that inclination is to be resisted or not. I do not wish to become embroiled in a debate on the status of logical truths. All I want to do here is point out that truth-attribution is not necessary for entertaining a proposition, nor is it necessary for engaging in the activity that normally characterizes entertaining, namely, imagining the consequences of the thought entertained. So here we have a case where substitution of p for 'p' is true seems natural, since attributing truth certainly does not in any way limit the activity of entertaining.

Let us look now at

I accept that p

and

I accept that 'p' is true.

Here the matter seems to me more straightforward than in the last case. "'p' is true" cannot be substituted *salva veritate* for 'p' in the context of acceptance. Recall "I believe that p". Suppose I believe that p, I have justification for p, and that p happens to be true. Now it seems I have a prime candidate (pace Gettier and some externalists) for knowledge; indeed, from these three facts, it can be inferred that I know p. So far I've said little. Now imagine that I assert that I believe p. This means that I believe 'p' is true. Suppose I also believe that I have justification for p. As far as I am concerned, I know p. It's true that if p were false that I would not, in fact, know p; however, as

far as *I* am concerned, I do know *p*. If I only accept *p*, have justification for *p*, and *p* happens to be true, I cannot infer that I know *p*.

Now imagine that I assert that I accept *p*. This does not necessarily mean that I take *p* to be true. Suppose also that I believe my acceptance is justified. As far as I am concerned, I do not know *p*. Moreover, even if it was a fact that my acceptance of *p* is justified, and *p* happened to be true, I still could not be said to know *p*.

Now if *p* were substitutable for “‘*p*’ is true” in this context, we would have a problem. For accepting that ‘*p*’ is true is semantically equivalent to believing *p*. Given the circumstances described in the last paragraph, if we substitute “‘*p*’ is true” for *p* we will have to conclude that, in the one case, as far as I am concerned I know *p* and in the other, that I really do know *p*. The substitution is not truth preserving because we can draw a conclusion in the case of accepting that ‘*p*’ is true that we cannot draw in the case of merely accepting *p*.

There are a few other worries that should be dispelled. The first is that it may appear that I have interpreted the epistemic attitude of *taking to be true* in an unnecessarily narrow fashion. In particular, it may seem as though my interpretation of this attitude depends on a form of infallibilism. An infallibilist maintains that we cannot be mistaken about propositions we accept, if they are suitably justified. A fallibilist, in contrast, maintains that we can always be mistaken about propositions we accept, even when our acceptance of those propositions is based on what we deem to be suitable justification.

Now it might appear, by insisting that our scientist cannot rationally revise her opinion of some claim she takes to be true, that she implicitly believes that she cannot be mistaken about the truth of that claim. I discussed the possibility of possessing a fallibilist attitude and yet still not being able to revise one’s beliefs earlier, but not in

detail. It is possible, I think, to coherently accept that some claim p is true, and also accept that one could be mistaken about p 's being true; that is, given sufficient evidence or argument, one could give up p . If I denied that this were possible, I could be charged with assuming that only an infallibilist conception of a truth-taking attitude is legitimate. What I wish to show is that, despite being a fallibilist, our scientist cannot give up her belief in the claim that she has taken to be true. In other words, although she can accept her own fallibility as a matter of principle, it will not affect her reasoning in practice. Let me explain why.

Suppose I am in a room which is completely devoid of light. A switch is pulled and a single illuminant whose source I cannot detect sheds light on a small sphere sitting on a table. The sphere appears red. Let us suppose that I am sufficiently satisfied that the ball is red; I take it to be true that the ball is red. Having reached this epistemic state, I now look around for and find the illuminant. It appears to me that the illuminant itself is radiating light which also appears red. Let us imagine that I have satisfied myself that the light *is* red; I take it to be true that the light is red. Need I revise my earlier truth attribution, to the effect that the ball is red? No, since the ball could be either red or white and still appear red under a red illuminant. So I can take it to be true that both the light and the ball are red. Now suppose I trundle in a piece of equipment designed to measure the reflectance properties of object surfaces. I apply the machine to the ball and receive a reading which indicates that the reflectance properties of the ball are such that it reflects virtually all wavelengths of light. In other words, it should appear to have the colour of whatever colour illuminant it is placed under.

Now the interesting question becomes, what is my epistemic attitude towards the data yielded by the machine? In order to answer this properly, I will also need to

talk about my attitude towards the machine itself. Suppose I had previously established to my satisfaction that the machine was reliable; I take it to be true that the machine is reliable. In this case, I can continue to believe that the machine is reliable and yet reject the data it is producing here; being reliable is not the same as never being wrong. Given that I take it to be true that the ball is red, I have no reason to reject that claim yet, because I have not been given sufficient evidence or argument against it. Why not? Because the epistemic status of the data from the reliable machine cannot be on a par with the epistemic status of the previously accepted claim that the ball is red. This is because the data entail that it is not the case that the ball is red. So it is not possible to allow the data to achieve the same epistemic status as the claim that the ball is red; for it would put me in the position of believing that the ball is red and that it is not the case that the ball is red. This has nothing to do with my putting more stock in my own visual system than in the accuracy of the machine; it is a matter of temporal priority. Since I first came to believe that the ball is red, I cannot now be dissuaded from the belief that it is red, unless I am given sufficient evidence. I might even be able to specify what would count as sufficient evidence; in this case, a demonstration that the ball is really not red. However this has not yet been demonstrated; the machine, I *must* conclude, is giving a faulty reading. Because I already take it to be true that the ball is red, it would be *irrational* of me to allow the evidence from the machine to dissuade.

Now suppose that, prior to establishing that the ball is red, I had also established that the machine was not only reliable, but that it never failed: its readings were always and everywhere 100% accurate. I take it to be true, in other words, that the machine is never wrong. In this case, when I take a reading from the machine which conflicts with my belief that the ball is red, it seems I must be driven to believe

simultaneously that the ball is red and that it is not red, for the readings from the machine cannot be inaccurate, and I already believe that the ball is red. This is not so, however. I can argue that I am not interpreting the data correctly; or that the machine is producing the correct reading, but I am not seeing it. I can continue to deny that the data conflict in this way as long as I do not consider my own visual systems to be 100% accurate in the way I consider those of the machine to be accurate. If I have that much faith in my eyes, as well as the measuring device, I will have to seek another ad hoc explanation of the evidence which allows me to hold onto all of what I take to be true: that the ball is red, that the machine is giving an accurate reading, that my perceptual system is flawless, and so on.

This situation should now look like what it is: a contorted and ridiculous caricature of what would actually happen. No person, scientist or otherwise, is going to go to the lengths I have supposed myself to go to above, just to preserve the claim that the ball is red. Yet if I take it to be true that the ball is red, the scenario described above must be accurate. In order to get myself into a position where I *could* rationally reject the claim that the ball is red, I need to recognize a conflict between that claim and some evidence, or between that claim and another claim. I can't recognize a conflict until the two elements of the conflict have equal epistemic weight. The two elements can't both have the epistemic weight of being taken to be true, because that isn't merely recognition of a conflict; that's belief in a contradiction.

What would be most likely to happen in a situation like that described above is that I would recognize and consider the conflict manifested in the claim that the ball is red, and the data from the machine. I would recognize the conflict, so long as I didn't also think the machine notoriously unreliable (and thereby regarded the data as very unlikely). That is, as long as I gave the same epistemic value to the claim that the ball

is red, and to the data from the machine, I could recognize the conflict. If I was merely accepting, rather than believing, these claims, then I could recognize the conflict; this is an epistemic valuation which allows conflict recognition without collapsing into belief in a contradiction. I cannot take these claims to be true, since this is a kind of epistemic valuation which leads immediately to belief in a contradiction, without allowing the intermediate step of recognizing a conflict and allowing the possibility of defusing it.

This brings us to an interesting problem. How should conflict be characterized? Intuitively, conflict between two claims can be characterized as saying that the two claims cannot both be true. It might appear that, if what I have been arguing is correct, this intuitive characterization of conflict is not correct. For suppose I recognize a conflict between two claims. If the conflict were characterized in terms of truth, then there would be no reason to regard the *acceptance* of two conflicting claims as unstable and undesirable. So characterized, one could not take two conflicting claims as true at the same time on pain of being irrational. One could, however, take two conflicting claims as accepted or acceptable without such pain, since acceptance is not a truth-taking attitude. Now it may turn out that my thesis is untenable precisely because conflict must be so characterized. However, I think that this can be successfully overcome.

To begin with, let us imagine that conflict can be characterized relative to acceptability rather than truth. We could say that two propositions conflict when they cannot both be accepted at the same time. Such a characterization would have unfortunate consequences for my argument. For in that case, we could play out the very same argument against an accepting attitude that I have mounted against a truth-taking attitude. One would not be able to attain the position of recognizing a conflict

between a theory and some evidence against it (evidence that conflicts with a consequence of the theory), because one could not ever actually accept both the theory and the conflicting evidence and so be able to compare them.

Instead, I venture the thought that the intuitive characterization relative to truth is correct. What must happen in the case where one has accepted two conflicting propositions is this: such acceptance, when conscious, must result in an immediate recognition of the situation's undesirability. It is undesirable, for continuing to think and act while maintaining an accepting attitude towards both will lead to potential problems. One might find that one is able to consciously assent to conflicting propositions, or be tempted to act in self-sabotaging ways. In any case, recognition of an undesirable situation leads naturally to a desire to resolve it; the resolution comes by way of maintaining acceptance of one while rejecting the other. The problem is now given over entirely to whatever processes govern decision-making on such issues, which processes I won't pretend to understand. Suffice it to say that such processes can only do their job effectively when the propositions which caused them to activate are both accepted rather than taken to be true. I've given arguments as to why they couldn't both be taken to be true, but it should also be noted that these processes wouldn't work properly even if only one of the propositions was taken to be true (while the other was merely accepted). If just one was taken to be true, the answer to the question of which one to adopt would be obvious and would require no intervention on the part of any decision-making process.

This discussion has revealed an important feature of the attitude of acceptance which distinguishes it from a truth taking attitude. While it is not rational to take two conflicting propositions to be true, it *is* rational to *accept* two such propositions. This simultaneous acceptance allows recognition of a situation which is potentially

problematic, and allows resolution of that situation to begin. If all we had was a truth-taking attitude, such problem-resolution could not even get off the ground.

Interlude: Van Fraassen, Belief and Acceptance

It might be helpful, before I continue, to compare the conclusion of this argument with an attitude espoused by Bas Van Fraassen. Van Fraassen is known for his separation of the epistemic attitude of *taking something to be true* from the attitude of *acceptance*.²⁷ However, Van Fraassen's specification of the difference does not quite accord with my own:

...there are two distinct epistemic attitudes that can be taken: we can accept a theory (accept it as empirically adequate) or believe the theory (believe it to be true).²⁸

The focus of Van Fraassen's concern is the epistemic status of *scientific* theories. Although I described my thought-experiment in terms of the deliberations of a scientist, I do not believe that the argument should be limited to that particular sphere of knowledge. I mean it to apply to all hypothetical reasoning; and this requires me to take a different interpretation of acceptance than Van Fraassen. He is right, I think, to treat acceptance within a scientific context as tied to empirical adequacy, since science is not concerned with hypotheses which are not at least empirically adequate, or in any case susceptible to empirical confirmation or

²⁷ Although his choice of terminology is identical to my own, I arrived at my choices before I was aware of Van Fraassen's work on these matters.

²⁸ *To Save the Phenomena*, p. 631.

disconfirmation. The conception of acceptance which I am after is broader – it must be possible to accept metaphysical propositions in isolation as well. Metaphysical propositions do appear in scientific contexts, but in isolation they are scientifically uninteresting. Newton’s proposal that the force of gravity varies inversely with the square of the distance is, on its own, unworthy of discussion – unless we combine it with some empirical hypotheses or initial conditions, we cannot scientifically dispute it. No scientist worth their certification will accept a metaphysical hypothesis of the sort mentioned without such auxiliaries, and only after testing them. So acceptance in the sense I intend it could, in scientific contexts, turn out to be *something* like Van Fraassen’s specification of it – although I have some criticisms below.

Philosophers, on the other hand, frequently dispute metaphysical proposals in isolation. The rules for acceptance or rejection of such proposals make no appeal to empirical evidence, however, as such evidence cannot have any bearing on the truth, falsity, likelihood or what have you, of such metaphysical claims. Instead, discussion proceeds by way of attempts to show how the claim in question is internally inconsistent, or vague, or perhaps without meaning. It would be just as dangerous to bring to such disputes a prior *belief* in the contested claim, as it would be in scientific contexts. For the attitude of belief – of taking to be true – will prevent the believer from making good sense of an objection based on inconsistency, incoherence, or a charge of meaninglessness – indeed, of any objection whatever.

It might be mentioned here that I do not actually have a different view from that of Van Fraassen; I merely extend his notion of acceptance beyond the scientific context - where susceptibility to empirical evidence is of central concern - into the realm of reasoning in general – where it might be said that evidence does not always matter. Perhaps Van Fraassen, if he agreed with such an extension of his concept at

all, would not find himself at odds with my views in any way. This is possible.

However, Van Fraassen's motivations for distinguishing belief from acceptance seem to arise from more traditional skeptical concerns than my own. As the argument of the last section was intended to show, my motivation arises not from a desire to avoid acquiescing in views for which we have insufficient warrant. Rather, it is because taking something to be true leads to – as far as I can see – insuperable difficulties with respect to changing opinion. Now it might be said that my argument, if successful, could hook into the skeptical tradition by giving a practical motivation for its worries. Let us take, for example, a typical skeptic, whose position constitutes a warning not to take as true anything which can be doubted – anything for which there could be contrary evidence, even if none actually exists. We might ask of this individual why it would be so bad if we did allow ourselves to treat as true some claims for which there remains some possibility for doubt. "Alright," we could say, "so you think it violates some canon of rationality to indulge ourselves this way. Why should we take that canon seriously? If we're wrong, then experience will eventually show us this; so our gamble will have failed to pay off. So what? Life is a gamble. Possibly discouraging, but if that's the case, what of it? So we need to devise a scheme on which what you call insufficient evidence will count as sufficient, at least in some cases. Where is the difficulty?" Something of this sort of response is commonly given nowadays to such skeptical caveats. It can be found in the probabilist schools of the structure of reason, and in general can be found in the fallibilist's tip of the hat to the skeptic when they say "I take this claim to be true, but I could be wrong". The skeptic could now reply with something like my argument – that taking something to be true will yield an inflexibility of opinion which may be undesirable. Its undesirability stems from the concerns which the skeptic has always had; for if we settle for belief in a falsehood,

we shall never be able to recognize it as such even in optimal evidential circumstances. Now, if there existed a justification for a claim which was sufficient, by the skeptic's lights, for believing that claim, then the inflexibility would obviously not matter; who cares if one's opinion has become rationally unrevisable in the light of further argument, if one has a hold of a genuine truth? The skeptic must agree with this; his worry about believing a claim on what he considers insufficient evidence is that the claim might be false, and one would no longer have the option of revising the belief in the light of further argument.

In any case, Van Fraassen makes no appeal to any argument of the sort I have given. Take, for example, his critique of inference to the best explanation [IBE]. Van Fraassen begins by saying that

...its [IBE's] purport is to be a rule to form warranted new beliefs on the basis of the evidence, the evidence alone, in a purely objective manner. It purports to do this on the basis of an evaluation of hypotheses with respect to how well they explain the evidence, where explanation again is an objective relation between hypothesis and evidence alone.

It cannot be *that* for it is a rule that only selects the best among the historically given hypotheses. We can watch no contest of the theories we have so painfully struggled to formulate, with those no one has proposed. So our selection may well be the best of a bad lot. To believe is at least to consider more likely to be true, than not. So to believe the best explanation requires more than an evaluation of the given hypothesis. It requires a step beyond the comparative judgement that this hypothesis is better than its actual rivals...For me to take it that the best of set X will be more likely to be true than not, requires a prior belief that the truth is already more likely to be found in X, than not.²⁹

Van Fraassen goes on to argue the merits of various sorts of response to this criticism, which he does not suggest is his own creation. Before launching into the fray, however, he notes that all of the responses *accept* the criticism, arguing that inference to the best explanation must be admitted as the correct characterization of our

practices notwithstanding.³⁰ It may be that Van Fraassen has unfairly characterized IBE as a sufficient condition for taking a hypothesis or theory to be true; Gil Harman, for instance, can be construed as offering the best explanation as a mere *constraint* on belief in a hypothesis – in other words, IBE turns out to be only a *necessary* condition for belief, not a sufficient one. However, some individuals do seem to treat it as a sufficient condition, David Wiggins among others; and this is certainly its more controversial form. Explanatory power can still be objected to as a necessary condition, but that form doesn't seem to be so troublesome for Van Fraassen.

The objection itself shows us what Van Fraassen's concerns are. He does not want to allow himself the luxury of taking a hypothesis to be true, or at least more likely to be true than not, when it isn't clear to him that he has obtained all possible evidence to enter into the computation. He doesn't go on to say that he's worried about the inflexibility of opinion such an indulgence might render; he is concerned only with avoiding an inference which is not justified, or at least insufficiently justified. Van Fraassen accepts the dictum that 'the evidence is never all in'; that is, there is always the possibility that there is more to be collected. Thus in his estimation, we are always in the position of having insufficient evidence for a claim. He resolves his difficulty by making a distinction between believing a claim, or taking it to be true, and merely *accepting* it as empirically adequate.

If we modify the dictum 'the evidence is never all in' to 'the arguments are never all in', we could extend Van Fraassen's concept of acceptance to include claims made and disputed outside the realm of the empirical sciences. If I were simply expressing Van Fraassen's views in this paper and merely trying to modify the distinction he has made to include disputes about such extrascientific claims, I could

²⁹ *Laws and Symmetry*, p.143.

stop here. However, I would then have to answer, as Van Fraassen himself must still have to answer, the response given above to the skeptic's injunction against taking something to be true. That response was just that there can't be anything intrinsically wrong with taking something to be true that might in fact not be; we can always revise later. For Van Fraassen's position does not illuminate why we should want to avoid taking claims to be true. The motivation for his distinction seems to rest entirely on wanting to satisfy the skeptic, but it does not help us understand *why* we might want to satisfy the skeptic. So the distinction doesn't seem to do any real work. Since Van Fraassen is a probabilist (although not a Bayesian), why doesn't he just say that probabilism gives us a way of deciding how and when we should believe certain claims? One might think he could respond by appealing to some sort of epistemological parsimony: not to admit to anything more than absolutely necessary. It's true that of the sorts of claims he has in mind, it is only necessary to admit that they are empirically adequate, if we are admitting anything at all.

The trouble is that we now face a second-order problem of the same form as we've been discussing. Suppose we accept some hypothesis as empirically adequate, but not as true. Are we not then taking it to be true that the hypothesis is empirically adequate? If so, then we might not be able to revise our opinion that the hypothesis is empirically adequate. Now the notion of empirical adequacy applies only to the evidence gathered so far (call this set of evidence E), so I would be wrong to suggest that given a larger set of evidence, we could not revise our opinion about the truth of the empirical adequacy of a hypothesis. For even if on the supposition of this larger set of evidence (call this E*) it turns out that the hypothesis is not longer empirically adequate, this should not disturb the truth of the hypothesis being empirically

³⁰ Ibid.

adequate with respect to E. What is worrisome is that, having taken it to be true that the hypothesis is empirically adequate with respect to E, if someone were to show that after all the hypothesis was not empirically adequate even with respect to E, the believer would not be in a position to adjust their opinion of the hypothesis accordingly.

So Van Fraassen's characterization of acceptance as recognition of empirical adequacy is not quite what is wanted, even in the scientific context. It isn't even possible to say that acceptance amount merely to recognition of some claim as acceptable, for the same reasons. For imagine that we said this of a claim. Without further qualification, it sounds as if we must mean that we are taking it to be true that the claim in question is acceptable, by whatever canons of rational acceptance we have established. Just as it might turn out that some claim which appeared to be empirically adequate to account for some set of evidence E was actually not so, it might also turn out that, by those very canons of rational acceptance previously laid down, the claim we are investigating is not after all acceptable. If we have already taken it to be true that the claim is acceptable in the light of such canons, we would not be in a position to alter that opinion of the claim given further demonstration that it is not acceptable.

If there is anything lacking in my distinction between acceptance and belief, it is that I have offered no positive characterization of it, as Van Fraassen at least attempts to do. All I can offer is a crude outline of the necessary and sufficient conditions for acceptance, as I see it. However, these conditions do not differ in any great degree from the conditions I would lay down for taking something to be true, except that in the case of acceptance, there is no condition of the form '... and the agent must be willing (in appropriate circumstances) to assert that the claim is true'.

I'm going to lay out the conditions as I currently see them. First, the necessary conditions. If someone accepts a claim, then they cannot also continue to accept a conflicting claim: the configuration is not stable. Now there is an immediate qualification that must be made, since people often do accept claims which conflict, without there being any loss of stability. However, when they do so, they are not aware of any such problem. It is unlikely that any of us will be able to do a pairwise comparison of every idea we hold at one time, both for the reason that the mental labour would be considerable, and that the set of claims we accept is constantly changing. So acceptance must be constrained relative to an awareness of conflict: we cannot accept two conflicting claims, if we are aware of that conflict.

Justification and Acceptance

The foregoing arguments, if accepted, must affect our understanding of the value of justification with respect to claims. Whatever positive claims we may want to make about the notion of justification, there are a few ways of understanding it that must be excluded.

Justification, as the first chapters of this work sought to show, has usually been construed as something which guarantees or supports the *truth* of a claim. There is a persistent strand of philosophical thought, beginning with Plato, which claims that reasons provide iron-clad grounds for taking a claim as true. Other philosophers, beginning (perhaps) with the academic skeptics, have held that reasons can only provide a measure of confidence in the claims they support. Moreover, these philosophers often claim that there is no greater support to be had; that is, it will never be the case that a justification could issue in an attribution of truth. Generally speaking, this view has the corollary that more reasons mean more support, and therefore greater confidence. The sorts of support which justifications can offer towards claims are subsumable under the titles fallibilism and infallibilism. These are,

of course, names for broad philosophical attitudes towards knowledge claims; but they carry attitudes towards the power of justifications in their train. The argument given in the previous chapter against taking a claim to be true is one that counts against both fallibilist and infallibilist approaches to the problem of knowledge. I will discuss each of these in turn.

If justification guarantees truth, then one who possesses a justified opinion (and knows that they do so) must also *take* that opinion to be true. Given that we generally do not take our opinions to be true, even though they are justified and we have accepted them on the basis of that justification, it follows that justification as we understand the term does not guarantee truth. This last was not a novel point; dissatisfaction with the infallibilist conception of knowledge became widespread long ago. Usually, though, the route to that dissatisfaction proceeds by way of skeptical arguments, like those discussed in chapter one. Such arguments try to show how treating justification as a guarantee of truth can lead to problems, and thereby warrant a normative claim: namely, that treating justification as a guarantor of truth is a bad way to proceed. Having shown why I think justification, conceived infallibilistically, is *normatively* incorrect with respect to theories accepted, I would now like to show how the infallibilist is wrong from a *descriptive* perspective as well, just by looking at how we are willing to treat our own acceptable opinions.

As I said before, I assume that justification is a sufficient condition for acceptance. If that is so, then it seems as though justification cannot be seen as guaranteeing truth. Justified claims form a subset of acceptable claims. There are claims which we find acceptable without justification, but there are no justified claims which are unacceptable. One might object that there are such claims. For instance, I may have a large body of evidence and argument, which would count as a

justification by my own lights, towards the claim that smoking causes lung cancer. Nonetheless, I go on puffing away. Doesn't it seem that I have not accepted the claim that smoking causes lung cancer, even though I would be willing to call it justified? I think not. It is true that my apprehension of the claim and the evidence for it have not changed my behaviour; but that isn't the same as saying that I have not accepted the claim. Without wishing to get into a protracted discussion on the philosophy of action, I seem to recall that the usual minimum requirements for action are a belief and a desire.³¹ Since all we have in this case is an opinion (for these purposes, a belief), a crucial element needed for an action is missing: a desire for (perhaps) a non-lung-cancerous state.³² Whatever else I might need in order to act, I need at least one belief and one desire; all I have is the belief that smoking causes lung cancer. It is perfectly possible to accept that smoking is bad for you and to keep on doing it anyway; although justification issues in acceptance, acceptance need not issue in action. There may be other counterexamples to my thesis that all opinions which one considers justified, one thereby accepts; in the absence of those counterexamples, however, I think the thesis should stand.

If it could be shown that all the opinions which one finds acceptable were open to revision, then the contention that justification does not lead to taking claims as true would fall out quite naturally. It is not quite as simple as that, however. There are some individuals who hold opinions which they treat as true;³³ moreover, as one might expect, they are unwilling to revise these opinions in the light of evidence. I

³¹ Clearly, these are not sufficient conditions for action; but on most accounts of action I have read, belief-desire pairs are treated as necessary conditions for actions. Presumably on a Churchland-style account of action (if one exists) they won't be, but I take it that this is an uncontroversial assumption to make.

³² Physical addiction can also play a role here; but the lack of a desire for a non-lung-cancerous state would be a sufficient explanation of my failure to stop smoking in this case, which is all that is needed for this argument.

have minimized my use of the word ‘belief’ in this essay, preferring the more neutral term ‘opinion’; I think that the word ‘belief’ connotes a sort of commitment which ‘opinion’ does not. As such, I think it is appropriate to reserve the word ‘belief’ for those opinions which are not merely accepted but also taken to be true. It might appear that the very existence of beliefs undermines my claim to have shown infallibilism to be descriptively false. For, if there are accepted opinions which are taken to be true, then it seems plausible that there may be justified opinions which are taken to be true – that is, justified beliefs. Indeed, it seems almost obvious that there should be such beliefs. Anyone who is willing to defend their belief against all comers is likely to have a defence already worked out; that is, they very likely hold all sorts of reasons which they can cite in support of this belief. Thus there may be justified opinions which are taken to be true – i.e., justified beliefs. Nonetheless, I think infallibilism is still descriptively false. It may be that there are justified beliefs; but that does not mean that the justification of the belief is what made it into a *belief* rather than just an accepted *opinion*. If it were, then all justified opinions should be justified beliefs. As I tried to bring out in the previous section, that doesn’t seem to be the case. Most justified claims remain opinions, although accepted opinions. If they did not, we could not change our opinions rationally – that is, in response to new arguments or evidence. These opinions are justified; and it is possible that an individual holding a justified opinion *may* decide to treat that justified opinion as being true. His doing so, however, does not suggest that it was the justification – or the acceptance, for that matter – which led him to take it to be true. He could have chosen to treat in the same way a belief which he accepted but for which he had no justification whatever; in that case he would have a belief, but no justification for it.

³³ It is arguable that we are all members of this class of individuals at one time or another, with respect

Someone who holds unjustified beliefs is usually called a *fideist* – someone who is satisfied with having faith in the truth of some claims, rather than reasons for holding them. Yet, given the relationship between justification and taking-to-be-true just discussed, it seems appropriate to call anyone who has a *justified* belief a fideist as well. Since taking something to be true is not a necessary condition for acceptance, then it cannot be a necessary condition for justification, either.³⁴ If that is the case, then whenever someone takes something to be true, they are doing so independently of the justification they may or may not have for it. Such an individual may cite a justification for the claim, if he is challenged to defend it; but that justification, whatever else it may be doing, cannot support his taking the claim to be true; so it does not explain why he takes it to be true. The origins of such an attitude may be psychological in nature, and it may turn out that the justification he accepts for the claim has, as a result, *some* role to play in the origin of his dogmatism.³⁵ But it cannot be the case that the justification in and of itself is the sole cause; if it could, then every instance of a justified opinion would be a case of justified belief, that is, each would be a case of a justified accepted opinion taken to be true. Thus a justification is not to be regarded as something which contributes to belief in a proposition.

Some treatment of a distinction between believing claims and merely assenting to them or accepting them, as I have mentioned, has been explored before; the earliest exploration extant is that of Sextus Empiricus in his discussions of ancient skepticism. The academic skeptics, in particular, tried to explain their ability to live

to some of our opinions.

³⁴ Or better, a necessary *consequence* of justification.

³⁵ We could, after all, be so constituted that the ways in which reasons concatenate in us will sometimes result in an attribution of truth with respect to a claim. An attribution which arose in this way, however, seems to me arational.

by their skeptical principles by appeal to ‘plausibility’.³⁶

What may be novel, however, is that the same argument applies to many fallibilist conceptions of justification as well. Fallibilists are usually described as those who take justified claims to be only probable rather than true; the position includes those who treat at least *some* justified claims as true, but at the same time admit that they ‘could be wrong’. The usual arguments adduced for these theses are skeptical in nature. Because we cannot verify our claims conclusively, we cannot treat our assertions as true; or, if we can treat them as true, we must admit that we could be wrong. It can be argued that such skeptical arguments have only a normative force, and that they do not capture our actual practices. Now it is clear from what has been said that I take the position that we do not treat our claims as true, so I do not agree with those fallibilists who assert that our justified claims are ‘true but possibly false’. Of the other collection of fallibilists, I take the first part of their claim – that our attitude towards our claims is not a truth-taking attitude - as descriptively accurate; however, the second part – that our claims be treated as probable - is not correct. In order to make this position more plausible, however, it will help to discuss probability in relation to justification, so that what I intend can be more clearly seen.

There are two ways in which I can conceive of probability entering into the relationship between justification and acceptance: either *consciously* or *unconsciously*. By ‘consciously’ I mean that the individual who wants to say ‘p is probable’ has come to that conclusion by way of a conscious calculation of the probabilities involved. This conscious calculation could be very rough and informal, or it could be as detailed as a calculation proceeding by way of Bayes’ theorem. By

³⁶ See the selections from Sextus Empiricus’ works in Inwood and Gerson, **Hellenistic Philosophy: Introductory Readings**, pp. 185-193. Plausibility does not seem to have much to do with probability

‘unconsciously’ I mean that there is some process which conforms to a calculation in terms of probability; here again I am thinking of the application of Bayes’ theorem. This process, however, would not be one to which the individual concerned had any epistemic access. The way the calculation is performed, incidentally, need not actually *correspond* to the probability calculus; but if it does not, then the results yielded by such an unconscious calculation must at least conform to the results which a calculation made by way of the probability calculus would yield. I wish to discuss both the unconscious and the conscious construals of probability, in order to see whether either of them can serve as an account of how accepted claims can be regarded as probable. The argument against the conscious construal is fairly short; that against the unconscious construal is rather more lengthy. I shall begin with the shorter of the two.

For convenience, let us imagine that the calculation which our agent consciously performs is, in fact, a calculation involving Bayes’ theorem.³⁷ Thus, given some hypothesis *h* to which the individual assents, let us imagine that they have assigned a probability value to it. Then they might reasonably be expected to assert that ‘*h* is probable’; indeed, they would probably be willing to specify the precise degree to which they consider *h* probable. This seems perfectly reasonable; but it won’t serve as a general account of the status of accepted claims. There are many cases in which we might want to say that *h* is probable; and it may even be that the way in which we arrive at that judgement of *h* is through a consciously performed Bayesian calculation. Can we say, though, that all propositions are judged in this way? In particular, what about those propositions which articulate the axioms of the

in these selections, which distinction deserves some attention; I will not, however, attempt any exegesis here.

probability calculus and Bayes' theorem, on which so much reasoning now *ex hypothesi* depends? Perhaps we would be willing to call those propositions probable too; if we do, however, it cannot be on the same basis as we have judged other propositions, like *h*, to be probable. If we used a Bayesian calculation in order to determine the probability value of the propositions on which the procedure for making Bayesian-style calculations here depends, we would engender an explanatorily unilluminating circle. How did we come to accept the axioms of the probability calculus and the elements of Bayes' theorem? If the answer is 'because the Bayesian calculations we performed told us that these propositions were probable', then the integrity of the justification has been violated. Nothing here turns on this difficulty being couched in terms of the Bayesian calculus in particular; any calculatory scheme, as long as it is pursued consciously, will yield the same difficulty.

Now it might be objected that this is not really a difficulty. In other circumstances, Bayesians in particular have been known to appeal to a certain sort of arbitrary introduction of probability valuations, in order to get their theory off the ground; I shall explain. In order to perform Bayesian calculations, the proposition under scrutiny must have a probability value in advance of being submitted to a calculation (often referred to as the *prior* probability). But many propositions are put forward as conjectures, with no prior probability attaching to them; in such cases, Bayesians can allow a probability to be chosen arbitrarily.³⁸ This is possible because it is a property of the Bayesian system that prior probability distributions make increasingly little difference to the final values attributed to propositions, as the evidence mounts. This is particularly apparent in cases where many different

³⁷ Nothing turns on this presumption, so long as we recognize that whatever procedure *is* being followed conforms, in ideal circumstances, to the results of the calculus.

individuals will come to attribute similar probability values to certain propositions, even though their original assignments were very much at variance with one another. Can this strategy be applied in our case? That is, could our conscious Bayesian have given arbitrarily high probability valuations to the propositions constituting the axioms of the calculus and Bayes' theorem? Let us suppose that they have done so. Having accepted these principles, then, this individual could go on to calculate the values of other propositions without worrying about the justifications for those principles being ultimately circular.

The solution is fine, so far as it goes; but it does not go very far. For now we must look at another principle which, it seems, is tied up with the use of a probability calculus, when that calculus is being used to explain the character of assertions. That principle is what may be called the principle of acceptance. In the case of our hypothetical conscious Bayesian, this principle must take a probabilistic form. Some sort of threshold will have been assigned, such that 'For any proposition p , p is acceptable iff the probability of p is $>x$ '. The principle must take this or a similar form if we are to characterize probability as a general account of how accepted propositions should be regarded. How is it possible for our individual to accept the principle of acceptance? She cannot do it on probabilistic grounds, on pain of begging the question; that is, it would be illegitimate for her to say that the principle of acceptance has a particular probability value, and that her acceptance of it issues from that valuation. Her acceptance of the principle could only follow from its having a particular probability if she had already accepted the principle. There must, then, be other grounds for her acceptance of that principle. Even if the acceptance of every

³⁸ There are difficulties associated with this procedure; it is difficult, for instance, to allow a prior probability of zero, because no amount of evidence can change such a valuation. But we'll ignore those difficulties here.

other proposition could be explained by the conscious application of the Bayesian apparatus, or at least by reference to an assigned probability, however arbitrarily determined, this one could not; and so the account fails as a general account.

This is not the only account on offer, however; we need to look at the possibility that acceptance is probabilistic, but unconsciously so. One of the advantages of such an approach is that we could avoid all the difficulties involved in explaining the acceptance of the first principles needed to get a probabilistic account airborne. Thus, we could assume that such principles as the principle of acceptance mentioned above were in some way unconsciously held,³⁹ so the question of an individual's conscious acceptance of such a principle need not arise. Such a principle could indeed be conditioning her acceptance; but she would not be aware of it.

This characterization of how we might work probabilism into an account of the status of accepted propositions seems to me acceptable, but it faces a difficulty. What motivation do we have for treating accepted propositions as probable? Presumably it is because we are apt to make assertions of the form 'p is probably true', or 'p is unlikely', and so on. If that is the motivation for introducing probability as a substitute for truth in the matter of acceptance, then the above suggestion for weaving probability into the explanation of these affairs will not help. It will not help, because that suggestion makes it a psychological fact that propositions we accept have a probability; but this is a probability we are not aware of. All we are aware of is our acceptance. In that case, our *conscious* ascription of probability values to propositions doesn't seem to be appropriately motivated by the psychological description.

We can scotch that particular difficulty, however, by observing that it need not be the case that the accepting individual is unconscious of the probability of the

proposition in question; all that is necessary is that she be unconscious of the processes by which the valuation came to be made. Then she can be aware of both her treatment of the accepted proposition as accepted, and her treatment of it as probable. With that caveat in mind, let us see if we cannot then make sense of acceptance as probabilistic.

If it aids understanding, one might think of the Bayesian schema as a (perhaps normative) description of the operations of the brain, where the brain is an analog computer. The mental states of accepting a hypothesis, believing a hypothesis, and other propositional attitudes, supervene on the brain states, which states are determined by the *operations* of the brain, which operations are *ideally* captured by the probability calculus by way of Bayes' theorem. Given such a picture, we can imagine the following scenario. An individual accepts some proposition *p*. This proposition also has a probability value, of which the agent is aware, although perhaps not to a precise numerical value. According to the Bayesian schema

$$P(h/e) = \frac{P(h).P(e/h)}{P(e)}$$

the value of the hypothesis should change as the evidence rolls in.⁴⁰ Let us suppose that the evidence is so rolling, and that it is of a sort that will adjust the probability value of the hypothesis in question; in particular, the evidence is adjusting the value of the hypothesis downwards. Now let us suppose that the hypothesis has gone from being likely ($P(h) > 50\%$) to unlikely ($P(h) < 50\%$).⁴¹ The individual might now be willing to say 'h was probable, but is now improbable'. I think that this probability

³⁹ Again, psychology intrudes on such an account: the 'principle of acceptance' may be just a manifestation of how we are put together.

⁴⁰ See Howson and Urbach, *Scientific Reasoning: The Bayesian Approach*, p.28.

ascription, although it may vary with the changes in the status of h on the neural level, will not explain acceptance. Although a probabilistic description of the operations of our brain may be correct, I do not think it can issue in our *accepting* probable propositions, even though it may appear to do so.

Just as we may treat some accepted proposition h as probable rather than true, so we must be willing to describe the status of our treatment of h . ‘ h is probable’ is an assertion; if we cannot say of the asserting individual that they are taking this assertion to be true, then, in advance of other options which might be put upon the table, we must see if they can be regarded as taking the assertion to be probable. If that is so, then they should be willing to assert that “‘ h is probable’ is probable”. The same goes for the treatment of the assertion just quoted: it too must be probable, and the asserting individual should be willing to agree that it is so.

Let us recall, however, the supposed motivation for regarding accepted opinions and assertions as items which should be treated as probable. What motivated it was this: that we are willing to say of certain propositions that we *consider* them probable. Under our current hypothesis of unconscious probabilism, the conscious ascription of a probability value to an assertion is supposed to be grounded on the unconscious operations of the brain and the resulting brain states. Now in our earlier discussion of probability and acceptance, we agreed that if acceptance is probabilistic and if the individual is aware of the probability, then they should be willing, in addition to assenting to p , to assent to ‘ p is probable’ and to “‘ p is probable’ is probable” and so on. So we must ask, if the current model is accurate, whether this can possibly be the case. Our asserting individual is willing to assert that h , and is also willing to assert that h is probable, and that ‘ h is probable’ is probable, and each of

⁴¹ This scheme is given purely by way of example; a scale this simple does not reflect the fine

these assertions and their attendant valuations must be products of the operations of the brain. In the case of *h* itself, the particular value consciously attached to it is a result of the interaction of those brain states which constitute *h* (or upon which *h* or its representation can be said to supervene) with some other brain states which must count as the neural analogue of some evidence. The resulting brain state is one upon which *h* still supervenes, but after the interaction with the evidential states *h* is considered differently: perhaps more probable than before, perhaps less.⁴² This consciously considered difference *ex hypothesi* reflects some difference in the subvenient brain state. So far, so good; yet we now have to tell some sort of story about the rest of the assertions, such as 'h is probable'. How did this assertion come to have its current consciously considered value? Just as with *h*, so too must 'h is probable' be probable relative to some evidence; but what sort of evidence could we adduce? The only sort I can imagine is something like the following. In the brain of our assenting individual, there is a brain state upon which *h* supervenes. This brain state possesses some characteristic which varies upon interaction with certain other brain states, and to which characteristic the individual concerned has some sort of conscious access. That access manifests itself in her willingness to describe *h* as probable. Thus she is willing to assert that *h* is probable. Let us call this assertion *h'*. What makes that assertion allowable must be an interaction between the brain state upon which *h* supervenes, and the brain state upon which the new hypothesis *h'* supervenes. Prior to their interaction, however, we shall have to assume that the individual assenting to *h* would not, at that point, necessarily assent to *h'*. Although *h'* may have been given an initial probability assignment, it need not have been one that issues in acceptance. This might seem to conform to what may actually happen if one

discriminations we make in assessing probabilities. It suffices, however, for this purpose.

played this scenario out with real subjects. After eliciting the assertion ‘h is probable’ from an individual, one could ask, “And what about ‘h is probable’? Is that probable, too?” In such a situation, one would expect some slight hesitation from the subject before they answered in the affirmative: this could be considered as an indication of brain activity in the service of calculating the probabilities.

In any case, after the interaction of the subvenient brain states for h and h’, our individual would presumably be willing to assent to h’. Thus we have here a solution for the regress first mentioned in the previous section. If accepting h entails accepting ‘h is probable’, then accepting ‘h is probable’ entails “‘h is probable’ is probable”. Whereas in the previous section this looked like a rather dubious way of getting on, it now seems perfectly explicable.

However, the tidiness of this strengthened articulation of unconscious probabilism comes at the expense of finding a general account of acceptance. Recall the characteristic which the accepting individual has some sort of conscious access to; this characteristic is the one which allowed the individual to say ‘h is probable’ or ‘h is improbable’. Presumably, then, this characteristic either constituted, or is tied to, the probability valuation which attaches to each proposition. This was, after all, what was wanted: it was supposed to be a virtue of the account that it should explain our normal habit of saying ‘h is probable’ for at least some of the h that we assent to. The fact that in some cases of acceptance, we are reluctant to say that we think that the proposition assented to is probable, we could explain away by saying that the degree of probability in such cases is so high that the individual is having trouble distinguishing their acceptance from a sort of conviction (which is not to say that in such cases we would want to say that those statements were taken to be true). This

⁴² It could, of course, have stayed the same, but that is not an interesting case.

would be merely an *epistemic* difficulty. The real trouble issues from the fact that we do not accept all of the propositions which we would want to call probable.

Consider the National Lottery. It is hard to imagine that those who play the lottery think that they have any significant chance of winning. If each person buying a lottery ticket were asked whether they thought the probability of losing the next draw was high, they could reasonably be expected to reply that they thought the probability was *extremely* high.⁴³ Yet it would seem unreasonable to expect that the same individuals would assent to the proposition that they will fail. If they accepted that they were actually going to lose on the basis of the high probability of their losing, it is hard to see why they would then buy tickets at all.⁴⁴ Instead, I conclude that those individuals who play are well aware of the low chance they have of winning the stake; in each case, however, they do not assent to the proposition that they will lose. If we replace 'This ticket will lose' with the symbol *p*, then we can see that lottery players will assent to 'It is extremely probable that *p*' While failing to assent to *p*. Evaluations of high probability, even when justified, are not sufficient for acceptance.

It could be objected that this is not the right sort of example, because the low probability of winning the Lottery is a purely formal result of the probability calculus. Since we have been discussing Bayesianism and the effect of evidence on the attitudes we take towards propositions, it seems only fair that an example which expresses the difficulty relative to that sort of justification should be included here.

⁴³ All those people who I asked during my own rather unscientific enquiry did, in fact, assent to the extremely high probability of their losing the next draw. Some were able, or at least attempted to, put a numerical figure on their chances.

⁴⁴ I assume that said individuals do not have a compulsive and conscious desire to throw money away; I assume also, perhaps unfairly, that the support that the Lottery gives to the arts is not the reason that people decide to participate in it. In fact some people may play for that reason, and that reason alone; yet I have had contact with individuals who play regularly, but exhibit no interest whatever in the sorts of cultural activity which the Lottery supports. Not does it seem reasonable to suppose that these individuals do accept that they will lose, but play only for the few hours of fantasy which purchasing a ticket affords them. In order for the fantasies to attain that plausibility which makes them enjoyable, the individual cannot also accept that they *will not win*; I speak from experience.

This is not to suggest that the low probability of winning the Lottery is accepted without reason; the reason for accepting it is that it is a mechanical result of the probability calculus. However, these sorts of propositions seem like those which belong to the class of logical and mathematical truths, which class I explicitly excluded from consideration in this paper.⁴⁵ Does that mean that all those punters who assented to the low probability of their winning, have already studied and absorbed (not to say accepted) the technical arcana of mathematical probability? It should not. Rather than saying that such evaluations are a formal result of the probability calculus, then, I ought to have said that they were the result of some intuitive form of reasoning which is partially expressed and refined by the formal calculus. One need know very little mathematics in order to conclude that, if one holds one ticket out of fifteen million, the chance of that ticket being the winner is low. In any case, the above objection still applies: an example is required which involves empirical evidence, rather than the mathematico-logical reasons that this one does.

Take the proposition 'It will rain tomorrow'. Looking out of my window, I see clouds massing under a grey sky, and as the light fades I see the trees thrashed by the wind. If someone were to now ask me how I would regard the proposition that it will rain tomorrow, I would naturally reply that it was highly probable that it will rain tomorrow. Yet do I want to say that I accept that it shall rain tomorrow? Somehow, that seems wrong; it seems wrong, even if I keep at bay the tendency to conflate *acceptance* with *taking to be true*. I don't want to assert that it will rain tomorrow; although I do want to assert that it is probable that it will rain tomorrow. In other words, I want to assert that p is probable (where p is 'it will rain tomorrow'), but not p.

⁴⁵ See introduction.

When we considered the conscious application of the probability calculus to the mechanics of acceptance, we discovered that acceptance of the probabilistic principle of acceptance could not itself be based on a probability valuation. In the unconscious scenario, I have been tacitly assuming that such a principle would somehow be 'hard wired'. That now seems impossible, since there are cases in which something can be judged probable without being accepted. That problem could be solved if we could find some internal demarcation criterion for probability ascriptions: one which would allow us to identify those instances in which some proposition which we judge probable is by that fact acceptable too, and one which hopefully would explain why it is that we use the same word to describe our attitude towards some propositions which we do not accept. But I do not see any way of so demarcating these uses of the word; so I must conclude that for us, conscious probability ascriptions and acceptance are not connected. If that is accepted, then the unconscious account as given cannot be correct; for it quite clearly ties conscious probability ascription to acceptance. Some sort of unconscious probabilism may still be true; perhaps the variety we ruled out of court earlier, in which the accepting individual was aware only of their acceptance and not of the probability valuation itself. If so, that would require a separate explanation of our practice of calling certain propositions probable. In any case, however, it seems that we cannot *consciously* regard accepted propositions as *probable*, any more than we could regard them as true, although for quite different reasons.

If propositions accepted on the basis of reasons cannot be themselves regarded as either true or probable, it seems that the justifications which led to their acceptance cannot be regarded as contributing to their being so regarded. That is, as I have said before, justification issues in acceptance, but not in attributions of truth, nor in

attributions of high probability.

The natural question to ask now, is what precisely acceptance *does* mean, if it does not generally indicate one of the options rejected above; an answer to this question is required, moreover, if we are to understand just what it is that justifications are supposed to provide. Justifications support, confirm, corroborate, and so on; these descriptions of justification, however well-established, are unilluminating descriptions if we accept that justification does not lead to attributions of truth or high probability. For support, high corroboration and high confirmation are often associated in the literature, if not with truth, then with high probability.⁴⁶

I confess that I am not prepared to say how the sort of acceptance which follows from a successful justification should be described; indeed I am not sure that I think it can be described. One might think that, if this sort of acceptance is an indefinable primitive concept, then it becomes difficult to describe how it should be distinguished from acceptance which follows from other routes: for example, the acceptance which arises by way of faith, or by way of unconscious influences, or by way of stipulation.⁴⁷ However, we may treat acceptance as an indefinable term, a property of propositions of which the various ways we come to accept propositions are modes. Just as Kant tells us that space and time are forms of intuition, and the categories are forms of experience⁴⁸, perhaps justification, faith, stipulation and so on are the forms which acceptance can take for us. There may be other forms of acceptance, but they are not available to us; and so a full definition of what

⁴⁶ Those who follow Carnap, in any case, seem to make this association. Popper notoriously disassociates corroboration from high probability.

⁴⁷ I am thinking here of a conscious act of will, a directive issued by the mind to itself, to the effect that the proposition should be accepted. This sort of act may turn out to always be an act of faith, and should be subsumed under that category; but it is not clear to me at present that such a subsumption should be made.

⁴⁸ Insofar as the categories are ways of synthesizing the manifold of a possible experience, they are forms of that experience.

acceptance is cannot really be given. But the different forms it takes can be distinguished on the basis of origin, and so no difficulty really arises.

Still, though, one would prefer to have some way of linking the various forms of acceptance. If what acceptance is cannot itself be defined, then at least some account might be given to explain how these various ways of arriving at acceptance share some pattern; something which could at least explain why they should all be brought under the umbrella term 'acceptance'. If there was some process that all propositions which an individual entertains can go through, and which processing results in their acceptance or rejection, then we would further our understanding of these matters.

I cannot promise success, but I have a sketch that may accommodate all of what has been said so far. I shall try to articulate rather than argue for it here.

All propositions of which an individual becomes aware⁴⁹ have their origins as speculations. When they first occur to us, they are nothing but neutrally regarded ideas. This status can be wilfully maintained; this *may* be the case with some propositions, the content of which involves only fictional objects.⁵⁰ Most of our speculations, however, are immediately submitted to test – most of them, I conjecture, without our even being expressly aware of it. The tests I have in mind are those of logic and consistency with previously accepted opinion, although there may be others. If a speculation is to become a candidate for acceptance, it must pass the logical test of non-contradiction first. If it does so pass, then it must be tested against the individual's currently held opinions This is, in a way, just a recapitulation of the

⁴⁹ The phrasing here is suggestive of a sort of strongly realist attitude towards propositions. Although I think a distinction should be made between sentences and propositions, and that our reasoning concerns propositions rather than sentences or other strings of symbols, I am not committed to any kind of Fregean 'third realm' (cf. *The Thought*, G. Frege). At least, I do not consider myself to be so committed; if it turns out that I must be, then so be it.

⁵⁰ I have in mind propositions in poetic or other literary works.

logical test, but with a wider scope; instead of testing for any internal inconsistency of the speculation itself, the test is of the speculation's consistency with all of the speculations previously accepted by the individual. Note that this is not a coherence condition, if coherence is taken as more than mere consistency.⁵¹

A conjecture which has passed these tests is still a conjecture. It may have received some initial justification for accepting it – for example, its consistency with propositions already accepted. That does not mean that it is accepted, however. There may be competing conjectures which have also passed these initial tests; deciding among them may require further investigation, or it may involve appeal to some standard which one conjecture achieves but which the others do not. Let us suppose the conjectures in this case are 'All emeralds are green' and 'All emeralds are grue'. Both of these conjectures will pass the initial tests of logic and consistency with previously accepted opinion. They will not both pass empirical tests continuing beyond the year 2000; nor will they both pass internal standards we have for predicate-construction.⁵² The predicate construction test would presumably come well in advance of any empirical tests; if 'grue' is ruled an illegitimate predicate, then no empirical test need be carried out. The one that survives the internal tests will be automatically accepted without any empirical tests. If there is, however, only one candidate conjecture to begin with, then after it has passed all the internal tests it will as a matter of course be accepted; future events, however, may interfere with that acceptance.⁵³

⁵¹ See chapter II.

⁵² There is a lengthy discussion in the literature about the legitimacy or illegitimacy of Goodman's gerrymandered predicates. Arguments that the 'time-sensitivity' of the predicates makes them illegitimate fail, because we can construct gerrymandered predicates 'blue' and 'green' if we begin with 'grue' and 'bleen' as primitives. However, given that we do start with certain primitives (i.e. 'green' and 'blue'), we can at least individually rule out predicates like 'grue' as failing to meet one or more standards of predicate construction.

⁵³ If, for instance, one imagines some alternatives; but I anticipate.

Clearly an objection can be made here. Surely, it may be said, individuals as credulous as I am suggesting could not make very effective people. In particular, the remarks above could be castigated for failing to yield an account of why it is that people actively undertake empirical investigation at all. Individuals who behaved in the manner I describe could never be scientists; as soon as a palatable conjecture had passed all of their internal tests, its acceptance would be guaranteed. There would be no motivation for continued tests.

It is a good objection, and I quite agree; if people in general conformed to the simple model presented, much of their behaviour would be totally inexplicable. But I am not trying to give an account which will actually explain all of that behaviour; I want rather to sketch a view of cognition which makes those behaviours possible. Although it is not obvious, I think that this model does make our normal behaviour possible; I shall explain.

Take the case of the active investigator. Having come up with a conjecture – perhaps to solve some problem in physics⁵⁴ – the investigator is not satisfied, merely because the conjecture is consistent with all known physical theory; she wants to submit it to empirical test. And so she should; she has been *trained* to have that reflex. It is not one that comes entirely naturally – although those who are disposed towards greater self-doubt and suspicion will be more likely to think about conducting such active tests. For most ordinary reasoning, however, when our critical skills are not being taxed, we tend to nod our heads at every proposition which accords with what we already have. But what *kind* of acceptability is this? Is it the same sort of acceptability which the active investigator will finally bestow upon her conjecture when she has exhausted her stock of tests? I think it is precisely the same sort of

acceptability. That is, I think that a proposition accepted on the basis of internal tests only, where those tests were deemed to be the only ones required, has precisely the same status as one which is deemed to require empirical tests as well.

I think that status is best described as a sort of unavoidability. The reason that some speculation graduates from being mere speculation to being accepted is that it seems an unavoidable or inevitable consequence of a train of thought or critical assessment. Suppose I were examining peaches in order to determine whether or not all peaches have pits. Having examined every one and found it to have a pit, and assuming that I have examined all the peaches that there are, were, and ever will be, the conclusion that all peaches have peach pits seems inescapable. But it may be that some of these peaches do not have pits: some have had clever wooden mockups, indistinguishable in every way from real pits, carefully sealed within, whereas before they had been pitless. Clearly I was incorrect when I examined these particular peaches and judged them to have peach pits; that is not the point. Given the premises, my reasoning was impeccable, and the conclusion that all peaches have pits inescapable. Yet it was false: the correct generalization is that not all peaches have pits. Now if someone were to say that I ought to have made a more careful inspection of these pits before making my judgment, I can reply that it had simply not occurred to me that someone would insert false pits into peaches that had grown without them. I made no errors in reasoning; my only fault was a failure of *imagination*.⁵⁵

It is here that the solution to our difficulty lies. We can characterize accepted theories or ideas as conjectures which, after passing all tests deemed necessary, have

⁵⁴ Not to suggest that scientists are the *only* active investigators, although I think that there is a higher incidence of such investigators to be found in populations of scientists.

⁵⁵ It would be a serious mistake to think that, simply because one had not been able to come up with the alternatives, that one had therefore committed some sort of error in reasoning. Some discussion of this topic can be found in an article by Lawrence Sklar, "Do Unborn Hypotheses Have Rights?", *Pacific Philosophical Quarterly*, vol. 62, pp. 17-29.

ipso facto been judged inescapable. But their appearing inescapable has more to do with a failure to imagine the many extra dimensions that may apply to an examined situation. In examining the peaches, the dimension of deliberate deceit had not occurred to me. Who knows what other considerations may later occur to me, on the basis of which I shall have to alter my judgment?

This is not to say that simply by *imagining* that I *may* have been deceived by some consortium of evil demons with respect to peaches, I should now judge that in fact not all peaches have pits. Further investigation would be required. However, without making this (admittedly paranoid) imaginative leap, I would have thought myself possessed of all the relevant information and assumptions needed to make my judgment that all peaches have pits an unavoidable and thereby acceptable one. The speculations one makes with respect to a certain situation will determine the extent of one's investigations, and the sorts of conclusions one might reach.

Thus the status of accepted ideas, theories, or statements can be treated as unavoidable conclusions of investigation. Returning to our original predicament, what may we say of the impact of recalcitrant data on such accepted conjectures? For on this model, the data themselves are, after all, not speculated but accepted – accepted as data.

We can give precisely the same status to the acceptance of such data as we have given to accepted speculations, such as theories. All data are, after all, interpreted – they are theory-laden. That one has interpreted some accepted data in a certain way does not mean that they could not have been interpreted in other ways, ways which may not have occurred to one. But in advance of the sort of creative speculation needed to think of other interpretations such data might be given, the interpretation which they have in fact been given will seem unavoidable: 'Behold – a

peach pit.’ --- ‘Maybe it isn’t a peach pit.’ --- ‘What do you mean? *What else could it possibly be?*’ In this circumstance, it isn’t even possible for me to see the instance of a peach-pit experience as a being an instance of a *false peach-pit* experience, and therefore as a possible refutation of ‘All peaches have pits’.

One could object, however, that these data must pass the same internal tests as our conjectures, and that this will lead us back to the original difficulty we had with taking something to be true. For, if the datum being subjected to these internal tests happens to conflict with previously accepted opinion, should we not have reason to reject the datum outright, just as before? To accept this objection would be to miss the point of distinguishing acceptance, and taking to be true; for it ought to be a feature of acceptance that such conflicts do *not* result in the sort of wholesale rejections that attributions of truth can engender. Instead, such conflict invites further examination; until resolution is reached, no acceptance is possible. The previously accepted opinion, which had a particular consequence with which the datum conflicted, will have to be compared with the concatenation of theories that informed the description of the datum that caused the conflict. They will have to be compared, that is, for their consistency with the rest of the individual’s accepted opinions; and, if this cannot decide the matter, some other tests and investigations will be necessary.

Human imagination is generally quite limited and slow. We call those people ‘creative’ who manage on a regular basis to produce ways of thinking which none of the rest of us has made. One of the reasons, I think, that coming up with genuinely new ideas is so difficult is that our imaginations are constrained by previous speculation. Thinking in certain ways for a long time can exacerbate such constraint. This condition is well-known: its name is narrow-mindedness. Being narrow-minded does not make any difference in many contexts. A physicist whose sole task is to

record the collisions of protons and antiprotons in an accelerator does not need an active imagination – which is not to say that it would do her any harm. If the same scientist, while studying (rather than merely recording) such collisions, obtained an unexpected result, then she would have to perform some sort of act of imagination in order to accommodate the result. How successful she is will depend not only on her ability to imagine, full stop, but also on the *sorts* of solutions she is able to imagine. There may be a resolution to her problem that she is unable or extremely unlikely to consider, which some other individual might well be able to come up with. Moreover, this scientist's inability to imagine the resolution to her problem may stem from a too rigorous training in the theoretical apparatus upon which his field of study is founded; and this may be so, even though our scientist is no dogmatist – that is, even if he treats the theories he accepts as speculations only.

Some will find it unlikely that someone unschooled in a particular subject would be able to resolve problems arising within that subject. There is, however, some weight to the idea that a professional may find herself hampered by her own professional training when faced with problems arising within her field. In my own experience I have discovered that bright individuals, unfamiliar with a problem and the way of thinking that accompanies it, can yet grasp the difficulty and suggest surprising solutions.

Certainly circumstances do not generally permit scientists to take heed of the solutions offered by amateurs or the opinions of those whose concerns lie outside their own interests. The fact that it does not often happen, however, should not suggest that it ought not to happen – that only individuals working within the profession will be able to solve its problems. That would be a dogmatism as damaging as allegiance to some particular theoretical apparatus and ought to be avoided for that

reason. The successful propagation of healthy ideas – just as with healthy organisms – sometimes requires cross-fertilization; and, although sometimes the result is a Minotaur, at others we may expect a Helen.

In any case, the ramifications of the model I have sketched here should now be somewhat more apparent. The distinction between internal tests and empirical investigations breaks down; they are better viewed as points on a continuum. It may be that we are preprogrammed to perform the logic and internal consistency tests I mentioned above.⁵⁶ As fully grown adults, however, thoroughly trained and socialized, we generally have an appetite for more tests. We have that appetite, because socializing and training introduces us to more possibilities than our imaginations on their own could conceivably yield.⁵⁷ Once we have acquired language, the torrent of suggestions we receive from our normal interactions with others push the boundaries of what is possible in all directions. Inevitably, a feedback cycle occurs. Conjectures made at an early stage, and accepted at that stage because of their presumed unavoidability, will later turn out to have been only one of a number of possible solutions; and so what appeared to be the only possible solution is now seen to be surrounded by competitors. When such a situation develops, the mind has a tendency to cast around for selection procedures – and to withdraw earlier acceptance. Sometimes selection procedures can be developed from internal standards which were previously insufficiently refined; much of the time, however, empirical selection procedures will be invoked.

Other routes to acceptance, it seems to me, can also be accounted for within

⁵⁶ In most cases, I think, without consciousness of the procedure; it is a kind of ‘knowing how’ rather than a ‘knowing that’, as articulated by Ryle in *The Concept of Mind*.

⁵⁷ Perhaps I should say instead ‘yield in the short space of time allotted to individual humans’. That is, there is nothing which seems to me problematic in principle about the suggestion that a single human being, given enough time, could construct a body of knowledge as extensive as the collective efforts of mankind.

this sketch. For example, many people accept claims on faith. Individuals are usually exposed to religious ideas at an early stage in their development: it is generally agreed that their lack of experience and impressionability makes them more susceptible to such teaching. My account, I think, helps to explain why. Exposure to the notion of, say, a *god* implants in the mind of each individual a conjecture concerning that god. If this conjecture passes all of the currently available internal tests, whether naturally occurring or themselves implanted, then the conjecture will be accepted. Experience, however, with more varied internal or external (empirical) tests may lead an individual to entertain doubts, especially about an idea central to much of human life. Those who find that the implanted idea of a god fails one or more tests that they afterwards acquired and retroactively applied, could end up withdrawing their acceptance of the idea. For example, one might begin to wonder why a being as powerful as a deity should never manifest itself. Many will, of course, continue to accept the existence of a god in the face of such doubts. To accomplish this requires some ways of making the doubts and the accepted conjecture compatible. So, if the doubt about the existence of the god in question arose from the spontaneous thought that such a powerful being should be in evidence round about, an ingenious reconciliation of the doubts and the conjecture would be to say that the god *does* manifest itself in the form of acts of nature, or in some natural land formation. I have in mind such cases as the purported connexion between Thor and thunderstorms, or between the goddess Gaia and the earth. Alternately, the individual may have consulted a spiritual advisor concerning her doubts, and received a similar explanation of the apparent conflict; in that case, she need not have come up with the solution by way of her own imagination. This process may continue indefinitely; at each stage a solution is provided to whatever doubts arise by means of insulating the accepted idea

from those doubts: it is in some way immunized against that particular criticism or family of criticisms. In other words, faith is sustained by *ad hoc* adjustment; to have faith in a conjecture is to try to make it irrefutable relative to any tests devised or devisable.

This concludes the exposition of my sketch. It seems to me to accommodate acceptability and taking to be true; I think that, with some modifications, a functional account of probabilistic evaluation should be possible as well; but I haven't the skills necessary to attempt that project. What the model does not explain, of course, is *why* we make attributions of truth to some propositions; that is something which requires further investigation, and not necessarily philosophical investigation at that.

Conclusion

I have tried to show that justification of claims cannot issue in attributions of truth, nor can it issue in attributions of probability; that is, justification is not a sufficient condition for either of these attributions to acceptable claims. To begin with, I distinguished the skeptical arguments which are usually given in support of the idea that justified claims should not be taken to be true. These arguments have always seemed to me forceful, but (perhaps unsurprisingly) not conclusive. This is because they have only a normative value; they counsel against such attributions because of the risk of error. It is possible, however, to argue that because these arguments do not describe but only evaluate our practices, we can question the wisdom of taking them seriously. Since it seems to one, one might say, that what one accepts, one thereby takes to be true, it is unfair (or even incoherent) to instruct against so taking it. Rather than going after these objections to the skeptical arguments, however, I try to show why the idea that acceptance and taking to be true are necessarily conjoined is descriptively false of our practices. This is so, even if it *appears to one that acceptance and attributions of truth are so conjoined*; for how we behave with respect to what we accept does not reflect such purported attributions.

I distinguish, then, attribution of truth to claims from acceptance of those claims. After making clear in chapter two that the target of my investigation was the internalist conception of knowledge in its various forms, both foundational and coherentist, I argue that a distinction between acceptance and truth-attribution is motivated by an examination of our practice. Since we do change our opinions about things we accept, and which we are moreover justified in accepting, it seems that

attributions of truth are simply not appropriate characterizations of those accepted justified claims.

Just as acceptance is divorced from truth-attribution, so too is justification divorced. This is due entirely to the assumption that justification issues in acceptance, which assumption I do not argue for. If justification does issue in acceptance, and acceptance does not issue in attributions of truth, then justification does not issue in attributions of truth either.

I have tried also to distinguish my own notion of acceptance from the one other occurrence of a distinction between belief and acceptance I found in the literature, namely Van Fraassen's. Apart from noting that he restricts his notion to the scientific context, I found some differences between his motivations for introducing the distinction and my own, and tried to show how this difference in motivation leads to a difference in our understanding of how the distinction should be characterized.

The other candidate for the status of accepted claims examined is the attribution of probability to such claims. As a result of the various skeptical arguments adumbrated in chapter one, it can be seen that probabilism was an attractive alternative to truth-attribution, although probabilism was also attacked on skeptical grounds for the connexion it assumes between high probability and truth.

Once again, I adduce non-skeptical arguments to show why acceptance is not accurately described as generally involving an attribution of high probability. To begin with, if p is accepted because it is probable, then it seems reasonable to assume that in addition to accepting p , the assenting individual should also accept *that* p is probable. If this second-order proposition is itself accepted as *true*, then probabilism fails as a general account, and moreover the same inflexibility of opinion which we sought to avoid before resurfaces. But if 'p is probable' is itself taken to be probable,

this engenders a troubling regress. Whether the regress is vicious requires a discussion of probability valuation relative to justification rather than acceptance alone.

Probabilistic justification, I argue, can be understood in one of two ways. The procedure of justifying claims is performed either consciously or unconsciously. If consciously, then we face the following problems. The principles of the calculus used must themselves be accepted. If all acceptance is probabilistic, it only makes sense to say that those principles are accepted probabilistically by the stipulation of a high probability value for each one. This must be the case at the beginning, at least. For, in advance of acceptance of these principles, we cannot adduce their valuations by attention to evidence and therefore by application of the principles to themselves. Once initially accepted, such self-reflexive evaluation is possible; but the initial arbitrary assignment of high probability to those principles makes such further evaluations of dubious value.

Even if that conundrum could be solved, however, it would not be enough; because there is still one principle that the conscious probabilist would need, and which could not be accepted probabilistically. That principle is the one by means of which an evaluation of high probability issues in acceptance. If this principle is itself accepted, then it cannot be because of an (arbitrarily) high probability valuation made of it. For, without the principle in place it cannot be accepted on such grounds; so its acceptance would have to be of a different sort. Thus the idea that acceptance entails an attribution of high probability fails as a general account.

But we need not have an account which makes necessary the conscious acceptance of all these principles. An unconscious probabilism could allow that the operating principles of the system are not themselves up for consideration as acceptable; they can be treated as a framework to which conscious access cannot be

had. Now within this model there are two possible ways for the accepting individual to treat the claims he accepts. On the one hand, there is no consciousness of the probability valuation which attaches to claims; the individual is aware only of their accepting attitude towards the claim. If that were the case, a separate account would be needed of our propensity to say of many claims that we do accept, that they are probable and that we accept them because of that. Even if such an account were forthcoming, under this model we would not be able to take attributions of high probability as a general account of acceptance.

The alternative, that the unconscious probabilist *is* aware of the probability valuation attaching to claims he accepts, seems more plausible. However, although a very strong case can be made for this model, it ultimately fails. It fails, because it makes a connexion between high probability and acceptance which does not reflect our actual practice of acceptance. In two examples, I provide reason to reject such a connexion. Lottery players will readily assent to the extreme unlikeliness of the proposition that they hold the winning ticket; but their behaviour would become inexplicable if such an attribution resulted in their rejection of the claim that they hold a winning ticket. That example has nothing to do with empirical evidence, however, and so does not seem in all ways appropriate to the overall problem being addressed: the connexion between justification in general and attributions of high probability.

But examples involving empirical evidence instead of merely logico-mathematical reasons are not hard to find. If I look out of my window to see clouds boiling across the sky, a wind whipping the trees and the sun slowly fading, I think it highly probable that it will rain tomorrow. But I don't thereby accept that it will rain tomorrow.

I conclude that, just as it is inappropriate to make attributions of truth a

condition of acceptance, so it is also inappropriate to make attributions of probability a condition of acceptance. Moreover, just as the first claim made it impossible to treat of justification as issuing in attributions of truth, so too the second claim makes it impossible to treat of justification as issuing in attributions of probability.

This was, of course, not to say that people *never* treat the propositions they accept as true, or probable; they do all the time. It seems natural to think that an attribution of truth itself issues in acceptance; so taking something to be true is sufficient for accepting it. Having a successful justification is likewise sufficient for accepting it; but, because taking something to be true is not a necessary condition for accepting it, it can't be a necessary consequence of a justification either. While I sympathize with the temptation to say of claims one accepts, that they are true, I think the temptation should be resisted. On the one hand there are the skeptical arguments which make such attributions an epistemological sin; on the other, there are the arguments that I made here, which suggest that such attributions simply do not, for the most part, accord with our practices. Thus even in many cases where one might be willing to *say* 'I accept this, and thereby take it to be true', it is very unlikely that such a connexion is meant seriously. In order to determine whether it is, one would have to wait until a change of opinion occurred, where the opinion was once treated as true.

Some claims we accept, we accept on the basis of their high probability; but it cannot be the case that probabilism provides a general account of acceptance. Many claims we accept do not seem susceptible to being described as 'highly probable'; and many claims to which we are inclined to attribute high probability, we are unwilling to accept.

The position this leaves us in is an uncomfortable one; for apart from these options there seems to be no other word in our language which could capture the

significance of the word 'acceptance'. I tried to give a description of it in terms of *unavoidability*: that acceptance is a function of consistency with previously accepted claims and a function of the imagination's capacity to articulate the possible alternatives. By way of fleshing this out, I tried to articulate a sketch of the process of acceptance which could conceivably allow for, if not explicate, all the different sorts of ways in which we come to accept claims. The sketch is in no way complete, and I have no doubt that it is mistaken in what little detail I have provided. But I trust that the approach seems to have merit, and is worth exploring further.

This paper has closely examined the significance of accepted claims. But as the introduction made clear, the ultimate goal of this investigation is to construct a theory of knowledge. Given the way I have separated attributions of truth from acceptance, and the way in which I have associated attributions of truth with belief, I think I can provide a matrix for the ways in which a claim to *know* something can be construed. The matrix may be simplistic, and makes little room, at first glance, for probabilism; nonetheless I think it illuminating.

The traditional characterization of knowledge is as a justified true belief. It seems to me that knowledge, depending upon whose mouth is uttering the word, can now be construed in a few other ways as well. Knowledge could mean mere justified belief; the actual truth or falsity of the belief would be immaterial to its status as knowledge. Knowledge could also be equated with *justified true accepted claim*; or, even more weakly, *justified accepted claim*. What definition of knowledge suits, can be determined by what an individual is willing to say about knowledge claims. Skeptics, for instance, generally deny that they have any knowledge. I think that the best construal of knowledge that suits their attitude is justified true belief. Skeptics do not want to deny that people may have a hold of some truth, or right opinion, but they

do want to deny that people can be reasonably confident of their opinions, or that their opinions are somehow possessed of a *mark* of truth.⁵⁸ Since belief, as I have construed it in the last chapter, seems to be an intersection of acceptance and an attribution of truth, this amounts to marking some propositions with a truth attribution. Thus the sort of knowledge which skeptics want to deny that we can reasonably have are justified true claims which are accepted and taken to be true; or, in other words, justified true belief.

Others are persuaded that we do know things: those that take seriously our common practice of saying things like ‘I know that there is a glass in front of me’. Now some of these individuals simply deny the coherence of the skeptical arguments, and legitimate the practice in that way. One might, however, accept the skeptical arguments, but argue that we are fallibly justified in believing some claims, and that these claims count as knowledge whether or not they are actually true. One could therefore construe knowledge as a matter of having reason to take some claim to be true, whether or not it is actually true. I know of no individuals who take this position; at least, I know of no individuals who think that such an understanding of knowledge adequately conveys our intuitive understanding of that term, although I think a number would agree that in many cases where we speak of knowing something, we have only justified belief, and that the truth of that belief makes no difference to our attitude towards it.

Nor do I know of any individuals who endorse the position which suggests that knowledge should be construed as mere justified accepted claims. All that remains is the idea that knowledge consists in justified accepted claims which also

⁵⁸ This was the form of the debate between the ancient skeptics and the Stoic and Epicurean philosophers, who felt that true opinions carried a mark of their truthfulness with them. See Inwood and Gerson, *Hellenistic Philosophy: Introductory Readings*, especially pages 161-173.

happen to be true, even though we do not take them as such. As I said above, which position on the matrix one takes depends on a number of factors, including whether or not one thinks that we do have knowledge when we speak as if we do. My own skeptical sympathies incline me to take the traditional definition of knowledge as 'justified true belief' as the right one. But which specification actually gets attached to the word 'knowledge' strikes me as an uninteresting problem; you pays your money and you takes your choice. It does not follow that I am uninterested in seeking truth; I hope very much that I have got a hold of some here. Even if I have not, however, I feel that I have in this investigation deepened my ignorance considerably; and this, like the collection of truths, is also a road to wisdom.

Bibliography

- Ayer, Alfred, **The Problem of Knowledge**, Penguin Books, Harmondsworth, UK, 1974.
- Bacon, Francis, *Novum Organum*, appearing in **The Philosophical Works of Francis Bacon**, R.L. Ellis and J. Spedding eds., Routledge, London, 1905.
- Bartley, William Warren III, **The Retreat to Commitment** 2nd edition, Open Court Publishing Company, La Salle, 1984.
- Burnyeat, Myles, **The Skeptical Tradition**, University of California Press, 1980.
- Carnap, Rudolf, **Logical Foundations of Probability**, 2nd edition, University of Chicago Press, Chicago, 1963.
- Dancy, Jonathan, and Sosa, Ernest, **A Companion to Epistemology**, Blackwell Publishers Ltd., Oxford, 1996.
- Descartes, R., **The Philosophical Writings of Descartes**, J. Cottingham, R. Stoothoff, and D. Murdoch trans., Cambridge University Press, Cambridge, 1985.
- Frege, Gottlob, *Begriffsschrift*, appearing in **From Frege to Godel**, J. Van Heijenoort ed., Harvard University Press, Cambridge, Mass., 1967.
- , *The Thought*, appearing in **Philosophical Logic**, P.F. Strawson ed., Oxford, 1967.
- Goodman, Nelson, **Fact, Fiction, and Forecast**, 4th edition, Harvard University Press, Cambridge, Mass., USA, 1983.
- Honderich, Ted, **The Oxford Companion to Philosophy**, Oxford University Press, Oxford, 1996.
- Hume, David, **An Enquiry Concerning Human Understanding**, Hackett Publishing Company, Indianapolis, 1993.
- Hume, David, **A Treatise of Human Nature**, P. Nidditch ed., Oxford University Press, Oxford, 1978.
- Howson, C., and Urbach, P., **Scientific Reasoning: The Bayesian Approach**, 2nd ed., Open Court Publishing Company, Chicago, 1993.
- Inwood, Brad, and Gerson, L. P. **Hellenistic Philosophy: Introductory Readings**, Hackett Publishing Company, 1993.
- Kant, Immanuel, **Critique of Pure Reason**, Werner S. Pluhar trans., Hackett Publishing Company, Indianapolis, 1996.
- Lehrer, Keith, **Theory of Knowledge**, Routledge, London, 1990.
-

Papineau, David, *Methodology*, appearing in **Philosophy: A Guide Through the Subject**, Oxford University Press, 1995.

Plato, **Collected Dialogues**, E. Hamilton and H. Cairns eds., Pantheon, New York, 1961.

Popper, Karl, **The Logic of Scientific Discovery**, Routledge, London, 1995.

Sellars, Wilfrid, *Empiricism and the Philosophy of Mind*, appearing in **Science, Perception and Reality**, Humanities Press, New York, 1963.

Sklar, Lawrence, "Do Unborn Hypotheses Have Rights?" in **Pacific Philosophical Quarterly**, pp. 17-29, January 1981.

Van Fraassen, Bas, **The Scientific Image**, Clarendon Paperbacks, 1981.

-----, **Laws and Symmetry**, Clarendon Paperbacks, 1989.

-----, "To Save the Phenomena", in **Journal of Philosophy**, pp. 623-32, 1976.
