

# **The Epistemological Significance of Reflective Access**

**Charlotte Emily Hanson**

Ph.D. Thesis

Department of Philosophy, University of Stirling

May 2010

---

## Abstract

---

This thesis is, in part, a defence of a broad-based approach to epistemology. We should be wary of taking too narrow a focus and thus neglecting important aspects of knowledge. If we are too focused on one methodology then we are likely to miss insights that can come about from a different perspective. With this in mind, I investigate two particular methodologies in detail: Kornblith's naturalism and Craig's 'genealogical' approach. Kornblith emphasises the importance of looking at knowledge in the context of the natural world, thus stressing the continuity between animal and human knowledge. Craig, on the other hand, focuses on a distinctly human aspect of knowledge: the importance of enquiry and the sharing of information. As such, the two theories of knowledge that are developed have different emphases. I argue that by bringing them together we can better understand what knowledge is.

This leads us to the other main contribution of this thesis, which is a defence of the role of reflection in epistemology. This has often been neglected in contemporary epistemology, primarily because of the effectiveness of externalist theories of knowledge. The focus on externalism has led to reflection being sidelined. I do not argue that reflection is necessary for knowledge, but rather want to bring back attention to the important role that it plays in human life. Reflectively accessible justification is necessary for our knowledge claims and therefore plays a vital role in enquiry. If we add reflectively accessible justification to knowledge then it is both more stable and more valuable. Even if it is not necessary for knowledge, reflection should not be neglected.

---

## Acknowledgements

---

Firstly, I would like to thank my supervisor Alan Millar for all of his support and assistance over the last few years. He has helped me immensely through the, often fraught, process of completing this thesis. I would also like to thank Duncan Pritchard who was a great primary supervisor for the first year of my thesis and who continued to support me even after he had moved to the University of Edinburgh. Many thanks also go to the rest of the Department of Philosophy at the University of Stirling, who have provided a fantastic environment to work in and the financial support that has made this possible. Thanks also to all the friends who have supported me on the many occasions that I have needed it during the writing of this thesis. Finally, I would like to thank Gareth for his unwavering support and encouragement. This really could not have happened without him.

---

# Table of Contents

---

<b>Abstract</b> .....	<b>i</b>
<b>Acknowledgements</b> .....	<b>ii</b>
<b>Table of Contents</b> .....	<b>iii</b>
<b>Introduction</b> .....	<b>1</b>
<b>Chapter 1: How Should Epistemology be Done?</b> .....	<b>8</b>
1.1 Arguments For Conceptual Enquiry .....	9
1.1.1 The connection between concepts and the phenomena they describe.....	10
1.1.2 Some concepts are inherently complex .....	15
1.1.3 Conceptual enquiry as a necessary first step .....	18
1.2 Arguments Against Conceptual Enquiry .....	19
1.2.1 Concepts are parochial.....	19
1.2.2 Concepts embody error.....	28
1.2.3 The gap between concepts and phenomena.....	32
1.3 The Importance of Actual Practice.....	36
1.3 Alternatives to Reductive Analysis.....	39
1.4 Towards a Comprehensive Epistemology .....	44
1.5 Conclusion .....	46
<b>Chapter 2: Internalism and Externalism.....</b>	<b>47</b>
2.1 Some Definitions .....	48
2.2 Internalism .....	49
2.2.1 What constitutes ‘reflectively accessible justification’? .....	50

a) How should we understand ‘internal’? .....	50
b) What do we need to have access to? .....	52
2.2.2 Arguments for internalism about justification.....	52
a) Justification and responsibility.....	52
b) The ‘New Evil Genius’ view.....	53
c) Conee and Feldman’s cases.....	54
2.2.3 Arguments against internalist justification .....	55
2.2.4 The case for internalism about knowledge.....	58
2.2.5 Problems with internalism.....	61
a) Scepticism .....	61
b) Gettier Cases.....	62
c) Barn façade type cases .....	63
d) Animal and infant knowledge.....	64
2.3 Externalism.....	65
2.3.1 The case for externalism .....	65
a) An answer to scepticism .....	65
b) Gettier .....	67
c) Animal and infant knowledge .....	67
d) Chicken sexers .....	68
2.3.2 Types of externalism.....	69
a) Reliabilism.....	70
b) Modal epistemologies .....	72
2.3.3 A problem for externalism.....	75
2.4 Two Possible Compromises: Sosa and McDowell .....	79
2.4.1 Sosa’s virtue epistemology .....	80
2.4.2 Disjunctivism.....	84
2.5 Conclusions .....	85
<b>Chapter 3: A Naturalist Approach.....</b>	<b>86</b>
3.1 Knowledge as a Natural Kind.....	87
3.1.1 What is a natural kind? .....	87
3.1.2 Knowledge as a natural kind .....	89
3.1.3 Kornblith and the value of knowledge .....	94
3.1.4 Bermúdez’s criticism of Kornblith.....	96
3.2 Only One Kind of Knowledge.....	98

3.2.1 Knowledge and Social Practices .....	98
3.2.2 Knowledge and Reflection .....	105
a) Is reflection necessary for human knowledge?.....	105
b) Is non-reflective human knowledge of the same kind as that of non-reflective animals?.....	110
3.3 The Consequences of Kornblith’s Perspective .....	114
3.4 Conclusions .....	115
<b>Chapter 4: A Genealogical Approach.....</b>	<b>117</b>
4.1 Craig’s Methodology .....	118
4.2 Knowledge as ‘Tagging Good Informants’ .....	122
4.2.1 What is enquiry?.....	122
4.2.2 What is needed to be a good informant?.....	123
a) True belief.....	123
b) Identifiability: Condition X.....	127
c) As likely to be right as needed .....	130
4.2.4 The process of objectivisation .....	131
4.3 Building on Craig: The Role of Reflective Access in Enquiry .....	137
4.3.1 The Importance of Reflective Access for Informants.....	138
4.3.2 The Importance of Reflective Access for enquirers .....	141
4.4 Contrast with Kornblith’s Approach.....	144
4.5 Conclusion.....	146
<b>Chapter 5: Value Problems.....</b>	<b>147</b>
5.1 Value Problems .....	148
5.1.1 The Meno, or primary value problem.....	148
5.1.2 Swamping Problem.....	150
5.1.3 Secondary Value Problem.....	152
5.1.4 Tertiary Value Problem.....	153
5.2 Responses to the value problems.....	154
5.2.1 Virtue epistemology.....	154
5.2.2 Pritchard’s Anti-luck Virtue Epistemology .....	157
5.3 The Importance of Internal Justification .....	159
5.4 An Explanation of Metaepistemological Scepticism .....	168
5.5 Conclusions .....	170

<b>Conclusion .....</b>	<b>171</b>
<b>Appendix: Real-life vs. Philosophical Chicken Sexing .....</b>	<b>175</b>
A.1 A Brief History of Chicken-sexing in Philosophy .....	175
A.2 Real World Chicken sexers.....	177
A.3 The similarities and differences .....	179
<b>Bibliography .....</b>	<b>181</b>

---

## Introduction

---

For a very long time, knowledge was understood in an internalist fashion; that is, it was thought that reflectively accessible evidence or reasons were required for knowledge. In order to know a proposition  $p$ , you must be aware of some evidence for  $p$ . However, in contemporary epistemology the general move has been away from such theories to a more externalist account of knowledge; that is, one that does not require reflectively accessible evidence. An externalist theory seems preferable in that it can deal with problems such as scepticism and Gettier cases and also fits better with many of our intuitions about knowledge, such as the idea that non-reflective animals and infants can have knowledge. Because externalist theories of knowledge are preferable, there has often been a downplaying of the importance of having reflectively accessible justification.

This thesis has two main aims. The first is to suggest that in order to develop a complete theory of knowledge we need to take a very broad perspective, encompassing intuitions from both sides of the internalist/externalist debate. If one concentrates purely on the continuity between animal and human knowledge, for example, you will develop a very different theory from someone who concentrates on the way that knowledge is used to communicate information between agents. The only way in which we will come to fully understand knowledge is to take into account the various



## *Introduction*

different perspectives. The second aim of the thesis is to suggest a possible theory of knowledge that takes both internalist and externalist intuitions into account.

The thesis will begin by examining the role of conceptual enquiry in philosophy in general and in epistemology in particular. Traditionally, philosophy has proceeded by manner of conceptual enquiry. If we want to investigate knowledge, for example, our investigation would begin by looking at our intuitions about knowledge and testing them in various thought experiments. From this testing of our intuitions we could, in theory, develop a set of necessary and sufficient conditions that govern the application of the concept of knowledge.

Many contemporary epistemologists, however, find fault with this approach and chapter one examines this debate. Firstly we shall look at why, historically, conceptual enquiry has been the primary method of philosophical investigation. This seems to come down to perceiving an intimate connection between concepts and the phenomena that they describe. It also seems that the concepts that are of philosophical interest are of a particularly complex kind and conceptual enquiry is needed simply to disentangle what is meant by a particular concept.

We shall then examine several arguments against conceptual enquiry, none of which succeeds in proving that conceptual enquiry is not a valid and useful form of investigation in philosophy. There is, however, a good case to be made for conceptual enquiry not being the *only* form of investigation in philosophy. Useful insights into both concepts and the phenomena that they describe can be gained from non-traditional methodologies. A naturalist methodology may give us an insight into how knowledge functions in the natural world and help us to focus on the continuity between animal and human knowledge; a socially focussed methodology may give us insight into the social role that knowledge plays; and, a concentration on our actual practices of knowledge ascription may prevent us from moving too far away from

## *Introduction*

reality and into the world of the philosophical abstract. Ultimately, there are insights to be gained from these different perspectives and we should be wary of becoming blinkered to other approaches. The best theory of knowledge will be one that takes into account insights gained by all these different perspectives (and whatever other relevant perspectives there may be).

Having cleared the ground, so to speak, with regard to methodology, we move on in Chapter Two to cover some background with regard to theory; the debate between internalists and externalists. This debate can be divided into two separate questions, one about justification and the other about knowledge. Internalism about justification claims that justification must be reflectively accessible and internalism about knowledge claims that internalist justification is necessary for knowledge. Externalism in both cases is the denial of internalism. There is a strong case for understanding justification in an internalist way: the intuition that two subjects who are identical with regard to beliefs are also identical with regard to justification; a perceived connection between justification and responsibility; and various thought experiments seem to point towards this conclusion. However, this does not mean that knowledge must also be internalist in nature.

In fact, the case against internalism about knowledge is very strong. Gettier cases, sceptical arguments, chicken sexers, and the fact that we happily ascribe knowledge to infants and animals all seem to point towards externalism. Internalism struggles to deal with these cases but externalist theories can give answers and explanations for all of them. It is not entirely plain sailing for externalism, however. Barry Stroud suggests that an externalist answer to scepticism can never be fully satisfactory. Chapter Two introduces Stroud's problem but there will be further discussion in Chapter Five. Another problem is that there will be something lacking in an externalist account of knowledge that does not take internalist intuitions into account or at least explain how

## *Introduction*

we come to have such intuitions. The chapter concludes by discussing two such theories, the virtue epistemology of Ernest Sosa and the disjunctivism of John McDowell.

In Chapter Three we move on to a detailed investigation of one particular methodology: Hilary Kornblith's naturalism. Kornblith argues that knowledge is a natural kind and as such should be investigated empirically like any other natural kind. He argues that knowledge plays an important causal and explanatory role in cognitive ethology – the science of animal cognition – and that it can only play such a role if it is a natural kind. He then argues that human and animal knowledge are of the same kind, meaning that the externalist understanding of knowledge developed from cognitive ethology must also be the correct understanding of human knowledge.

There are problems with both stages of Kornblith's argument, but even if they can be resolved, there is a larger problem. His focus on animal knowledge underplays the importance that reflection plays in human knowledge. By using animal knowledge as his starting point it skews the focus of his investigation onto only that which is possible for non-reflective creatures to have. However, humans have reflective capabilities and this cannot fail to have an effect on knowledge. It may be the case that our knowledge is, at heart, identical to that of non-reflective creatures, but that cannot be the whole story. For a full understanding of knowledge, we need to take into account how we, as human beings, use and interact with knowledge.

Chapter Four examines an approach that puts the social importance of knowledge at the heart of its investigation. Edward Craig argues that the origins of the concept of knowledge lie in a need for us, as information-sharing beings, to be able to tag reliable informants. Human beings are not always reliable and sometimes make false claims, both intentionally and unintentionally. Thus, we need to have a way of telling which informants are the good ones. A good informant as to whether or not *p* is one that is,

## *Introduction*

accessible to the enquirer, recognisable by the enquirer, as likely to be right as the enquirer needs and someone with whom the enquirer can communicate. This protoconcept of knowledge has undergone a process of objectification away from the needs of a specific enquirer on a particular occasion to someone who is potential of use to someone in some situation, meaning that to be a knower is to be potentially a good informant rather than an actual good informant.

Craig's focus on the social development of knowledge also skews the way that he looks at knowledge. He neglects the continuity between human and animal knowledge because he is focussed on the human context. Also, because he is looking at the issue from the perspective of the enquirer, he places little importance on reflection. However, given the importance that is placed on being able to give grounds for your belief, it seems that reflectively accessible reasons must take a central role in the social development of knowledge. If we are to believe the Gricean claim that it is inappropriate to claim knowledge unless one has sufficient grounds, then it seems that reflectively accessible evidence is of great importance when it comes to sharing information. Taking this into account, we can extend Craig's view such that, to be a good informant one must also have good reflectively accessible grounds for one's belief.

That is not to say that we should ignore the insights gained from Kornblith's investigation. It is important to bear in mind both the evolutionary history that has made us capable of having knowledge and the cultural evolution that has brought about the concept of knowledge as it is used today. We should acknowledge that there is a continuity with non-human animals, a shared core phenomenon that is properly called knowledge. In addition, because humans are both reflective and social, the role of reflectively accessible justification cannot be underplayed. Chapter Four concludes with sketching a suggestion for incorporating both of these insights. Knowledge at

## *Introduction*

heart must be externalist, something that we share with non-reflective creatures. However, we also need to make space for reflection. Rather than arguing for a different *kind* of knowledge, it is suggested that there is knowledge simpliciter and knowledge-plus; that is, knowledge plus justification.

Chapter Five goes on to discuss the implications of this theory, in particular with regard to the various value problems that are much discussed in contemporary epistemology. The *Meno* problem asks how knowledge can be more valuable than mere true belief if they have the same practical value. The answer that I give is that, in fact, knowledge is not more valuable than mere true belief but that knowledge plus justification is more valuable. Justification adds value beyond the value of having mere true belief or even knowledge. Of course, an explanation is needed as to why we believe that knowledge is more valuable than mere true belief. My suggestion is that in most cases, our knowledge is accompanied by internalist justification. This, combined with the fact that we require reflectively accessible grounds for knowledge *claims* means that in almost any situation in which we would claim to know, we would have justification for our claim. This explains why internalism had a hold in epistemology for so long and why we have the intuition that knowledge is more valuable than mere true belief. Blind knowledge, however, is not more valuable than true belief.

A second advantage of this theory is that it can explain why the externalist answer to scepticism is unsatisfactory, despite being adequate. We cannot have a knowledge-plus answer to scepticism as there is no way that one can have reflectively accessible evidence that one is not a brain-in-a-vat. Thus, one can only know that one is not a brain-in-a-vat but only in an brute externalist fashion. Thus, although we know, we know in a way which is not valuable.

Thus, taking into account the different perspectives and methodologies that are suggested in the literature can lead to a more rounded theory of knowledge which is

## *Introduction*

able to deal with some of the problems faced by more narrowly focused theories. I do not pretend to have covered every possible perspective, nor every perspective that can be found in the literature. However, we can take this theory as a starting point and further add to our understanding by taking in to account further insights gained elsewhere. The danger of remaining too narrowly focused is that we will miss out on some important insight that would lead to greater understanding of knowledge.

---

## CHAPTER ONE

### How Should Epistemology be Done?

---

This chapter aims to do two things: firstly, to defend the validity of conceptual enquiry as a method of enquiry in epistemology; secondly, to suggest that conceptual enquiry alone is too narrow a methodology. Traditionally epistemology, like all philosophical investigation, has been conducted through conceptual enquiry; that is, we examine our concepts in order to understand better the phenomena that they describe. This enquiry has standardly taken the form of reductive analysis – a search for necessary and sufficient conditions for the application of our concepts. For example, enquiry into the concept of knowledge has usually resulted in the claim that knowledge is true belief plus another condition, such as reliability or justification. However, in recent times there has been a move against conceptual enquiry in general and reductive analysis in particular, with alternative methodologies, such as experimental philosophy, being advocated instead (e.g. Nichols *et al* (2003), Weinberg *et al* 2008).

This chapter begins by looking at the reasons why, historically, conceptual enquiry was thought to be an appropriate form of enquiry for philosophy. The main motivation appears to be the perception of an intimate connection between concepts and their associated phenomena. It also seems that the concepts that are of interest to philosophers are ones that are inherently complex in a way that other concepts are not.

We shall then examine the arguments against conceptual enquiry: the claim that we should directly investigate phenomena, not just concepts; the claim that our concepts are inherently complex; and the claim that our concepts are in some sense parochial. None of these criticisms succeeds in showing that conceptual enquiry is not a legitimate methodology for epistemology, but they do provide interesting suggestions for alternative methodologies. Reductive analysis, however, is more problematic than conceptual analysis in general, and we will examine one suggested alternative: the ‘practical explication’ of Edward Craig (1990). A practical explication does not require that we rely on our intuitions alone; if our intuitions are inconsistent then we can reject them in favour of a more consistent account. In addition, it does not ask for necessary and sufficient conditions that define all cases of knowledge, but rather conditions that define a prototype of knowledge. This will allow us to accommodate more perspectives on knowledge, making our theory more comprehensive.

## **1.1 Arguments For Conceptual Enquiry**

Although conceptual enquiry was for a long time the standard form of investigation in philosophy it may not be immediately obvious why this should have been the case. One might think, as Hilary Kornblith (2002) suggests, that our time would be better spent investigating the phenomena that we are interested in directly. In this section I shall look at the reasons that ‘traditional’ philosophers have given for pursuing such an investigation using A. J. Ayer (1957/1973, 1973) and Peter Strawson (1992) as two examples. Firstly, I shall examine the idea that there is an intimate connection between concepts and the things that they describe. To be an instance of knowledge, for example, is to meet the conditions to be classed as knowledge. Secondly, the concepts that interest us as philosophers are particularly complex and as such they require



conceptual investigation. As philosophers we are not interested in all concepts, and complexity may be one distinguishing feature of philosophically interesting concepts. Because of this complexity it seems that even in pursuing non-conceptual enquiry we will need to begin with conceptual enquiry.

### **1.1.1 The connection between concepts and the phenomena they describe**

According to Hilary Kornblith:

Not so long ago, philosophy was widely understood to consist in an investigation of our concepts.... The idea that philosophy consists in, or at a minimum, must begin with an understanding and investigation of our concepts is, I believe, both natural and very attractive. It is also, I believe, deeply mistaken. On my view...the subject matter of epistemology is knowledge itself not our concept of knowledge. (Kornblith 2002:1)

Traditional epistemology at the very least begins with conceptual enquiry.. It seems that to some extent this will be true of any philosophical enquiry whether or not concepts are the main focus of investigation. If nothing else, enquiring into the concept has the effect of clearing the ground. We need to identify what exactly it is that we are attempting to enquire about. However, this is likely to be true of all investigations. The role that conceptual investigation has played in philosophy is obviously much more central. However, this is not, as Kornblith seems to suggest, because philosophers were preoccupied with the concepts themselves. Instead, concepts and the phenomena that they describe were seen as intimately and inherently connected; by examining the concepts it was thought that we could find out about the phenomena.

Ayer suggests that ‘already knowing the use of certain expressions [philosophers] are seeking to give an analysis of their meaning’ (Ayer 1957/1976: 8). Although this claim explicitly concerns the meaning of linguistic expressions there is a parallel claim

about concepts: linguistic expressions express concepts, and the concepts expressed are fixed by the meanings of linguistic expressions. Analysis of meanings, therefore, goes hand in hand with the analysis of concepts. Philosophers are not merely investigating the meaning of words but rather are trying to uncover 'the work that the word...does' (Ayer 1957/1976: 28). One may talk of investigating meanings or of investigating concepts but the aims of such investigations would be similar. In understanding the meaning of a word we should be uncovering what it would be for something to fall under that concept.

According to traditional philosophy, investigating the concept should tell us something about the phenomenon that it describes:

It is therefore indifferent whether... we represent ourselves as dealing with the words or dealing with facts. For our enquiry into the use of words can equally be regarded as an enquiry into the nature of the facts which they describe. (Ayer 1957: 29)

Ayer also claims that examining concepts:

...throws light not only on the workings of our language but also on the character of the world it serves to describe. There is in any case no sharp distinction between investigating the structure of our language and investigating the structure of the world, since the very notion of there being a world of such and such a character only makes sense within the framework of some system of concepts which language embodies...[O]ur experience is articulated in language, and the world which we envisage existing at times when we do not is still a world which is structured by our method of describing it... we cannot detach ourselves from every point of view...The idea that we could prise the world off our concepts is incoherent; for with what conception of the world should we then be left? (Ayer 1973: 49)

What Ayer seems to be saying in these passages is that by investigating our concepts we will better understand the facts that those concepts describe. We can only understand the world in terms of our concepts and we cannot move completely away

from them. As concept-using beings we can only view the world through the prism of our concepts. There may be facts that exist independently of our conceptual structures but in our thinking and talking about those facts we must deploy our concepts. Conceptual enquiry will therefore give us greater insight into how we understand the world, which in turn should give us insight into how the world works. Presumably our concepts have come about as a reaction to the world around us and should tell us at least something about how the world works. Our concept of a cat has presumably come about through the observations of and interactions with cats that people have had throughout history and so we would expect that the concept can tell us something about what cats are.

One natural way of understanding the connection between concepts and phenomena is that a concept determines what falls under that concept. Certainly this is very simply the case for some concepts. To be a bachelor just is to be an unmarried man. In cases such as these it is clear that conceptual investigation will uncover something about the phenomenon that the concept is describing.

Natural kinds are usually thought to be determined by the world and not by our concepts. Water is the substance with chemical structure  $H_2O$ , irrespective of whether any concept of ours represents this fact. That is not to say that we cannot find out interesting things about water simply by examining our concept of it. Examining our concept of water will not tell us what the underlying structure of the substance is but we can discover that it is usually liquid under normal, room-temperature conditions, that it is thirst-quenching and so on

Strawson (1992) suggests a similar justification of conceptual analysis to Ayer's. In investigating the concept we can find out about the actual phenomenon that we are investigating. In fact, he goes as far to say that it is preferable to investigate via conceptual analysis rather than trying to investigate the phenomenon directly:

By talking about our conceptual structure, the structure of our thought about the world, rather than, as it were, directly about the world, we keep a firmer grasp on our philosophical procedure, a clearer understanding of what we are about. (Strawson 1992: 33)

He goes on to suggest that there are certain concepts or types of concepts, such as our concepts of ‘material object’ and ‘states of consciousness’, that are so pervasive in our thinking and talking about the world that it is:

...quite inconceivable that these concepts should have this universal employment unless we took it for granted that there were, or existed, in the world things to which those concepts, or concepts of those concept-types applied. So the question ‘What are our most general concepts, or types of concept, of things?’ and the question: ‘What are the most general types of things we take there to *be*, or *exist*?’ really come to the same thing. (Strawson 1992: 33)

What Strawson is saying is that our most general concepts and issues about ontological categories go hand in hand. This is not to say that we can assume from the fact that we have such concepts that the things that those concepts identify must exist, but rather that if we want to investigate the general types of thing that exist the best starting point that we have is to look at the concepts of the things that we take to exist.

The very fact that we have such concepts should carry some theoretical weight, claims Strawson:

...it is a fact that we must give great weight to, that the pervasiveness and generality of certain concepts or concept-types carries ontological implications in the undisputed sense; that is, implications about what we ordinarily and quite generally *take* to exist...[It] would surely be reasonable to get a clear grasp of how they do function before trying to evaluate the reasons which some philosophers might have given for challenging our general accepted working ontology. (Strawson 1992: 34)

He goes on to say that philosophers who put forward such challenges must surely start from somewhere, such as ‘some abstract considerations about existence and identity, or unity and plurality’ (Strawson 1992: 34), and so cannot avoid calling on those concepts. We will be best equipped to engage with, and assess, such challenges if we have a good understanding of the concepts involved. Presumably, if opponents stray too far from our normal usage of such concepts then we would have good reason to challenge their suggestions.

However, according to Ayer, philosophy should not be restricted to describing the way in which our concepts operate. Our investigations can also be normative in nature:

Philosophers do not limit themselves to uncovering the criteria which we actually use in assessing different types of statement. They also question these criteria; they may even go as far as to deny their validity. (Ayer 1957: 30)

In questioning a set of criteria we may produce paradoxes, such as the Cartesian sceptical argument, but these are not just philosophical perversions. Instead, discussion of what lies behind these problems and paradoxes can be informative about the concept under discussion.

However, even if philosophy goes beyond simple conceptual enquiry, it is still a vitally important starting point to our investigations:

If someone wishes to convince us that he has a better way of describing the world, he has to make it intelligible to us, and this means that he has to relate it to concepts that we already have... This is not to say that philosophy is restricted to the practice of conceptual analysis, but it is only there that it can profitably begin. (Ayer 1973: 43)

Although there might be deficiencies in our conceptual scheme we must start with the concepts that we have, and we cannot move too far from them. If someone suggested an analysis of knowledge that was completely removed from our ordinary conception

of it, we would be right to question whether we were actually talking about the same thing. Any radical departure from a commonly held concept would have to be justified and it would have to be recognisable that the concept applies to what we were investigating in the first place, unless of course the revision was linked to an error theory about the latter.

### **1.1.2 Some concepts are inherently complex**

In fact there is further reason that conceptual enquiry is important. We do not feel the need to analyse all concepts philosophically; we have no interest in philosophically analysing the concepts of ‘apple’ or ‘taxi driver’, for example. The kind of concepts that are traditionally subject to philosophical enquiry, such as ‘knowledge’, ‘justice’, ‘truth’ and so on, seem to be inherently complicated. Although we are very competent users of these concepts, it is not immediately obvious what unites the cases that fall under them. There is no easy way of explaining why all cases of knowledge count as knowledge or why all just acts can be gathered together under the one concept. This is not the case with, for example, ‘apple’. We can see that all apples have a similar shape, taste, and smell and come from the same kind of tree. We may not know what it is that underlies these similarities but we can identify features that make each individual similar to others of the same kind.

Secondly, in exploring such concepts we often become more confused rather than less. It seems that as soon as we commence a philosophical investigation into such a concept we are led into confusion and often outright contradiction. Some concepts seem to have in-built conflict; for example, there seems to be an inherent tension in the concept of free-will that is revealed with very little reflection: neither a determinist nor and indeterminist account appears to be compatible with it. The problem seems to be

built into the concept from the very start. That is not to say that we cannot solve such problems but rather than there are inherent tensions in certain concepts, which is exactly what makes them suitable candidates for conceptual enquiry.

A similar complexity or tension also exists in the concept of knowledge and emerges as the problem of scepticism. That the concept of knowledge is inherently complex in this way goes a long way to explaining the preoccupation with scepticism that is found in traditional epistemology. The sceptical problem can be understood as a paradox: from premises that seem intuitively correct we are led to an unacceptable conclusion. A brain-in-a-vat (BIV) is a disembodied brain floating in nutritional solution and connected by electrodes to a very powerful computer. Via these electrodes the BIV is 'fed' experiences. The experiences are indistinguishable from the experiences that the BIV would be having if it was actually experiencing the real world. If the experiences are subjectively indistinguishable then it seems clear that I cannot know that I am not a BIV, because if I were I would be having ostensibly the same experience as I am now. This, however, seems to undermine my knowledge of the 'external' world:

- 1) If I know that I have hands then I know that I am not a BIV.
- 2) I cannot know that I am not a BIV.
- So, 3) I cannot know that I have hands.

Although the two premises appear to be intuitively plausible, together they lead to the unpalatable conclusion that I cannot know that I have two hands, or any other proposition about the external world for that matter.

Some philosophers (e.g. Schiffer 1996, Spicer 2008, Weatherston 2003) have gone as far as to suggest that the concept of knowledge is inherently flawed. Spicer, for example, argues that the sceptical puzzle demonstrates that we have 'inconsistent

intuitions of extension' (Spicer 2008: 51). He then argues that the best explanation for this inconsistency is an inconsistency in our folk epistemology, which in turn means that there are in fact no conceptual truths about knowledge. I shall discuss this in more detail in section 1.2.2, below but what is clear is that if there is an inherent flaw in the concept of knowledge then this will only be uncovered by conceptual investigation of some kind. In order to discover whether a concept is flawed we must first engage with that concept.

If it is the phenomenon of knowledge that we are interested in, then why should it be a problem, or even be of interest, that there is a tension in our concept? Surely this is only of interest if we are preoccupied with concepts. This is not the case for two reasons. Firstly, if there is tension inherent in a concept then it is going to be problematic to even identify the phenomenon that we are discussing. Take the case of scepticism, for example. Beliefs that seem to be paradigmatic examples of knowledge, such as the belief that I have two hands, seem to be undermined by the sceptical argument in such a way that they also seem intuitively *not* to be examples of knowledge. Unless we have some good, independent, reason to dismiss the sceptical hypothesis then our investigation will struggle from the start. Secondly, as discussed above, our concepts describe the world from a human perspective. If there is a tension in a concept then this may very well come down to some complexity inherent in the phenomenon under investigation. That is not to say that complexity in our concepts necessarily means complexity in the phenomena that they describe but rather that it is a possible indicator of such complexity.



### **1.1.3 Conceptual enquiry as a necessary first step**

If philosophically interesting concepts are inherently complex in this way then it seems that it may be necessary in any investigation of such phenomena to engage in some level of conceptual enquiry. Even Kornblith, a staunch opponent of conceptual enquiry, admits that appeals to intuitions are a reasonable starting point for a naturalist investigation of knowledge (Kornblith 2002: 12-14). Whilst Kornblith does not talk of conceptual enquiry, he does claim that, in the first instance, we need to rely on our intuitions about knowledge. Kornblith draws an analogy between the job of an epistemologist and that of the rock collector. At the beginning of the investigation we need to rely on our intuitive judgements in order to collect rocks of the same kind, or examples of knowledge. Our intuitive judgements should pick out the most obvious cases. Once we have collected enough clear examples then we can move away from intuitive judgements and into empirical investigation about the underlying structure of the thing we are interested in.

If, as suggested above, the concept of knowledge is inherently complex then the cases that Kornblith suggests will not be analogous. Even cases that seem to obviously count as knowledge throw up sceptical problems, which raises the question of whether they are cases of knowledge after all. It is not enough to simply collect obvious cases of knowledge. Further conceptual enquiry needs to be done to clarify which cases count as obvious cases of knowledge. If there is tension in our concept then it may not be as simple to collect cases of knowledge as it is to collect specimens of a particular type of rock.

## **1.2 Arguments Against Conceptual Enquiry**

Here we shall examine three of the objections against conceptual enquiry: the claim that concepts are parochial, that they embody error and that there is a gap between our concepts and the phenomena that they describe. In each case, a problem is suggested and then it is claimed that the conclusion we should draw is that we need a change in philosophical methodology. However, although these problems suggest that conceptual analysis is not perfect they do not show that it is not a valid form of philosophical enquiry.

### **1.2.1 Concepts are parochial**

One challenge to conceptual investigation comes from adherents of experimental philosophy, who claim that there may be systematic differences between the intuitions of people from different cultures and socioeconomic groups. The suggestion is that if there really are such systematic differences along class or cultural lines then in relying on the intuitions that we – that is, Western philosophers – have we may unintentionally be doing ‘ethno-epistemology’ (Nichols *et al* 2003), that is, an investigation into what a certain restricted group of people mean by the term ‘knowledge’. There is no reason for us to assume that all people have exactly the same intuitions, or that philosophers occupy some privileged position in this regard. If intuitions about scepticism, for example, are not shared by all, or at least the majority of, people then it is hard to see why it should be considered so worrisome.

Experimentalism can be divided into two kinds, characterised by the following two theses:

**The negative thesis:** Armchair reflection and informal dialogue are *not* reliable sources of evidence for (philosophically relevant) claims about folk concepts.

**The positive thesis:** Survey studies *are* a reliable source of evidence for (philosophically relevant) claims about folk concepts. (Kauppinen 2007: 97).

Here I shall concentrate on the first thesis as it is this that directly attacks conceptual analysis. I see no reason why, in principle, surveys can not be a valid source of evidence in philosophy provided that they are carried out well. Potentially, even if concepts are shared by the majority, there may be insights that philosophers have missed or downplayed that could be brought to our attention by such surveys. They may also reveal a different emphasis or focus amongst the general public compared with the ‘experts’. There are also many cases in which intuition appears to be divided amongst philosophers, such as in the debate between internalists and externalists about knowledge, and so to get data on people’s pre-theoretic opinions might be able to shed some light on such debates.

Traditional philosophy can be divided into four projects: the normative, the descriptive, the evaluative and the ameliorative (Weinberg *et al* 2001/2008: 18-19). Here I shall concentrate on what is described as the descriptive project. This project aims to analyze either epistemic concepts or epistemic language. In particular it aims to understand what ‘our’ concepts are and how ‘we’ use epistemic terms. The concepts are, claim Weinberg *et al*, ‘characterised rather vaguely by using the first person plural’ (2001/2008: 18) The implication of this statement is, of course, that in characterizing the project in these terms, something inappropriate is being done. Weinberg *et al* see the traditional epistemic project as being guided by what they call ‘Intuition Driven Romanticism’ or IDR (2001/2008: 19-20). This method uses intuitions as the basis for coming to a reflective equilibrium about the concepts under discussion; an intuition is ‘a spontaneous judgement about the truth or falsity of a proposition’ (Nichols *et al*

2003: 2<sup>1</sup>). In epistemology, for example, Gettier cases are used to test our intuitions against a particular theory of knowledge. The assumption underlying IDR is that ‘our’ intuitions are everyone’s intuitions and this is an empirical claim that needs to be tested.

Weinberg *et al* conducted a set of studies in which they set about testing (1) the intuitions of people from different cultural and socioeconomic backgrounds and (2) the intuitions of students who had been exposed to more or less philosophy instruction. In their studies they presented participants with various different scenarios based on important epistemological thought experiments. For each case they asked whether the subject in the scenario had knowledge or a true belief regarding a particular proposition. These judgements were recorded as the participants’ intuitions. For example, participants were asked to judge a ‘Truetemp’ case:

One day Charles is suddenly knocked out by a falling rock, and his brain becomes re-wired so that he is always absolutely right whenever he estimates the temperature where he is. Charles is completely unaware that his brain has been altered in this way. A few weeks later, this brain re-wiring leads him to believe that it is 71 degrees in his room. Apart from his estimation, he has no other reasons to think that it is 71 degrees. In fact, it is at that time 71 degrees in his room. Does Charles really know that it was 71 degrees in the room, or does he only believe it? (Weinberg *et al* 2001/2008:26)

Weinberg *et al* discovered that there seem to be systematic differences between people of different ethnicities and socioeconomic statuses. In particular, the reactions to Gettier cases seemed to be very different depending on cultural background. A majority of subjects in the East Asian group claimed that the subject did know in a standard Gettier case whereas in the Western group the majority thought that the subject didn’t know, which tallies with the standard intuition as presented in the

---

<sup>1</sup> The two papers, Nichols *et al* (2003) and Weinberg *et al* (2001/2008), are written by the same three authors, Jonathan Weinberg, Shaun Nichols and Stephen Stich.

literature. Similar results were found when comparing Westerners to subjects with a background in the Indian Sub-Continent.

There were also differences with regard to socioeconomic status. Weinberg *et al* (Weinberg *et al* 2001/2008: 33), following usual practice in social psychology, divided subjects into two groups depending on whether they had a college education. Those that had not attended college were labelled low-SES and those that had were labelled high-SES<sup>2</sup>. The subjects were presented with a case based on Dretske's 'cleverly disguised mule' case (Dretske 1970: 1015-6). Pat is visiting the zoo and points out a zebra to his son. If the zebra had been a cleverly disguised mule then Pat would still have thought that it was a zebra.<sup>3</sup> However, the animal that Pat is looking at is a real zebra. As with the Truetemp case, the subjects were asked whether Pat 'really knows' or 'only believes' that the animal is a zebra. The results seem to show that people in a lower socioeconomic group are more likely to say that a subject knows than people from a higher socioeconomic group. 33% of the people in the low-SES group thought that Pat knows, compared to 12% in the high-SES group. Although the majority in both cases thinks that Pat 'only believes' the proportions appear to be strikingly, and significantly, different.

The experimenters also looked at whether studying philosophy has an effect on one's intuitions (Nichols *et al* 2003: 241-242). They took two groups, one of which consisted of students who had taken two or less philosophy courses; the 'low-

---

<sup>2</sup> Although they are following usual practice in the social sciences, different labelling may have been useful here. The relevant factor distinguishing these two groups seems to be college education, not socioeconomic status. If, as they attempt to show, the study of philosophy changes people's intuitions then perhaps study *in general* will also have an effect. Undoubtedly, college education and SES generally go hand-in-hand, but perhaps it would be more informative in this case to label the groups according to their level of education rather than their SES.

<sup>3</sup> Interestingly, there have been actual cases of painted or dyed donkeys or mules being used in place of zebras. For example, in October 2009 the BBC reported that a zoo in Gaza was using donkeys dyed to look like zebras after their real zebras died ([http://news.bbc.co.uk/1/hi/world/middle\\_east/8297812.stm](http://news.bbc.co.uk/1/hi/world/middle_east/8297812.stm), accessed 24/05/10). Also 'Tijuana Zebras' are a popular tourist attraction in Tijuana, Mexico. The story goes that white donkeys were originally painted with black stripes in order to show up better in black and white photographs.

philosophy' group. The other group was made up of those who had take three or more philosophy courses; the 'high-philosophy' group. When presented with a BIV case, it was found that while the vast majority of the high-philosophy group thought that the BIV case rules out knowledge. In the low-philosophy group, on the other hand, there was a slight majority in favour of the conclusion that we can know in spite of BIV cases. Their suggested interpretation of this is that either philosophers are unintentionally teaching the standard intuition or that, again unintentionally, we are weeding out students with conflicting intuitions (Nichols *et al* 2003: 232). Interestingly, Nichols *et al* note that the BIV case is the only one in which they saw a significant difference between the intuitions of the low- and high-philosophy groups (Nichols *et al* 2003 242).

In addition to the finding that intuitions about certain matters differ according to group the study also found that there seem to be some shared intuitions. The vast majority in every group thought that you could not come to know something on the basis of a 'special feeling' that it was going to happen when this special feeling was only as accurate as guessing (Weinberg *et al* 2001/2008: 36). If this is the case then we may not share all of our intuitions about knowledge but there may be some core intuitions about knowledge that are universally shared.

As the authors themselves claim, these results are preliminary. The experiments had relatively small samples and have not yet been replicated. Many questions have been raised over the quality and effectiveness of the trials and I will not go into all of them here. Rather, I think there are two general hypothetical questions to be asked. Firstly, if it were shown that there are significantly different intuitions between different ethnic or class groups then what are the implications for epistemology? The suggestion of Weinberg *et al* is that this difference in intuitions undermines the project of using intuitions to make universal claims about knowledge. We cannot guarantee that our

intuitions are the only intuitions (or the only correct intuitions) about knowledge as it seems that intuitions are affected by culture in ways that we were unaware of.

The second question is of whether the kinds of experiment suggested can actually get at the intuitions we are interested in. Certainly, we should be wary of assuming that our intuitions are the only intuitions around, but that is not to say that we need to abandon conceptual enquiry altogether. Rather, it would seem to say that we need to do further investigation into our concepts. We need to check if our concepts are the same as the concepts of other groups of people. There seems to be no good reason to think that this could not be done through traditional methods of conceptual enquiry but with a larger set of people. The problem, it seems, is not so much that conceptual enquiry is suspect but that philosophers, as a group, are not sufficiently diverse. It is not that the method is at fault, but that we need to expand the range of people who take part in philosophical debate.

Kauppinen claims that the type of experiments that have been carried out only test the subject's 'surface' intuitions and not our 'robust' intuitions, which are what we should be interested in, philosophically speaking. Thus any apparent disagreement does not necessarily mean that there are actually cross-cultural differences. The best way to find out these robust intuitions, he claims, is by the traditional Socratic method of dialogue between the experimenter and subject. Presumably, in principle, we could uncover differences in robust intuitions – this should not be surprising as we encounter clashes of intuitions amongst philosophers relatively frequently – but we cannot discover these via the methods that the experimental philosophers suggest. If Kauppinen is correct, then the apparent differences that the experimental philosophers have suggested may not actually mean that intuitions vary cross-culturally. It also shows that there is a case for conceptual enquiry of a more traditional sort; discovering our robust intuitions will go hand in hand with enquiring into our concepts.

Kauppinen (2007: 97) defines ‘robust’ intuitions as intuitions

represented by those responses of non-specialists that are *stable* under arbitrary increases in consideration of relevantly similar circumstance, ideality of circumstances, and understanding of the workings of language. (Kauppinen 2007:110)

These are the intuitions that should be the focus of philosophical enquiry but we can only guarantee getting to these if the following conditions are met:

1. The subject is a competent user of the concept in question.
2. They have considered the case in sufficiently ideal conditions.
3. Their answer was influenced only by semantic considerations. (Kauppinen 2007: 101)

If these conditions are met then we will have access to the intuitions that are of most philosophical interest. They will not be met, however, in the type of experiment that Weinberg *et al* advocate.

To be a competent user of a concept is to generally apply the concept in line with the norms of your linguistic community. This demonstrates that there is a normative dimension to meaning; it is possible to make a mistake about the content of the concept. If this is the case then we cannot simply deduce the rules guiding a subject’s use of a concept from the cases in which she applies the concept. Nor can we be sure that someone who is minimally competent will be able to deal with ordinary cases but not with the far off possible worlds that are often in play in philosophical thought experiments.

Ideal conditions are those which are ‘conducive to avoiding performance errors’ (Kauppinen 2007:103). The robust intuitions that we are interested in are not ‘gut reactions, but simply...pre-theoretic judgements that may require careful consideration’ (Kauppinen 2007: 104). The more unusual or remote a case is, the more difficult it is to



bring forth an answer that truly reflects the concept. It may not be at all simple to work out what one thinks about a particular case and we may need to distinguish between different considerations that weigh on the case and even perhaps consider different cases to get things clear. If our judgement in one case does not fit with our other uses of that same concept, then it would not count as a robust intuition.

The third condition that Kauppinen suggests is that the subject should only be influenced by semantic concerns. However, semantic considerations are not the only ones that influence proper usage of a concept. There are also pragmatic considerations. For example, it would be misleading to say that yesterday you had lunch voluntarily when there is no special reason for saying so. It is true, and semantically appropriate, but conversationally misleading as it carries the implicature that there is something notable about the fact that you dined of your own volition. If we are to identify people's *semantic* intuitions then we need to somehow isolate them from pragmatic concerns.

Kauppinen argues that these conditions have not been met in the surveys completed by experimental philosophers so far. As such the subjects of those experiments may

only appear to understand the question, who may have an imperfect grasp of the concept in question, who may or may not think hard about the application of the concept in circumstances that may or may not be conducive to avoiding conceptual mistakes, who may or may not rush in their judgements, and who may or may not be influenced by various pragmatic factors. (Kauppinen 2007: 105)

Thus, the surveys may very well not be getting at the philosophically interesting robust intuitions but rather at the participants' gut reactions, and so apparent evidence that intuitions about knowledge vary along social and cultural lines may not be as problematic as it first seems.

It is also not clear how, in the context of a survey of the type the experimentalists want, one could check that the conditions were ideal, and that subjects were giving properly considered answers, without conducting dialogue with the subjects. Nor is it clear how we could determine that a subject was focussing on purely semantic intuitions. The best way would presumably be to conduct a dialogue with the subject, asking why each answer was chosen and encouraging further reflection. However, this would not be the kind of hands-off observational research that the experimentalists want. Rather, it would be ‘a return to the good old Socratic method’ (Kauppinen: 2006:106). Through dialogue we can make sure that distracting conditions and distorting factors have been removed in order to reveal the subjects’ robust intuitions.

If, after conducting this kind of research it turned out that Americans and South-East Asians had different *robust* intuitions then it would be correct to conclude that there are two different concepts in play but this has not yet been tested. It should be no surprise that those subjects who had had some philosophical training gave different answers to those without. Those who had some experience of philosophy were likely to be better at sorting through the distorting factors that affect our gut reactions and more used to considering the kind of far-off possibilities that philosophers tend to talk about. This does not necessarily mean that doing philosophy changes one’s intuitions but rather that it helps one become better at uncovering one’s, possibly well-hidden, robust intuitions.

However, even if it were shown to be the case that robust intuitions did not vary according to social factors, any differences in surface intuitions are still interesting. Why would surface intuitions differ in an apparently systematic way if our robust intuitions are in fact the same? This in itself would appear to be an interesting area for investigation. That is to say, there may be space for the kind of investigation that the experimentalists suggest even if it does not do the job that they want it to. A full

investigation of knowledge would include investigations of our surface intuitions as well as our robust ones.

Even if there are some differences in our intuitions about knowledge there is good reason to think that different cultures would develop similar concepts of knowledge. Edward Craig suggests analysing the concept of knowledge from a state-of-nature perspective. If we do this then we may come up with reasons why a concept at least very similar to our concept of knowledge would be likely to arise in any society. Craig's suggestion is that the concept of knowledge arose in response to the need to tag good informants. Presumably this need would exist in any group of beings capable of sharing information and capable of making mistakes. Even if we do not accept Craig's approach or his theory of knowledge we can see that the concept of knowledge plays a significant and important role in human interaction. This, in itself, may be a reason to think that different societies would develop at least a similar concept of knowledge.

It seems, therefore, that the case put forward by experimental philosophers is not as strong as it first appears. There is good reason to think that dialogue is a better way to discover what people's intuitions are than simply asking questions in a survey. The moral is that we should encourage as many people as possible from as many backgrounds as possible to join in the practice of philosophy. If we also vary our methods of investigation to include experiment as well as conceptual enquiry we may develop a more complete understanding of all the aspects of knowledge.

### **1.2.2 Concepts embody error**

Finn Spicer argues that our folk epistemology is inconsistent, meaning that there can be no conceptual truths about knowledge. If this is the case then conceptual enquiry will not be an effective way of finding out about knowledge. Spicer first outlines a view

of what concepts are that, he claims, underlies many common theories of conceptual truth; inferential descriptivism (ID) (Spicer 2008: 44). ID claims that the identity of concepts is determined by their connection with other concepts, which are embodied in our dispositions with regard to those concepts. ID consists of three claims:

- **Universality Claim:** there is a set of inferences that all possessors of the concept knowledge are disposed to draw (let us call these core inferences);
- **Identity Claim:** the identity of the concept is fixed by its core inferences;
- **Semantic Claim:** the semantics of the concept knowledge is fixed by its core inferences. (Spicer 2008:44-5)

To claim that something is a conceptual truth about knowledge is just to claim that it is one of the core inferences with regard to knowledge. The set of all of our core inferences about knowledge form a folk epistemology.

However, because this folk epistemology is a tacit theory, it will not be straightforward to get at. Spicer (2008: 48-9) suggests that the best method would be to use an experimental method like the one suggested by Weinberg *et al.* We need to uncover people's intuitions, but in a scientific manner. Once we have discovered what intuitions are held we can then deduce what the best explanation for each intuition is. These explanations taken together constitute folk epistemology.

However, Spicer (2008: 50-2) is doubtful of the possibility of developing such a coherent theory. He thinks that there is evidence that our folk epistemology is inconsistent. When faced with a sceptical argument it seems that our intuitions are pulled in two contradictory directions. We have the intuition that we do know everyday propositions, but when we are faced with the BIV possibility we have the intuition that we do not know such propositions. It seems that we have inconsistent intuitions about the extension of knowledge. The best explanation of this tension is that there is an

inconsistency in our folk epistemology. The sceptical argument, therefore, should be taken as *prima facie* evidence that our folk epistemology is inconsistent.

It should be possible to test this hypothesis empirically. Spicer (2008: 53) suggests that if we do have an inconsistent folk epistemology then we would also have a mechanism such that we would not generally be drawn into outright contradiction. Thus, it is likely that we would have a mechanism such that we only think about a few folk epistemological principles at once and our intuitions about knowledge would exhibit framing effects. Spicer suggests that the literature on contextualism suggests that this may be the case, but he would rather see some empirical data. To this end he has begun experiments to test this hypothesis but has not yet published any results.

If folk epistemology is indeed inconsistent then there can be no conceptual truths about knowledge. According to the semantic claim, what knowledge is fixed by folk epistemology. If folk epistemology is inconsistent, then we have the choice to either deny the semantic claim or deny that there is any such thing as knowledge. One way to give up the semantic claim but to still hold on to something like ID would be to claim that some of the principles of folk epistemology are in fact defeasible. Spicer describes this position as ‘dual-status descriptivism’ (Spicer 2008:56). This view distinguishes between logical and empirical principles and of the two only the logical principles determine the reference of knowledge. These are the only conceptual truths about knowledge. Empirical principles are not conceptual truths but rather useful rules of thumb for ascribing knowledge.

The problem for this theory, however, is that there is no obvious way of determining which principles are logical and which empirical. If we cannot distinguish between the two kinds of principle, then, although we can save conceptual truth, we can no longer consider it *a priori*. That is, we will no longer be able to infer from the fact that something is part of our concept of knowledge that it is therefore a

conceptual truth about knowledge. Thus, doing conceptual investigation will not be able to reveal the nature of knowledge as it will uncover both conceptual truths and rules of thumb with no way of distinguishing between them.

Spicer thinks that there is good *prima facie* evidence that our folk epistemology is inconsistent. However, it looks as if there are alternative ways of explaining the same data. One way of explaining the intuitions is just to say that the intuitions are not about when one knows but about when it is right to claim that one knows. It may be that, although one does know that one is not a BIV, it is conversationally inappropriate to claim that one knows. When we make a claim to knowledge it carries with it the implicature that we have grounds to back up that claim. It could be that we confuse intuitions about when we have knowledge with intuitions about when it is appropriate to claim that we have knowledge (For further discussion of this see Pritchard 2005).

An alternative explanation of the apparently contradictory intuitions of extension is the claim that there are in fact two senses of knowledge involved in the sceptical argument; for example, Sosa's animal and reflective knowledge (Sosa 1997; 2007). In one sense a subject knows as long as her belief is formed in the right way. In another sense of 'knows', however, she only knows a proposition provided that she has formed it responsibly, that is, if it is warranted by her reflectively accessible evidence. If this were the case, then, when assessing a sceptical scenario, it could be that these two senses of knowledge are getting confused. The subject cannot know that she is not a BIV in the sense of having a responsibly held true belief because there is no reflectively accessible evidence that she could have that would show this to be the case; it is an integral part of the sceptical scenario that the subject cannot distinguish between the good and the bad case. That does not mean, however, that she cannot have a grounded belief, and therefore know, that she is not a BIV. If this belief was in fact formed in the right kind of way then she can know this.

There are, therefore, ways of explaining the apparent contradictions. That is not to say that it would not be useful to do empirical research into whether certain intuitions are universally shared or not, but that even if we discovered that there was an apparent conflict in our intuitions, that would not mean that our folk epistemology is inconsistent. If we can give an explanation for why there is an apparent inconsistency then we will be able to answer Spicer's criticism. Thus, he would need to show that it was not only an apparent conflict but an actual one. And, it seems that the only way that we will be able to determine whether it is a real conflict or not is by more conceptual investigation. Thus, Spicer's argument does not show us that conceptual investigation is illegitimate.

### **1.2.3 The gap between concepts and phenomena**

A third objection to conceptual investigation is the claim that there is a gap between our concepts and the phenomena that they describe (See, for example, Kornblith 2002). As such, if we are interested in what *knowledge* is then we should investigate it directly rather than looking at our concept of knowledge. Hilary Kornblith argues that knowledge is a natural kind and, therefore, investigating the concept alone will not get us very far.

In his "Twin Earth" thought experiment Putnam (1975b: 223-227) argues that the meaning of natural kind terms is not determined solely by the ideas that we associate with them. He imagines that there is a far-off planet, 'Twin Earth', that is almost exactly like Earth to the extent that some of the people speak a language called 'English'. However, there is one respect in which Twin Earth differs from Earth and that is that the substance that speakers of Twin Earth English refer to as 'water' is not H<sub>2</sub>O but a substance whose chemical structure we shall abbreviate to 'XYZ'. XYZ and

water are indistinguishable from each other at normal temperatures and pressures and have the same thirst-quenching and lake-filling properties as water. If scientists from Earth visit Twin Earth and examine the water then they will discover the difference in structure and report that 'On Twin Earth the word 'water' means 'XYZ'' (Putnam 1975b:223). The word 'water' means different things on Earth and Twin Earth.

If we think back to a time before the discovery of the chemical structure of water then it would still be the case that there were two senses of 'water'. Even though Oscar<sub>1</sub>, on Earth, and his Doppelganger on Twin Earth, Oscar<sub>2</sub>, would be in the same mental state when thinking about water, the extensions of the term 'water' would be different for each of them. On Earth 'water' refers to H<sub>2</sub>O and on Twin Earth it refers to XYZ, even if the people on those planets do not know the relevant chemical information. The reason for this, Putnam argues, is that when I point to a glass of liquid and describe it as water then I am presuming that the liquid I am identifying is the same as the stuff that I and the other people in my community describe as water. If that empirical presupposition is fulfilled then we have a necessary and sufficient condition for what it is to be water, it is to be the same as the stuff that I am pointing to. We may not know what the underlying structure is that makes water what it is but nonetheless that structure in some sense defines what "water" means.

If we agree with Putnam that meanings 'ain't in the head' (1975: 227) then analysing the meaning of 'knowledge' or the associated concept will only tell us about our concept and not about the phenomenon that falls under it. If Kornblith is right that knowledge is a natural kind term then the same is true of 'knowledge'. We will not be able to find out what knowledge is by examining only our concept of knowledge. Instead we should engage in an empirical investigation as we would for any other natural kind.



A natural kind ‘defines...a well-behaved category, a category that features prominently in causal explanations, and thus in successful inductive predictions.’ (Kornblith 2002: 62). As such, anything that features in causal explanation in the natural sciences should be considered a natural kind. Knowledge, claims Kornblith, plays an important causal and explanatory role in cognitive ethology – the study of animal cognition – and as such should be considered a natural kind. There is an underlying structure to knowledge – reliably formed, true belief – that explains the role that knowledge plays in animal survival: knowledge makes animals better equipped to survive. Kornblith claims that it is common in the cognitive ethology literature to use the term ‘knowledge’ to describe animal behaviour and cites various examples such as Herman and Morrell-Samuels work on dolphins which talks about ‘knowledge acquisition’, ‘knowledge acquiring abilities’ and ‘knowledge requisites’ (Herman and Morrell-Samuels 1990/1996:289, cited in Kornblith 2002: 56).

Kornblith’s argument appears to turn on his claim that knowledge is a natural kind. However, it is not clear that he needs to go that far. For the central claim that knowledge should not be investigated by conceptual analysis, all he needs to claim is that the most important truths about knowledge are not conceptual truths. If it were the case that every truth specifying the nature of knowledge were a conceptual truth then it would be appropriate to investigate it by conceptual analysis. It is, however, inappropriate to investigate something about which there are conceptual and non-conceptual truths in a way that only took account of the conceptual ones. If this is the case for knowledge then a purely conceptual enquiry would be insufficient to find out what knowledge is; we should also do some empirical investigation.

It is not clear, however, how Kornblith’s method actually differs from traditional epistemology, except that it advocates a different evidence base. In traditional epistemology we also collect various cases that clearly are and are not knowledge and

try to derive what knowledge is from those cases. Gettier cases and fake barn cases are supposed to be clear examples of where we lack knowledge and chicken-sexers are suggested by externalists as a clear case of where someone does have knowledge. The main difference between Kornblith's methodology and traditional epistemology seems to be that he advocates taking our knowledge samples from cognitive ethology whereas traditional epistemology takes them from thought experiments. Kornblith appears to still be interested in the concept of knowledge; it is merely that he is interested in the concept held by cognitive ethologists, not philosophers.

In taking this approach Kornblith seems to move too far away from the traditional data of epistemology. Whilst cognitive ethology offers a wealth of animal-based examples, Kornblith offers no examples of human knowledge behaviour. This is because he argues that there no important distinction to be made between human knowledge and the knowledge of non-human animals. If this is the case then cases of animal knowledge will suffice. However, having a broader base for our theory would surely be better, so it would be preferable to pick a theory that is based on evidence from both human and non-human behaviour, giving traditional philosophy the advantage over Kornblith's method as he sets it out. I shall discuss this in more detail in Chapter 3 but it is important to note here that it seems that we would need to do some conceptual investigation to settle this question. If it turns out that we have two different concepts of knowledge, for example, then Kornblith's claim that knowledge is whatever cognitive ethology says it is does not hold much weight. If we are at all interested in human knowledge, and Kornblith only discusses non-human knowledge, then it looks as if he has changed the subject. Conceptual enquiry, therefore, will still have some role to play even in Kornblith's methodology.

It seems, therefore, that none of the criticisms of conceptual enquiry that we have discussed here is successful in showing that conceptual enquiry should be abandoned. If there is a close connection between our concepts and the things that fall under those concepts then we should expect to find out some things about our objects of enquiry just by examining our concepts. If the concepts are complex or even apparently incoherent then that may tell us that the situation in the world is a complex one. An inherent tension within a concept may give us reason to think that there is more than one thing being grouped under that concept, for example. Conceptual enquiry may not be the whole of philosophical investigation but it is certainly an important part.

### **1.3 The Importance of Actual Practice**

Mark Kaplan (2000. See also Kaplan 1991) suggests a problem of a different kind. In some cases traditional philosophical methodology appears to elicit intuitions that are at odds with our actual practice. For example, one assumption that appears to underlie the Cartesian sceptical argument is that in order for S to know that p, S needs to rule out *any* hypothesis that she is aware of that is incompatible with p (Kaplan 2000: 279-80; c.f. Stroud 1984: 29-30). However, this is completely at odds with our everyday epistemic practice. In everyday cases, we do not expect people to be able to rule out every possible case, but only relevant cases (Kaplan 2000; see also Austin 1946/1979). We do not expect a witness in a court case, for example, to rule out the possibility that she was dreaming, unless there is a reason to think that this is likely.

Barry Stroud (1984: 56-64) agrees with this description of ordinary practice but denies this means that the requirement that we rule out incompatible hypotheses is not a condition on knowledge. Rather, he claims that what is required for knowledge-*claims* is different to what is required for knowledge itself. Whilst it may be completely

appropriate for me to claim that I know that p, this doesn't mean that I actually do know p. There are circumstances such that I can legitimately claim to know without actually knowing. If I have incredibly good reason to think that I know that p then it may be appropriate to claim that I know that p even if it later becomes apparent that I did not actually know. What we can be said to know for practical purposes is different from what we know strictly speaking.

Stroud distinguishes between the real-life context where one needs to make a decision based on one's belief, and the context of philosophical investigation where we can sit back and examine the situation at leisure (Stroud 1984: 64-66). The fact that the witness cannot rule out the possibility that she was dreaming makes no difference for practical purposes. However, when we sit down and reflect upon the case we can see that it is not *actually* true that she has knowledge. In the context of a philosophical investigation we are not interested in what is good enough to do the job of knowledge but what knowledge actually is. Stroud believes that the conditions for knowledge that are developed in the context of philosophical reflection are the ones that we should take note of, irrespective of whether they match how we actually act or not. The conditions, he claims, remain the same whether in the philosophical context or not. However, practical concerns prevent us from applying them in real life situations. Although the conditions are constant, whatever the situation, there are pragmatic and social pressures that cause us to adopt a lower standard in most cases.

Stroud gives an example to illustrate the way in which our investigation of knowledge can be shielded from normal practice in this way: the plane-spotter case (Stroud 1984: 67-8). The plane-spotters have a book that tells them the following about enemy planes:

1. If a plane has features x, y and w it is an E.
2. If a plane has features x, y and z it is an F. (Stroud 1984: 67)

In the first scenario the careful plane spotter (CPS) spots a plane that has features x, y, and z, checks the book and concludes that the plane is an F. The irresponsible plane spotter (IPS) only sees that it has features x and y but doesn't check whether it has z, but still claims to know that the plane is an F. In the second case the scenario is the same except that there is a third type of plane G, not included in the book because it is rare or not dangerous, that has features x, y and z, the same as F-type planes.

There are three judgements that Stroud suggests are correct.

- i. In the first case the CPS knows the plane is an F but the IPS doesn't.
- ii. In the second case neither plane spotter knows but the CPS is blameless for his lack of knowledge.
- iii. If the CPS was told of the existence of Gs he would withdraw his claim to knowledge. (Stroud 1984: 69-70)

When you read the dreaming argument, you are like the CPS in (iii), you become aware of something that undermines your knowledge. Although you thought that you knew that p, and are blameless for that, you do not, in fact, know that p. You can still claim to know for practical purposes but it is not really the case that you know.

Kaplan (2000: 290-2), however, claims that the reactions we have are much better explained by our everyday practices of knowledge attribution than Stroud's explanation. It is only when a relevant hypothesis – the existence of Gs – is suggested that we no longer think that the CPS has knowledge. Only in the second case, where we have special reason for thinking that there is a relevant hypothesis do we think the CPS does not know. The CPS may say, in the second case, that he knows only for practical purposes but actually he doesn't even know by our ordinary standards.

Kaplan's argument is that if we do not judge our epistemological theories by our ordinary practice then it is mysterious as to what we are judging them against. It is not

clear why our philosophical intuitions should take precedence over how we actually use the concept. That is not to say that there is not some use for conceptual enquiry and examining our intuitions about knowledge, but if these intuitions pull in a different direction to ordinary practice then our practice should be the default that we appeal to. If our intuitions are at odds with this then there is something odd about our intuitions, not about our ordinary practice. It is important that, in our philosophical theorising, we do not lose sight of the way in which our concepts are used in normal life.

### **1.3 Alternatives to Reductive Analysis**

There seems to be no reason to reject conceptual enquiry as a form of philosophical investigation. There are, however, good reasons to think that reductive analysis is not the best way forward. For a start, it is not clear that all concepts can be analysed into neat sets of necessary and sufficient conditions. In particular, if the concepts that we are investigating in philosophy are inherently complicated, as argued above, then it seems that they are likely not to be analysable in this way. If we add to this the fact that despite the work of many hundreds of epistemologists for many years we have still not come up with a completely satisfactory analysis of knowledge – at least not one that is widely accepted – then it seems that we may be heading on the wrong track. So, if we are unlikely to achieve a satisfactory analysis of knowledge, what are the alternatives? In this section we shall examine what reductive analysis is and how one might go about conceptual investigation without analysis, focusing in particular on Edward Craig's idea of 'practical explication' (Craig 1990:8) of concepts.

The 'traditional' method, which is often rejected in contemporary philosophy, is reductive conceptual analysis. This consists of the search for necessary and sufficient conditions for the correct application of a concept. As Ayer suggests:

[The theory of knowledge] aims...at establishing criteria for knowledge; criteria which may possibly set limits to what can be known...[it] is primarily an exercise in scepticism; the advancement and the attempted rebuttal of arguments that are intended to prove that we do not know what we think we know. (Ayer 1973: 1)

Even if we know how to use the verb 'know', that does not prevent us from asking the philosophical question, 'What is knowledge?' Ayer suggests that this question can only be answered by asking further questions, such as, 'Do cases of knowledge have something in common?' and, 'Is knowing a person or an object the same thing as knowing that something is the case?' and so on. Such questions should then be 'put...to the test of particular instances' (Ayer 1957/1976: 27).

For Ayer, methodology is very important. Indeed, he writes:

It is by its methods rather than its subject-matter that philosophy is to be distinguished from other arts or sciences... Philosophical theories are not tested by observation.' (Ayer 1957/1976: 7)

It is not that philosophy looks at a particular range of topics distinct from other disciplines but rather that it looks at those topics in a different way. We could, presumably, conduct a scientific investigation into what knowledge is but that would not answer the philosophical question. When conducting scientific investigation, we investigate the objects that fall under a concept, it would not tell us about the concept or how to apply it. The traditional method by which concepts are investigated philosophically is to ask general questions about a concept and then see how particular examples suggest that we should answer them. For example, in asking what knowledge is, we may come up with the suggestion that it is true belief. However, one can then suggest a counterexample of a belief that is just accidentally true, leading us to conclude that knowledge should not be accidental, and so on.

That is not to say, however, that philosophical investigation is completely inward looking. Although the investigation is not conducted using experimentation, it is still world-directed. As Jessie Prinz says:

It is a bit misleading to describe [philosophical investigation] as nonempirical... The philosopher examines her own mental states, monitors patterns of behaviour and pays close attention to how language is used. (Prinz 2004: 27)

Nor is philosophy a solitary pursuit; discussion of ideas is vital to making progress. For a start, discussion with others may bring up errors in reasoning or other mistakes that have been made. However, this is not the only role that discussion plays. If I am interested in the concept of knowledge then I am not just interested in what knowledge means in my idiolect. I am interested in a shared concept that plays an important role in my interactions with others. Although we are interested in intuitions, unique intuitions are not the focus of interest. We are interested in concepts that are shared, which enable us to communicate and which pick out important phenomena in the world. We cannot discover what these concepts are simply by sitting alone and reflecting on them. This may be a first step but it is only through discussion with others and observation of the world that we can fully examine our concepts.

We have still not answered the question of what exactly traditional conceptual analysis is, however. Strawson suggests two analogies for conceptual analysis. The first analogy is the view of analysis held by Wittgenstein and the logical positivists: analysis as therapy (Strawson 1992: 3-4. See also Wittgenstein 1953). When we try to think philosophically we are likely to get ourselves tied up in knots and to be led to conclusions that are unacceptable, if not paradoxical. The philosopher's job is to help us out of these confusions and set us back on the right course. We should also be trying to simplify things, to reduce our concepts to more basic ideas. Wittgenstein also



suggests why we get into such a mess. In thinking philosophically we lose sight of how the concepts actually work. When the concepts are actually in use such problems do not arise: I usually have no problem identifying when knowledge is present or absent, but as soon as I start thinking philosophically I lose sight of that. The cure is to bring ourselves back to the way in which these concepts are used in normal, everyday language (Wittgenstein 1954: §116)

The second analogy is that doing conceptual analysis is analogous to investigating the grammar of a language (Strawson 1992: 4-5). Instead of the grammar of a particular language we are investigating the logical grammar of our concepts. One can be a fluent native speaker of a language but still struggle to express the grammatical rules that underlie the structure of the language. It is similar with concepts: we can be very competent users of the concept 'knowledge', but that doesn't mean that we will be able to express what it is for something to be a case of knowledge.

The grammar analogy implies that there is a hidden structure to our concepts that we can uncover through analysis. The therapy analogy, on the other hand, implies that the aim of analysis is to undo confusions that we get ourselves into when thinking philosophically. There is some truth in both of these accounts. Philosophical thinking does lead to confusion about certain concepts which we need to get out of but this is not to do with the manner in which the concepts are investigated. Instead, it is that the concepts themselves are inherently complex.

Can there be useful conceptual investigation that does not aim for a reductive analysis? Strawson suggests that instead of reduction we should aim for 'elucidation' (Strawson 1992: 19). Instead of aiming to reduce a concept into simpler ones we should try to understand how it connects other concepts, aiming for a web of understanding rather than a series of reductive analyses. We shouldn't worry about

circularity provided that we move in a ‘wide, revealing and illuminating circle.’ (Strawson 1992: 20-21)

This gives us a suggestion for how we may proceed should we find that reductive analysis is untenable. There are ways in which we can come to a greater understanding of our concepts, and of the phenomena that they describe, even if we cannot give a reductive analysis. We can gain clarity without having to reduce our concepts to basic concepts. In some ways the task that Strawson suggests is harder than the traditional project. In order to fully elucidate what knowledge is, I may also need to investigate truth, belief, perception and a whole host of other concepts. However, this would give a greater understanding not just of what knowledge is but of the place that it holds and the role that it plays in relation to our other concepts.

A more detailed description of how we might carry out a non-reductive investigation into our concepts is suggested by Edward Craig. We shall discuss Craig’s methodology in detail in Chapter Four, but in short Craig suggests that we should attempt a ‘practical explication’ (Craig 1990: 8) of knowledge. Instead of investigating the concept directly, Craig looks to the possible origins of the concept; why would we have developed such a concept, what use is it to us? In doing so, Craig hopes to shed light on the concept that we have *now* through looking at why and how it may have originated.

We should not expect our definition of knowledge to fit with all cases of knowledge. Rather, if we achieve anything like an analysis of knowledge, we should expect it to define the prototypical cases only. Craig focuses on the purpose for which a concept has been developed and as such ‘We are asking not so much: when is the ascription of a certain concept correct, but rather, why is it applied? In freakish circumstances, a purpose may be achievable in unusual ways’ (Craig 1990:14). As such, we should not expect a neat set of necessary and jointly sufficient conditions but rather

conditions that apply in prototypical conditions, some of which will apply in some cases and others in other cases. Many cases of knowledge, on Craig's view, will not meet *all* of the conditions that we would traditionally see as being necessary.

Although we may not want to adopt Craig's methodology completely, we may want to take on board the flexibility of conditions that he suggested. We could still search for conditions for knowledge, but instead of rigidly defining all cases of knowledge, these conditions may describe conditions for prototypical cases, allowing for some variation in what it takes for something to count as knowledge. There is also something to be said for allowing that our intuitions can be mistaken. If we have confusing, or even contradictory, intuitions about knowledge then we should not worry excessively about having to bite the bullet and deny one of our intuitions, or claim that one or two intuitions are not true for all cases of knowledge. This is not too far a move from 'traditional' philosophy. As we saw earlier the traditional job of philosophy is not merely descriptive. Rather, if we uncover inconsistencies in our concepts then we can adapt the intuitive concept to be more consistent. We do not have to be constrained to make our philosophical investigations exactly in line with our intuitive understanding.

#### **1.4 Towards a Comprehensive Epistemology**

As we have seen, conceptual enquiry remains a legitimate form of enquiry in epistemology, although it need not take the form of reductive analysis. However, as we have seen above, there are several alternative methodologies have been suggested. What of these? There is no reason why these methodologies could not also properly be included within a philosophical investigation of knowledge. Although the role of our intuitions and concepts is important, there should also be space for naturalism, experimentalism and facts about ordinary practice.

However, there is one obvious problem with being so broad: there are likely to be considerable differences between the theories of knowledge that are produced by these different approaches. A naturalist will emphasise how knowledge fits in to the natural world, perhaps emphasising the continuity between human and non-human animal knowledge. A view that focuses on social practice, on the other hand, is likely to focus much more on the human side of things. The perspective with which one approaches the question 'What is knowledge?' is bound to have an effect on the theory of knowledge that is ultimately produced. In Chapters Three and Four we shall examine two different approaches, those of Kornblith and Craig respectively, and look at how each leads to a very different theory of knowledge.

Although different theories will be produced by having different approaches to the questions of epistemology, this does not mean that the different theories are irreconcilable. Knowledge is a complex concept and it is likely that we have competing intuitions about it. If we bear in mind that we are not looking for a reductive analysis of the concept but rather conditions that define a prototypical case, then we can allow that something that is required for knowledge in one case may not be required in another. There may even be more than one prototype of knowledge, for example: one for human knowledge and another for non-human animal knowledge. If we are no longer seeking a reductive analysis then we have more flexibility in the theory that we finally develop. Thus, we may be able to take on board insights from various different methodologies. The more we can take on board, it seems, the more comprehensive our understanding of knowledge will be and we should aim for as comprehensive a theory as possible. We should not limit ourselves to conceptual enquiry alone but take on board insights from as broad a base as possible.

## **1.5 Conclusion**

In this chapter we have discussed various suggestions for how we might go about philosophy, and what philosophy should answer to. Although it has faced criticism in recent times, conceptual analysis will still be an important part of any investigation into knowledge. However, we should not think that this alone will be able to give us a complete epistemology. We need to look not only at the concept that we have but also at the associated phenomenon, the social role of knowledge and take on board any experimental data that there may be. It is only once we have taken all these different factors into account that we will have a comprehensive account of knowledge.

---

## CHAPTER TWO

### Internalism and Externalism

---

The debate between internalist and externalist theories has come to dominate contemporary epistemology. Internalism about knowledge claims that we need to have reflectively accessible justification for our beliefs, if we are to have knowledge. Externalism, on the other hand, denies that we need to have any such reflective access in order to know, instead the knower simply needs to be connected to the facts in the right way.

There seems to be some intuitive plausibility to internalism but once examined in more detail it appears problematic. It seems impossible to give an answer to scepticism from an internalist point of view, for example. Externalism, on the other hand, can give a relatively simple answer to scepticism. Provided that we are not in fact being deceived then there is no reason that we cannot know many things about the world. In fact, provided that the belief is formed in the right way, we can even have knowledge of the denials of sceptical hypotheses (See, for example, Pritchard 2005: 69).

Externalism is not without its problems, however. The fact that it eschews reflectively accessible justification leaves its answer to scepticism looking rather unsatisfactory. According to the externalist a subject lacks any reassurance that he is not being deceived. From the subject's perspective he could just as easily be a brain-in-a-vat as not. I suggest that the reason for this dissatisfaction is precisely that we cannot

have reflectively accessible justification that we are not being deceived. The dissatisfaction is not proof that the externalist is wrong but it is rather an indicator of the importance that we place on reflective accessibility. Thus, any adequate theory of knowledge should take account of internalist intuitions as well as externalist ones. The chapter will end by examining some theories that appear to take into account both sets of intuitions; McDowell's disjunctivism and Sosa's virtue epistemology.

## **2.1 Some Definitions**

Before discussing the debate it is necessary to define some terms. 'Internalism' and 'externalism' are not used in a uniform manner across the literature but I shall use them in the following way:

- **Internalism about justification:** justification is provided only by what is reflectively accessible.
- **Internalism about knowledge:** knowledge requires justification conceived internalistically.
- **Externalism about justification:** justification need not be provided only by what is reflectively accessible.
- **Externalism about knowledge:** knowledge does not require justification conceived internalistically<sup>4</sup>.

Sometimes the epistemological debate is couched in terms of whether or not one should be an internalist about justification but in other places it is set out in terms of

---

<sup>4</sup> This way of drawing the distinctions can be seen in BonJour and Sosa (2003) and Pritchard (2005), amongst others.

whether or not one should internalist about knowledge. Where it is taken for granted that justification is understood as inherently internalist, the debate is essentially about whether or not justification is necessary for knowledge (e.g. Craig 1990; Gettier 1963; Kvanvig 2003; Lewis 1996). Some, however, appear to see ‘justification’ as a placeholder for whatever condition differentiates knowledge from ‘mere’ true belief (e.g. Chisholm 1989: 75; Goldman 1980; BonJour 1980, 1985; Greco 2005; Kornblith 2002). Here, I have tried to separate the questions of justification and knowledge as they seem to be independent of one another. Even if we were to decide that justification was an internalist notion it is not clear that such justification would be necessary for knowledge. Perhaps, as Robert Audi suggests, ‘an internalist account seems preferable for justification and a[n] ...externalist one seems preferable for knowledge.’ (Audi 1998: 244)

## **2.2 Internalism**

This section will examine the case for internalism. Firstly, we shall look at two possible interpretations of what it means for justification to be ‘internal’, accessibilism and mentalism. We will then discuss the case for internalism about justification, concluding that there is one understanding of ‘justification’ that is inherently internalist. Whether such justification is necessary for knowledge, however is a different matter. We shall look at the case for and against this claim. The case against internalist knowledge appears to be overwhelming and gives motivation to the externalist alternative.



### **2.2.1 What constitutes ‘reflectively accessible justification’?**

*a) How should we understand ‘internal’?*

Whether or not there will be a good case for internalism about knowledge or justification will, in part depend on how exactly one spells out what it is for justification to be internal. There are two main options in the literature with regard to this: accessibilism and mentalism. There is also a third, more unusual, possibility – disjunctivism – but we shall leave discussion of that for later in the chapter. An accessibilist account of knowledge claims that for a belief to be known it must be justified in a way that is reflectively accessible. That is, you can confirm whether the belief is justified or not by reflection alone. The claim ‘I am in pain’ is a prime example of a belief which is internalistically justified. Mentalism, on the other hand, claims that the justification for a belief must be part of the mental life of the subject (Conee and Feldman 2001: 233).

The two positions are very similar and may, in practice, come down to the same thing. However, Earl Conee and Fred Feldman (2001:234-5) claim that mentalism is the preferable option. More specifically, mentalism can be understood in terms of the following supervenience thesis:

**S** The justificatory status of a person’s doxastic attitudes strongly supervenes on the person’s occurrent and dispositional mental states, events and conditions. (Conee and Feldman 2001: 234)

This echoes the internalism/externalism distinctions in philosophy of mind and in ethics. In philosophy of mind, for example, internalism is also construed as a supervenience thesis, where a subject’s mental contents supervenes on states and conditions internal to that subject. Conee and Feldman argue that internalism in philosophy of mind and epistemology are concerned with similar issues. They are both

interested in explaining how 'being in some condition which is of philosophical interest...is settled by what goes on inside of cognitive beings.' (Conee and Feldman 2001: 235). Equating epistemological internalism and mentalism makes this feature very clear; mental states, events and conditions are clearly internal to the subject.

Accessibilism does not bring out this internality as clearly. In order to do so, it needs to be supplemented by the additional thesis that 'what is relevantly accessible is always internal to something, presumably the mind' (Conee and Feldman 2001: 235). There is nothing inherently internal in the concept of 'access' in the way that there is with mental states. So, argue Conee and Feldman, the mentalist thesis is preferable. It seems that in practice the two views will be extensionally identical, so perhaps we should take Conee and Feldman as pushing mentalism as the more fundamental of the two ideas (For further discussion of this see Pritchard *forthcoming* b: 2-3).

However, it is not clear that the mentalist thesis fully captures the internalist intuition. One of the most important things about internalist justification is that it should be reflectively accessible to the subject but this is not necessarily the case for all of one's mental states. For example, one can imagine a form of content externalism that would mean that this is not the case (Pritchard *forthcoming* b: 3). If the subject's relevant mental states were not accessible to him then it seems unlikely that we would want to classify such a theory as internalist. Although the things that justify the belief are internal to the subject as far as mentalism is concerned he would not be aware of them.

It is this very lack of awareness that often underlies internalist criticism of externalism. If Conee and Feldman's theory can likewise make justification completely opaque then it seems that it has lost a large part of the motivation for internalism but without gaining any advantage that an externalist theory may have. The natural understanding of 'internal' justification is something that we can access by reflection

and this, rather than the mentalist conception, appears to be the preferable understanding of 'internal'. As the two theories are likely to converge anyway, there should be no significant disadvantage to adopting this view.

*b) What do we need to have access to?*

The internalist claims that justification needs to be reflectively accessible, but what is justification in this case? Traditionally justification has been understood as grounds or evidence for one's belief. Your belief is justified, according to the internalist, if you can back up your belief. Let's take a simple case: John knows that the apple is on the table. To back this up John might appeal to current perceptual experience that there is an apple on the table; his memory that the apple is on the table or the testimony of a reliable informant that there is an apple on the table. Any of these things, along with many others, would act as justification for a belief.

## **2.2.2 Arguments for internalism about justification**

*a) Justification and responsibility*

One reason for claiming that justification is internal, is the connection between justification and responsibility (e.g. Bonjour 1980:55, 1985: Ch.1; Ginet 1975: Ch. 3). One's belief is only justified if it is held responsibly – if you cannot be reproached for holding it. Responsibility appears to be tied up with internally accessible justification. Your beliefs are held responsibly if they fit well with the evidence that you have available to you. You are not to blame for believing falsely that the capital of Australia is Sydney if all of your evidence points in that direction. If, for example, you have been told by more than one (usually) authoritative source that this is the case. Even though

your belief is false, you are blameless. If, on the other hand, you have the same belief but that is because you persistently ignore the excellent evidence presented to you that the capital is in fact Canberra then you do not hold the belief responsibly; your belief is not justified.

Responsibility is tied up with duty. If it is our epistemic duty to believe truly and avoid believing falsely, then we have an obligation to try to believe truly. It is irresponsible to form beliefs on the basis of little or no evidence or to reject beliefs in the face of good evidence. There is, of course, a parallel with the moral case. It is your duty to save the life of a drowning child if you can. If, however, you are unaware that there is a child drowning in the nearby lake then you will not be held responsible for not saving that life. If, however, you hear the child's cry for help but ignore it then you will be blamed for not trying to help. If you try your best to save the child but fail then we will not blame you. In the epistemic case, you will not be blamed for holding a false belief if there was no way that you could find out the truth, for example, if you are the subject of a complicated and extensive practical joke which fools you into thinking the capital of Australia is Sydney.

If this conception of justification is correct then it seems that justification must be an internal matter. You cannot be held accountable for not taking into account things that are outside your ken. If all of the evidence, as it appears to you, points to the idea that Sydney is the capital of Australia then your belief is justified, even if it is false. If, on the other hand, your reflectively accessible evidence shows that this is not the case then your belief is not justified. Thus, justification is an internal matter.

*b) The 'New Evil Genius' view*

A second justification for internalism about justification is the 'New Evil Genius view' (Neta and Pritchard 2007: 381. See also, Conee and Feldman 2001:234; Lehrer and

Cohen 1983: 92-3). This claims that two subjects who are exactly alike with regard to beliefs have the same justification. In other words, you cannot be more justified than your recently 'envatted' counterpart. The brain-in-a-vat, a modern version of Descartes' demon-victim, is a disembodied brain that is hooked up to a supercomputer which is feeding it experiences. These experiences are indistinguishable from the experiences of a normal human being. Thus, if you were a BIV you wouldn't be able to tell you were and, conversely, normal human beings cannot rule out the possibility that they are BIVs. The new evil genius intuition is that whatever justification you have your recently envatted counterpart has the same. You are exactly alike with regard to beliefs and reasoning and so you must be alike justificationaly. To hark back to the previous point, the BIV is epistemically blameless and so must be justified in his beliefs.

*c) Conee and Feldman's cases*

Conee and Feldman (2001) argue for an internalist account of justification with a series of thought experiments. In each case, one person appears to be more justified than another but the only difference between them is something internal. Here we shall examine two of their cases. The first case is that of two tourists, Bob and Ray, sitting in an air-conditioned hotel lobby (Conee and Feldman 2001:236). They have both read in the paper that today will be warm and consequently believe this to be the case. Bob then exits the lobby and walks into the blazing sunshine. At this point his belief becomes better justified than Ray's. Conee and Feldman's claim is that the only difference between the two men is an internal one and so the difference in justification must come down to an internal difference.

A second example regards two birdwatchers. On spotting a particular bird alighting on a tree, the expert bird-watcher and the novice bird-watcher both form the belief 'That bird is a woodpecker' (Conee and Feldman 2001: 237). The expert has good

reason to believe that it is a woodpecker. He's seen many a woodpecker before and knows what characteristics to look out for. The novice, on the other hand, doesn't have a good reason. He just thinks all birds of a certain size are woodpeckers. In this case it seems that the expert is justified in his belief but the novice isn't. If their visual experiences are the same then it seems that the difference can only be an internal one, having to do with their differing beliefs about woodpeckers. So, conclude Conee and Feldman, it is internal differences not external ones that make a difference to justification and so justification must be an internal matter.

### **2.2.3 Arguments against internalist justification**

John Greco argues that not only justification but in fact '*no interesting or important kind of epistemic evaluation...is internalist*' (Greco 2005: 266, my emphasis). He argues that there are two types of evaluation with regard to people and their beliefs; subjective and objective evaluation. The objective point of view evaluates how well beliefs fit with the world and the subjective point of view evaluates how subjectively appropriate it is to believe a particular proposition. However, claims Greco, it is not the case that the former is externalist and the latter internalist, rather they are both externalist forms of evaluation.

Responsibility, just like reliability, depends in part on etiology and accuracy and this is external to the subject's perspective (Greco 2005:266). Whether a belief is responsibly held or not depends in part on whether it was responsibly formed and this may not be internal to the believer's perspective. For example, Maria formed her belief that Dean Martin is Italian in an irresponsible way, on the basis of testimony that she knew to be unreliable. She now seems to remember that Dean Martin is Italian and has no reason to doubt that this is the case (she cannot remember how she formed the

belief). The fact that she formed this belief irresponsibly means that she is not justified, claims Greco, but this is an external not an internal matter; she is not aware of the fact that she formed the belief irresponsibly.

Having good reasons, claims Greco, is not the same as believing *for* good reasons. You might have good reasons available but believe a proposition for other, bad reasons. For example:

A math student knows all the relevant axioms but doesn't see how the axioms support a theorem that must be proven on the exam. Eventually he reasons fallaciously to the theorem, and believes it on the basis of his fallacious reasoning. (Greco 2005:261)

The fact that your belief is based on bad reasons may not be available to you and therefore not internal to your perspective, so justification is not purely an internal matter. Greco extends his argument to claim that *no* form of epistemic evaluation is internalist. One could stipulate a kind of justification that removed any externalist aspect but such evaluation would be entirely uninteresting. Etiological concerns will always have a role to play (Greco 2005:266).

However, this misses something important about our epistemic position. We are often in the position of having to evaluate our own beliefs and knowledge. In doing this, the only information that we have available to us is what is internally accessible. *As far as you can tell* your belief is held in a responsible way. This is surely an important sense in which you are justified. Greco claims that such justification is uninteresting but surely it will be highly interesting to the subject who is trying to assess how well justified his belief is.

Greco, as with many externalists, concentrates on an objective perspective of evaluation. 'We are social, highly inter-dependent, information-using, information-sharing beings' (Greco 2005: 226-7), he claims, and so it should be no surprise that we

are interested in the history as to how beliefs are formed. Whilst this is undoubtedly true, it does not show that external factors should be of primary importance. In order to be part of a network of information-sharing beings we have a responsibility to identify which beliefs seem the most likely to be true from our own perspective. We filter and assess our own information before sharing it and it would be irresponsible to pass on information for which we have no reflectively accessible justification. The very fact that we are part of a network of beings who share information shows that there is an interesting type of evaluation that is purely internal.

With this in mind, we can re-examine Greco's Maria example. It is not clear that Maria holds her belief responsibly. According to Greco, she simply 'seems to remember clearly that it is so, and... she has no reason for doubting her belief' (Greco 2005:261). It is not clear that Maria could count as blameless in holding this belief even if we take an internalist perspective. She cannot remember what evidence she had for the belief that Dean Martin is Italian. If you cannot recall any evidence, or any reason that you might have for believing that  $p$  then it would not be responsible to pass on such a belief, or claim it as knowledge. In terms of asserting knowledge, at least some qualification would be required. For example, 'I'm pretty sure that Dean Martin is Italian, but I can't remember why I think that.' Whether or not such justification is necessary for knowledge, it is of vital importance in our interactions with others. This significant aspect of justification appears to have been rather overlooked in contemporary epistemology and we shall examine it in detail in Chapters 4 and 5.

It is not clear, however, that the only way of understanding 'justification' is internalist. Williams (2001:21-3) suggests that there are two ways of assessing someone's belief, subjectively and objectively. Unlike Greco, however, he has an internalist understanding of subjective justification. Whether a belief is responsible is dependent on the evidence that the believer can access. We could say that this



dimension of assessment concerns whether or not the *person* is justified. We can also ask whether a belief itself is justified. Is the belief in actual fact well-formed, irrespective of what the believer thinks. Williams describes these types of justification as ‘personal justification’ and ‘adequate grounding’ respectively (Williams 2001:22).

There is, therefore, an important sense in which ‘justification’ is an internalist notion, but it is far from clear that it is the only way of understanding it. Indeed, as discussed above, some writers appear to understand ‘justification’ as a place-holder for the third condition on knowledge. If we adopt Williams’ position on this then we can allow for both the intuition that justification is an internalist notion and the intuition that it is whatever turns true belief into knowledge. For the sake of clarity I shall not use Williams’ terms but rather refer to internalist and externalist justification.

#### **2.2.4 The case for internalism about knowledge**

The arguments for internalism about justification often yield arguments for internalism about knowledge. Underlying this understanding is the idea that knowledge is, or implicates, justified true belief, whatever justification may turn out to be. This confuses the issue as there does seem to be a natural understanding of ‘justification’ which is internalist, and, as argued above, there may in fact be more than one way to interpret the term. In this section we shall look at how well the arguments for internalism about justification transfer over to claims about internalism with regard to knowledge.

In the preceding section we looked at the connection between knowledge and responsibility. Through his clairvoyant cases, Bonjour (1980: 59-65) argues that such responsibility is indeed required for knowledge. In the most interesting case, Norman is a completely reliable clairvoyant but has no reason to think that he is, or to think that such a power is even possible. One day, due to this reliable clairvoyant faculty, he

comes to believe that the President is in New York but he has no other evidence either way. According to BonJour this cannot be knowledge because it is not responsibly held.

Norman's belief satisfies the criteria for knowledge according to a reliabilist account of knowledge; he has a true belief formed by a reliable belief forming process. It is also likely to count as knowledge on other externalist accounts of knowledge. However, claims BonJour, it is irrational and irresponsible for Norman to hold such a belief. From Norman's internal perspective the belief about the president could just as well be an accident or a hunch, neither of which would be knowledge (BonJour 1980: 62). The problem is that there is no way for Norman to tell whether or not his belief was formed in a reliable way. Beliefs based on his faculty of sight are responsible because he has had enough experience to know that his sight is reliable. However, he has had no chance to test whether his faculty of clairvoyance is reliable or not and so does not hold beliefs responsibly based on that faculty.

In fact it seems odd that Norman is even forming beliefs under these circumstances. If he is unaware of his clairvoyance then presumably it simply strikes him as being the case that the president is in New York. Perhaps in the same way that it simply strikes me that the cup is on the table when I'm looking at the cup. However, if a fact simply strikes one, without any obvious cause, it would be very odd to accept it. If Norman was accustomed to suddenly form ideas about the president's whereabouts and had established some kind of connection, then it would make sense that he would take this being struck as meaningful. If, however, it comes out of nowhere then it seems odd that he should simply accept that the president is in New York on the basis of nothing more than a feeling that it is the case.

Certainly, when it comes to *assertions* of knowledge it is usually assumed that someone who claims to know can also put forward some evidence to back up their

claim and if they can't then it would be a reasonable response to say that they don't actually know. Grice (1975), for example, argues that the most important conversational maxim is that of quality, which breaks down into two further maxims:

1. Do not say what you believe to be false.
2. Do not say that for which you lack adequate evidence. (Grice 1975: 46)

It is conversationally inappropriate to violate either of these maxims, by stating as fact a proposition that you either believe to be false or for which you do not have adequate evidence, even if it is true. If you make a claim to know, either explicitly or implicitly, then you generate the conversational implicature that you have adequate grounds for your belief. If evidence is required in this way for knowledge claims then we might think that this gives us *prima facie* reason to think that knowledge itself requires evidence.

However, the fact that it may be inappropriate to claim to know does not mean that the subject does not in fact know. The conditions for having knowledge may very well be different from those for *claiming* knowledge. It is inappropriate to claim knowledge if you do not have evidence to back it up, but that does not mean that you do not have knowledge. As suggested in the previous section, it may be important from a social perspective that we do not claim knowledge when we lack evidence, but we may nevertheless have knowledge. Basically, the appropriateness or not of knowledge claims does not necessarily tell us anything about the conditions under which one can be said to have knowledge.

It could be, therefore, that the inappropriateness of Norman's belief is not so much to do with whether or not he has knowledge but more to do with whether it would be appropriate for Norman to claim knowledge. If he does indeed form the belief (which seems questionable), then nonetheless it would be inappropriate for Norman to claim

that he knows where the president is as he has no way of backing up his claim. This, however, is independent from the question of whether or not he has knowledge. Although there is a clear connection between knowledge assertions and responsibility, there is nothing in the Norman case itself that shows that this is the case for knowledge.

### **2.2.5 Problems with internalism**

#### *a) Scepticism*

Sceptical arguments pose a big problem for the internalist. In fact, the sceptical problem can be understood as arising from an internalist perspective on knowledge. The standard sceptical argument in the contemporary literature begins with the BIV case as outlined above. We can use this case to form the following sceptical argument:

1. If I know that I have two hands then I know that I am not a BIV.
  2. I do not know that I am not a BIV.
- Therefore,
3. I do not know that I have hands.

If I cannot rule out the possibility that I am a BIV, then I cannot know any everyday propositions as seemingly obvious as the proposition that I have two hands. As the BIV case is set up in such a way that it is impossible for a subject to rule it out then it seems that it will be impossible to know anything about the 'external' world.

The sceptical argument is often understood as a paradox; from two premises that seem intuitively plausible we can deduce an unacceptable conclusion. The conclusion is particularly devastating because it is not just the claim that we do not, as a matter of fact, have knowledge but rather that there is no possible way that we can have knowledge. Because of the way that the BIV case is set up there is no way that you

would ever be able to be sure that you are not a BIV and so the very possibility of having knowledge is ruled out.

The sceptical problem appears to arise out of an internalist idea of knowledge. The idea that we need to rule out alternative hypotheses, such as the BIV scenario, only makes sense if we think in terms of reflectively accessible justification. Your reflectively accessible justification may be undermined if you cannot rule out the possibility that you are a BIV because *as far as you can tell* you may be a BIV. This isn't problematic for an externalist account because no such justification is required for knowledge.

*b) Gettier Cases*

Gettier cases are situations in which a subject has an internalistically justified true belief that *p* but fails to have knowledge that *p*, demonstrating that internalistically justified true belief is not sufficient for knowledge<sup>5</sup>. For example, Jane comes down for her breakfast in the morning and, by looking at her (usually reliable) clock forms a true belief that it is 8.07. However, unbeknownst to Jane, the clock stopped exactly 12 hours ago and it is simply luck that she happened to look at exactly the right moment. If she had been even a minute earlier or later she would have formed a false belief. Although she has an internalistically justified true belief it seems that she cannot have knowledge; she's just too lucky<sup>6</sup>.

It might be thought that the problem in the above case is that Jane is in some way reasoning from a false premise. She is acting on the assumption that the clock is functioning as normal. However, Gettier cases have been suggested that supposedly involve no false lemma. For example, John looks into a field, sees a fluffy white animal and comes to a true belief that there is a sheep in the field. However, the animal he

---

<sup>5</sup> Gettier cases are named after Edmund Gettier who first proposed them in his 1963 paper 'Is Justified True Belief Knowledge?'

<sup>6</sup> This example is originally (pre-Gettier) due to Bertrand Russell (1948: 170-1). It has since been adopted as a Gettier case by Alvin Plantinga (1997) and Duncan Pritchard (2005), amongst others.

spotted was actually a big, white, fluffy dog which happens to be standing in front of a sheep<sup>7</sup>. There doesn't seem to be an error of reasoning in this case but John still has a justified true belief that does not amount to knowledge. That is, it is not clear that John is making an inference from a false assumption. Rather the error is of misperception rather than misreasoning.

Such cases demonstrate that internalistically justified true belief cannot be sufficient for knowledge, as internalist justification could always be undermined in this way. This means that any internalist theory of knowledge will need, in addition to internalist justification, an externalist 'anti-Gettier condition and such a condition runs the risk of looking gerrymandered and ad hoc. Gettier cases demonstrate that a purely internalist theory is not possible and that adding an extra condition may be problematic. This gives externalism a distinct advantage.

*c) Barn façade type cases*

Barn façade cases are similar to Gettier cases in that it is lucky that the subject forms a true belief, but luck is involved in a slightly different way. In Gettier cases something 'intervenes' between the subject and the belief that they form. Their beliefs are justified but not by the thing about which they are forming the beliefs. The barn façade case is different. Here the luck involved is 'environmental'. That is, nothing gets in the way of his belief forming processes, but rather something odd about the environment makes it impossible for the subject to have knowledge (Pritchard 2009; Pritchard *forthcoming a*).

In the barn façade case, Henry is driving along the road and watching the scenery. He spots a barn in the distance and consequently comes to the belief that it is a barn. It is, in fact, a barn and so he has a justified true belief based on a veridical perceptual experience. However, Henry is unaware of the fact that he is in barn façade county

---

<sup>7</sup> This case originates in Chisholm (1989:93).

where they pride themselves on their extremely realistic barn façades. As it happens, Henry is looking at the only real barn in the county. Again, luck comes into play here. It doesn't look like Henry can know because he's just too lucky. If he had looked at any other apparent barn, he would have formed a false belief.<sup>8</sup>

What is interesting about such cases is that it's not as obvious what has gone wrong with the subject's belief as it is in a Gettier case. After all, the belief is formed in a normal way and is true but we want to say that it is not knowledge. It seems that Henry's belief is too lucky but at the same time, he's come to that belief in a good way. Internalism cannot give an answer to this problem as there is no way that an internalist condition of knowledge could rule out such cases.

*d) Animal and infant knowledge*

There is also a problem for internalism when it comes to non-human and infant knowledge. We are generally happy to ascribe knowledge to animals and infants; we claim that the cat knows that there is food in her bowl or that the baby knows his mother. According to internalism this should not be the case. Neither dogs nor babies have the capacity to have reflective justification for their beliefs and so, according to internalism, they cannot have knowledge. By making reflection essential to knowledge, the internalist makes knowledge something that is impossible for any creature except a fully grown human (or at least an older child), thereby raising the bar for knowledge too high. One could, perhaps, suggest that animals do not really have knowledge or only have something analogous to knowledge but this doesn't seem to fit with the way that we use the term knowledge. We do not hesitate to ascribe knowledge to non-

---

<sup>8</sup> The barn façade case first appeared in Goldman (1976: 772-773), although Goldman has since acknowledged that it was originally due to Gettier, who suggested it in correspondence.

reflective creatures and so there is no obvious reason to exclude animals and infants from the sphere of knowledge and so this is another problem for internalism.

The case for internalism about knowledge, therefore, seems to be quite weak. It cannot answer Gettier or sceptical problems and nor can it account for our intuition that knowledge is not just restricted to fully grown human beings. As we shall examine in the next section, externalism does have an answer to these problems and can more easily fit with our intuition that knowledge can apply to animals and infants, which gives it something of an advantage.

## **2.3 Externalism**

### **2.3.1 The case for externalism**

The externalist claims that internalist justification is not necessary for knowledge. That means that a subject can have knowledge even if he doesn't have reflectively accessible justification. Instead having knowledge is a case of being connected to the relevant facts in the right way, for example, by having a reliably-formed true belief. This means that the externalist can give an answer to scepticism and has a more straightforward answer to Gettier problems than the internalist. Externalism also fits better with our intuitions regarding infant and animal knowledge and our intuitions about the knowledge of so-called 'chicken sexers'. This section shall examine these arguments which constitute the case for externalism.

#### *a) An answer to scepticism*

According to externalism, you do not need to be able to rule out the possibility that you are a BIV in order to know that you have hands. Instead, you can have knowledge of the 'external' world provided that your relationship to the world is of the right kind.



Thus, the fact that you cannot rule out the claim that you are a BIV is entirely irrelevant to whether you have knowledge or not. Having knowledge is simply a matter of having the right connection to the facts.

The sceptical problem only arises if we approach knowledge from an internalist perspective. The BIV scenario is set up in such a way that it is impossible to have reflectively accessible evidence one way or the other because the two situations are subjectively indistinguishable. However, as soon as you deny that reflectively accessible evidence is needed for knowledge then it is no longer a problem. The sceptical problem doesn't even get started. I can know that I have hands, provided that I am connected to the fact that I have hands in the right way. It may even be possible for me to know that I am not a BIV, provided that I am connected to that fact in the right way. I cannot have reflectively accessible justification for my belief that I am not a BIV but that is not an issue because such justification is not needed for knowledge.

One might object that this externalist answer has not proved that we do in fact have knowledge. Rather, it has only shown that we have knowledge provided that we are in the good case. That is, we are normally embodied human beings not BIVs. However, even if this is the case, the externalist is still in a much stronger position than the internalist. Firstly, the sceptical argument appears to show that knowledge is completely impossible but the externalist has demonstrated how it is possible. Secondly, given that that the world is pretty much as we assume it to be, then we do know a lot of the things that we take ourselves to know. It may not be possible to prove to the internalist's satisfaction that this is the case but provided it is then we have lots of knowledge.

*b) Gettier*

As we discussed above, in order to combat Gettier examples, any internalist account of knowledge will therefore have to add an externalist ‘anti-Gettier’ condition to knowledge, excluding the possibility of a ‘purely’ internalist theory of knowledge. The externalist, however, potentially has the advantage of simplicity here. Rather than having to add an extra condition to knowledge, an externalist condition can do two jobs at once. If we take knowledge to be reliably-formed true belief, then we do not need to add something extra to protect knowledge from Gettier cases as beliefs in Gettier cases do not appear to be reliably-formed; looking at stopped clocks is not generally a reliable way to find out what time it is.

However, although externalism has a ready answer to some Gettier cases, it is not clear that it can answer all the problems that they raise. In the dog/ sheep Gettier case outlined above, for example, it is not clear that any unreliable process is being employed. The subject forms his belief through his, presumably reliable, perceptual faculties. Nor is it clear that the externalist has the advantage in barn façade cases. Henry forms a true belief that what is in front of him is a barn using his reliable perceptual faculties and he is actually looking at a barn. In this case, it is not clear that the externalist has the advantage.

*c) Animal and infant knowledge*

Internalism, it is argued, sets the bar too high for knowledge, meaning that infants and animals cannot rightly be said to have knowledge. Externalism, on the other hand, has no such problem. If reflectively accessible justification is not needed for knowledge then there is no barrier to non-reflective creatures having knowledge. If, for example, having a reliably-formed true belief is sufficient for knowledge then this can be

achieved by any animal that can form beliefs reliably. Externalism fits our intuitions much better in this regard than internalism.

*d) Chicken sexers*

The chicken sexer example varies from writer to writer, and is sometimes presented as a real-life case and sometimes fiction, but in all cases it is supposed to show that externalist knowledge is possible (See appendix for a detailed comparison between real-life and philosophical chicken sexing). In the original chicken sexer case, Douglas Gasking (1962: 158-9) claims that the chicken sexer is trained to identify chicks by looking at photographs of chicks and making a judgement about whether they are male or female. He is then told whether he is right or wrong. After a period of time the trainee chicken sexer is increasingly correct in his judgements. Even though he may become completely reliable in distinguishing between male and female chicks, he cannot explain the basis for his belief. The example has since come to be used to demonstrate that a subject can have knowledge in the absence of reflectively accessible reasons. If the chicken sexer has knowledge in this case then it shows that we do not need reasons in order to have knowledge.

However, the internalist might object that the experienced chicken sexer does indeed have reflectively accessible evidence for his beliefs about the chicks. He has tested his ability on many occasions in the past and knows that he is good at distinguishing between male and female chicks. So, if we want to make it a true case of externalist knowledge then we need to remove this factor. If we take the trainee chicken sexer at a point before he has been tested but after he has become good at chicken sexing, then he does not have any reason to think that he even has the skill, never mind knowing how he does it (see, for example, Foley 1987: 168; Zagzebski

1996: 300-1). If we still think that the chicken sexer has knowledge in this case then it knowledge must be externalist.

A slightly different version of the example can also be found in the literature. Brandom (1998: 175-6. See also, Pritchard 2005:43) claims that it is not only the case that chicken sexers do not know how they form their beliefs but that they are actually mistaken about how they do it. Specifically, the chicken sexers think that they are using some visual cue to distinguish the chicks but in fact they discriminate on the basis of smell. This development in the story means that even the limited grounds that the chicken sexer may have had – “I know it’s male because it looks male” – is undermined, meaning that there are no internally accessible grounds available to the sexer.

The chicken sexer case, at least in its final version, is supposed to show that one can have knowledge without having reflectively accessible grounds for that knowledge. If we agree that the chicken sexer knows that a particular chick is male, or, even more minimally, that he knows that chicks in group A are different from those in group B, then reflection cannot be necessary for knowledge and therefore externalism is true.

### **2.3.2 Types of externalism**

The general guiding intuition behind externalism is that knowledge should not be accidental, or lucky. As with internalism, the way in which the guiding intuition is spelled out will make a huge difference to the plausibility of an externalist account. Here we shall examine two ways in which this has been done: reliabilist theories and modal theories of knowledge.

*a) Reliabilism*

The core intuition behind reliabilism is that knowledge is true belief formed in a reliable way. In particular, here we shall focus on ‘process’ reliabilism. That is, the idea that knowledge is true belief formed by a reliable process, a theory held by, for example, Alvin Goldman (1975; 1976; 1980; 1986).

According to Goldman, one needs to look at the history of a belief to find out whether or not it is justified. If it was formed by a reliable belief forming process then it is justified. Bill is a birdwatcher who has good hearing and is familiar with the different songs of common birds. In the garden he hears a familiar ‘cuckoo’ noise. From this Bill comes to believe that there is a cuckoo in the garden. As his belief is formed by a reliable process, Bill has knowledge. Kate, on the other hand, is prone to auditory hallucinations of the sound ‘cuckoo’ and so when she hears the bird and forms a belief on that basis she does not have knowledge. The method by which she formed her belief is unreliable even if it leads to a true belief in this case and so she does not have knowledge.

There are, however, some big problems with simple process reliabilism. Firstly, there is the ‘generality problem’ (Conee and Feldman 1997); the problem of specifying the process by which one comes to hold a particular belief. Reliability is not something that applies to one-off events. It is only ‘enduring mechanisms’ (Conee and Feldman 1997: 2) that can be reliable. This raises the question of how to define a particular process. There are various ways to describe the process that is at work when Bill identifies the bird. For example, it is a case of auditory-initiated belief formation; a case of identifying a bird based on the noise that is heard; a case of classifying an unseen bird on the basis of its song and so on. There are innumerable ways to describe the situation. The question is, which one we should focus in on to find out if Bill knows. Presumably the reliability will vary depending on how the process is specified. Unless

we have a systematic way of understanding what process is relevant we will not know which description is the correct one.

Another accusation is that reliabilism is guilty of 'bootstrapping' knowledge. Jonathan Vogel, for example, demonstrates this with his example of Roxanne (Vogel 2000: 613-4). Roxanne believes whatever her petrol gauge says, even though she has no reason to think that it is reliable. On the basis of this Roxanne often forms beliefs of the form 'The gauge currently reads 'Full' and the tank is full.' From this she deduces that the further belief that the gauge is accurate. So, assuming that the gauge is reliable, Roxanne has formed both of these beliefs through reliable processes and so is justified in her belief that the gauge is reliable, without forming any further checks or gathering any further evidence. With additional deduction she could even come to the conclusion that she is justified in believing that the gauge is accurate. As bootstrapping could be applied to many different processes, both reliable and unreliable, this is highly problematic for reliabilism.

Another problem is that of the Keith Lehrer's 'Mr. Truetemp' case (Lehrer 1990: 163). Lehrer imagines a man who has a microchip implanted in his brain such that he develops a completely accurate ability to tell what the temperature is. Whenever he thinks about the temperature he will come to a true belief about the temperature. However, he is unaware of his new ability and the device has the side-effect that he never checks whether his belief about the temperature is true or not, he just blindly accepts it. Lehrer claims that Mr. Truetemp does not know what the temperature is but according to reliabilism he must know – he has a true belief formed by a reliable mechanism.

The Truetemp case is analogous to both Bonjour's Norman case and also to the chicken sexer case. It seems that there is a difference in intuitions here. Internalists claim that Mr. Truetemp, Norman and the chicken sexer cannot know and the

reliabilists claim the opposite about all of them. It is not clear that there is any obvious way to defuse this clash of intuitions, but perhaps if there were an externalist theory that made space for some internalist knowledge that would help. We will examine two such theories below in Section 2.5.

*b) Modal epistemologies*

Simple process reliabilism seems to be too problematic to be correct. However, there are other ways of interpreting the core externalist intuition which are more successful. One such way is to describe reliability in terms of accuracy across a range of possible worlds. Such theories don't concentrate on history, i.e. on how beliefs are formed, but rather on how stable they are modally. That is, they focus on how easily a proposition could have been mistaken. A first suggestion of this nature is the idea of 'sensitivity'.

A belief is **sensitive** iff had it been false S would not have held it (Sosa 1999: 146).

If you test a child's temperature with a fully functioning thermometer, you would not believe that the temperature was normal unless it was. If, however, the thermometer was stuck at 37°C then you would still believe that the child's temperature was normal, even if he had a fever. So, if you are checking the temperature with a broken thermometer, you don't know that the temperature is normal, even if it is (Dretske 1971: 2).

Robert Nozick (1981:172-6) expresses sensitivity in terms of possible worlds: a belief that p is sensitive if in the closest possible world in which p is false the subject would not believe p. Nozick also adds a second truth-tracking condition: in the closest possible world in which p is true, s believes that p. That is, in a possible world which is slightly different from this one but in which p remains true, s would still believe that p. Nozick adds this extra conditional because there are cases in which belief is sensitive

but we would not say that the belief was known. For example, a dictator is killed and his death is reported in the newspapers. However, a false retraction is then published by all the news sources. Geoff fails to hear the retraction, and so continues to believe that the dictator is dead. He has a sensitive true belief – if the dictator had not died, he would not believe that he had – but we are reluctant to grant Geoff knowledge in this case (Nozick 1981: 177).

Sosa (1999) argues that sensitivity is not necessary for knowledge and that we should instead adopt an alternative modal condition, the safety principle. Sosa suggests the following counter-example to the sensitivity account of knowledge:

On my way to the elevator I release a trash bag down the chute from my high rise condo. Presumably I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) were not to arrive there? That presumably would be because it had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance. But none such could affect my predictive belief as I release it, so I would still predict that the bag would soon arrive in the basement. My belief seems not to be sensitive, therefore, but constitutes knowledge anyhow, and can correctly be said to do so. (Sosa 1999:145)

What this example appears to show is that beliefs can fail to be sensitive but still be knowledge. Sosa's belief that his trash bag will arrive in the basement fails to be sensitive because in the nearest not-p world – i.e. the one in which the bag snags – the subject would still believe that it was on its way to the basement. However, that possibility is so far off that intuitively it should not affect knowledge.

Sosa suggests that what is required for knowledge is not sensitivity but rather safety:

A belief is **safe** iff S would not easily have held it without it being true.



That is, in a range of nearby possible worlds in which S believes that p, p is true. Unlike sensitivity, safety is formulated in such a way that a wide range of nearby possible worlds is relevant to the assessment of a belief, rather than just the closest possible world. Safety does not face the same problem as sensitivity when it comes to the rubbish chute example. In the nearby worlds in which the subject believes that the rubbish is heading to the basement, it does end up in the basement.

The safety principle must be formulated in such a way as to require us to look at a wide range of close possible worlds, not just at the closest world[s] in which S believes that p. The reason for this is that there are cases in which a close, but not closest possible world undermines knowledge. For example, let's take another thermometer example, this time from Pritchard (2005:72-3; c.f. Goldman 1986:45). In this case there are several thermometers in the medicine cupboard which are all apparently identical but in fact most of them are unreliable. Sue cannot tell the difference between the reliable and unreliable thermometers and just picks one at random. As it happens, the unreliable thermometers are closer to the back of the cabinet and Sue is more likely to pick one from the front of the cabinet. She does in fact pick a fully functioning thermometer and forms a true belief about her son's temperature. If safety is restricted to only the nearest possible world[s] where Sue believes that p, then Sue's belief would be safe, because in the nearest possible worlds she chooses another functioning thermometer. However, we wouldn't say that she knows because it seems she could have all too easily picked an unreliable thermometer. However, if safety is formulated to encompass a wide range of close possible worlds in which Sue believes that p her belief fails to be safe and thus would fail to be knowledge under such an account.

Safety has another advantage over sensitivity. Sensitivity encounters a problem when it tries to answer scepticism (Sosa 1999:141-2). Your belief in an everyday proposition such as 'I have two hands' is sensitive and so constitutes knowledge

according to a sensitivity account. Your belief that you are not a BIV, on the other hand is not sensitive, because even if it were not true, you would still believe it. However, this means that we have to deny the closure principle, which states: if S knows that p, and S knows that p entails q, then S knows that q. This principle seems, intuitively, to be true but Dretske (1976), for example, decided to bite the bullet and deny closure. However, the safety theorist doesn't need to make any such difficult decision because your belief that you are not a BIV is safe. Unlike with sensitivity, the safety principle does not force you to look to distant possible worlds and in all nearby worlds in which you still believe that you are not a BIV, you are in fact not a BIV.

Duncan Pritchard (2002) suggests that the safety condition can be strengthened even further and calls this new condition 'super-safety' (2002:297):

An agent has a *super-safe* belief in the contingent proposition  $\varphi$  if, and only if that agent has a true belief in  $\varphi$  and her belief in  $\varphi$  'tracks' the truth as to whether  $\varphi$  across a wide-range of near-by possible worlds (i.e. where  $\varphi$  is true, she believes it; where  $\varphi$  is false, she does not believe it. (Pritchard 2002: 297)

Like a standard safety account, Pritchard's account limits the range to close possible worlds but the advantage is that this condition tracks both belief and non-belief. Whereas a simple safety condition could allow that there is a nearby possible world in which the target proposition is true but in which the subject doesn't believe it, Pritchard's account would rule out such a world.

### **2.3.3 A problem for externalism**

Although externalism about knowledge is in a better position than internalism, it is not without criticism. Barry Stroud (1989/2001) argues that externalism cannot provide a fully satisfactory answer to scepticism. Although externalism can restore the *possibility*

that we have knowledge it gives us no guarantee that we do in fact have knowledge. The best that we have is the conditional statement that *if* a certain type of theory is true then we can have knowledge. Beyond wishful thinking, it is not clear that we have any reason to accept such a theory. Nor do we have any reason to pick one kind of externalist theory over any other. Stroud suggests an externalist version of Descartes who claims, like the real Descartes, that his clear and distinct impressions are guaranteed by God to be genuine and truthful (Stroud 1989/2001:138-9). Unlike the historical Descartes, he does not need to resort to a circular argument justifying his belief in God. Rather, he can just claim that provided that God exists then he does know the things that he clearly and distinctly perceives.

Do we have any reason to reject this theory as opposed to reliabilism? Both explain how it is possible to have knowledge and Descartes' theory has the added benefit that it also explains why at least some of our belief-forming processes are reliable and guaranteed to yield knowledge; they are guaranteed by God. If the externalist Cartesian is in the stronger position then other than prejudice there doesn't seem to be any reason to choose the reliabilist account.

Stroud argues that the externalist has not, therefore, fully explained how knowledge is possible. In order for a theory to be a satisfactory explanation of how knowledge is possible, we need to have a reason to believe that the theory is true. Stroud claims:

...having an explanation of something in the sense of understanding it is a matter of having good reason to accept something that would be an explanation if it were true.  
(1989/ 2001:140)

The externalist has no such reason and no reason to accept his theory over another explanation, such as the externalist Cartesian. A externalist theory can explain how knowledge is *possible* but this gets us no closer to knowing whether we do in fact have

knowledge. If we are in the good case – i.e. we are BIVs and we are connected to the world in the right way – then we have knowledge but we can have no reassurance that this is the case. This lack of reassurance is what makes the externalist answer unsatisfactory; as far as I can tell, I might be a BIV.

Sosa is unconvinced by Stroud's criticism of externalism, claiming that he demands too much of a theory of knowledge. Stroud imposes the following metaepistemic requirement:

In order to understand one's knowledge satisfactorily one must see oneself as having some reason to accept a theory that one can recognise would explain one's knowledge if it were true. (Sosa 1994: 272)

However, even given Stroud's assumptions about the nature of epistemology, we do not need to go this far, claims Sosa. Stroud claims that we want to understand how we know the things that we know and also that understanding 'is a matter of having good reason to accept something that would be an explanation if it were true' (Stroud 1989/2001:140). However, this simply says that we need to have a good reason to accept the explanation; it says nothing of whether we should be aware of that good reason. According to an externalist theory, one can know that the theory is correct without seeing oneself as having a reason. So, even if the externalist does not *see* himself as having a reason to accept his theory, that does not mean that he does not *know* his theory. Stroud's insistence on needing to see oneself as having a reason just uncovers a deep-seated internalism that he holds on to.

However, Sosa's defence doesn't seem to get quite to the heart of Stroud's problem. Stroud can accept that the externalist knows that his theory is correct but what he wants to deny is that he has a full understanding of his knowledge. In order to *understand* how he has knowledge he needs to have a reason to accept that the theory is

true. Stroud's claim is not about what is necessary for knowledge, but rather for understanding, which we may take to be more demanding than knowledge. We won't be in a satisfactory position with regard to understanding our knowledge until we have an explanation that we also have good reason to accept as true, even if we know it to be true. The problem that I take Stroud to be getting at is that to know externalistically is not sufficient for understanding; understanding requires having a good reason to accept the theory as true.

Crucially, the type of reason that Stroud is thinking of here is a motivating reason. That is, a reason that the believer is aware of, a reason that is capable of motivating one to act. The mere fact that my flight is cancelled will not prevent me going to the airport, even if it constitutes a good reason not to go. It will only stop me driving to the airport once I become aware of it, in which case it stops me from setting out. In order for an explanation to be adequate one needs to have a motivating reason to accept it. That is, one has to be aware of a reason for accepting that theory as true.

What seems to be at the heart of what Stroud is saying is the idea that in order to rationally accept one theory rather than any other of the same explanatory power, you need to be aware of a reason to accept that theory. If, as Stroud claims, the externalist Cartesian theory is just as good, if not better, at explaining how knowledge is possible then why should we accept contemporary externalist accounts as opposed to a Cartesian version of them. It is all well and good that one may know that if one's externalist theory is correct then one knows. It is different to claim that your theory *is* correct without being able to give reasons for why it is *your* theory that is the correct one.

The problem is, in part, that you are in a very weak position dialectically. You can bring forward no reason for why someone should favour your theory over others. You cannot even give a good reason as to why *you* should favour your own theory. We

expect people to be able to justify their beliefs. If they cannot do this then it is not clear that their understanding is complete. A school pupil may know that  $E = MC^2$  but if he can bring forward no reasons for believing that to be the case, beyond the fact that his teacher told him so, then we would think that he does not understand the theory of relativity. It is important that people do not just have knowledge but that they also have reasons that back up their knowledge. Although we may not always require these reasons it is clear that at least in some instances a lack of such reasons will count as a deficiency.

What is lacking in externalism is not that it cannot explain how knowledge is possible but that it leaves us in an unsatisfactory position. The externalist may know that she has knowledge, in an appropriately externalist fashion, but he cannot give a reason why he, or anyone else, should accept his theory as opposed to some other. He is in the strange position that he may know that his theory is correct but cannot give any reason why he, or anyone else, should believe his theory to be the case. I think that what this brings out is that we have conflicting intuitions about knowledge. On the one hand there is the externalist intuition that reflectively accessible reasons are not necessary for knowledge. On the other hand is an internalist intuition that we need to have reasons to back up at least some of our knowledge for it to be satisfactory.

## **2.4 Two Possible Compromises: Sosa and McDowell**

Despite its own problems externalism is less problematic than internalism, but since we seem to have some internalist intuitions about knowledge then the best theory would be one that is externalist but that takes into account internalist intuitions in some way. There is not necessarily anything contradictory about this. As an externalist theory is one that claims that internalist justification is not necessary for knowledge, but that is

not to say that there is not any role for internalist justification, it is just that it is not necessary for knowledge. In this section we shall sketch out two different theories which take into account, to some extent, both internalist and externalist intuitions.

### **2.4.1 Sosa's virtue epistemology**

Sosa's safety account of knowledge has since developed into a form of virtue epistemology (Sosa 2007)<sup>9</sup>. Virtue epistemology comes in several varieties but what ties them together is a focus on the knower rather than the belief that is known. The idea is that a proposition is known when it arises out of an epistemic virtue. This parallels the virtue-theoretic ethical position that says that an act is good if it arises from a moral virtue. The question then is, how do we define what an epistemic virtue is? Sosa defines it in terms of competences. There are various ways in which we can assess a performance. First, it may or may not be accurate; that is, it may or may not hit its target. Second, it may or may not be adroit; that is, it may or may not manifest a competence on the part of the performer. Thirdly, and most importantly, a performance can be apt or not. An apt performance is one that is successful (accurate) *because* it is competent (adroit). When it comes to belief, a belief is apt – and so is known – if, and only if, it is true due to the exercise of an epistemic competence in appropriate conditions.

Relativising to appropriate conditions is an important part of the theory. Our competences are only competences in some conditions. The fact that I have the ability to make a good cup of tea does not mean that I would be able to make a cup of tea on a ship in a storm or on the top of Mount Everest. This distinguishes aptness from

---

<sup>9</sup> For other versions of virtue epistemology see Zagzebski (1996, 1999) and Greco (1999, 2003, 2007).

safety. A performance would be safe, according to Sosa, if it could not easily have failed. However, that the relevant conditions obtain could be a matter of luck so it could be the case that a performance is apt but that it's not safe. For example, it could be the case that the archer luckily chose the one un-drugged drink amongst lots containing a sedative. This would mean that his subsequent bulls eye shot is not safely accurate – he could very easily have taken a drugged drink and missed – however, it is an apt shot as it is because of his ability that he scores the bulls eye.

As with the shot, a belief can fail to be knowledge in several ways. One could have a true belief that was not *adroit*, as in cases where one trusts a source that is known to be unreliable but just happens to be right on this occasion. One could also have an *adroit* belief that is inaccurate. For example, when I come to believe that I have spotted a famous actress in a shop only for it to turn out to be a lifelike cardboard cut-out. Finally, I could have an accurate, *adroit* belief that fails to be knowledge, as in Gettier cases. If that same actress is standing behind her cardboard cut out then my belief that she is in the shop is true but not through the exercise of an epistemic virtue.

Sosa's theory of knowledge also has space for knowledge construed internalistically. Sosa wants to distinguish between 'animal' and 'reflective' knowledge (Sosa 1997: 422; 2007: Ch2, Ch5; BonJour and Sosa 2003: 2). Animal knowledge is just apt belief i.e. any belief that is true due to the exercise of a competence is a known belief, and this is true of adult human beings, infants and (at least some) animals. However, for normal adult human beings there is also the possibility of having reflective knowledge or 'apt belief aptly noted' (Sosa 2007:32). Reflective knowledge involves a meta-competence; not only has the knower formed a belief in a competent way but they are also competently aware that they have formed a belief in this way.

Animal knowledge is externalist in nature. The subject does not need to know that her belief is apt. That it is, in fact, apt is enough for it to count as knowledge. Human



beings, however, are capable of more than just animal knowledge; we can also have reflective knowledge. According to Sosa, reflective knowledge is ‘apt belief aptly noted’ (Sosa 2007: 32, 34). That is, the subject not only has an apt belief that *p* but also that she has an apt belief that she has an apt belief that *p*. For example, Geoff might think that he has forgotten all the French that he learnt at school, but when pressed on what the French for ‘hello’ is he answers ‘bonjour’. He knows the answer but lacks an apt belief – in fact, any belief – that he knows it. However, Yvette, who is bilingual in French and English, has not only apt belief but also an apt belief regarding that apt belief. She knows that she knows that ‘bonjour’ means ‘hello’.

By drawing the distinction between animal and reflective knowledge, Sosa can be seen as taking into account both internalist and externalist intuitions. Animal knowledge is clearly externalist, whereas reflective knowledge could be seen as taking into account more internalist intuitions. Internalists emphasise the need for reflectively accessible reasons for belief and it seems that to note aptly that one has knowledge would be inherently internalist, although not in a strong sense. Precisely what is involved in aptly noting one’s apt belief, however needs some more definition. Sosa claims that reflective knowledge:

requires...an apt apprehension that the object-level belief is apt. What competence might a believer exercise in gaining such meta-apprehension? It would have to be a competence enabling him to size up the appropriateness of the conditions. (Sosa 2007:108)

This competence, however, may just be ‘the default competence of taking it for granted that conditions are appropriately normal, absent some specific sign to the contrary’ (Sosa 2007:111). We will lack reflective knowledge, therefore, in situations in which we do not exercise such a competence. Presumably in situations where the conditions are not appropriately normal and there is no sign that this is the case.

Here we can see a similarity with what internalists claim knowledge to be. Crucially, reflective knowledge requires ‘apt apprehension that the object level belief is apt’ (Sosa 2007:108). This implies that some sort of realisation on the part of the subject that she does in fact have knowledge, which is exactly the kind of thing that internalists say. Reflective knowledge, according to Sosa, requires not only that your belief is appropriately formed but that you are aware that it is appropriately formed. For example, I have reflective knowledge that there is a cup of tea on the table because I can see that it is there and I am aware that the conditions are appropriately normal.

It must be noted, however, that Sosa’s reflective knowledge is not strongly internalist in the way that some internalists would demand. It requires only that you have an apt belief that you have an apt belief. That is, you must have a second order belief that your first order belief is well formed. This second order belief, however, just needs to fulfil the same requirements as a first order belief. That is, it must be true because well-formed. It is not clear that this brings in reflection of the sort that is often in play in internalist accounts of knowledge. You do not necessarily need to be able to bring forth evidence for your belief, for example. All you need to do is to have a well-formed belief that your first order belief is well-formed, and that these beliefs must be true because well-formed. Thus, from the perspective of the knower you may still not be able to give any reasons as to why you have knowledge other than that you know that you are reliable. Although this may not be considered properly internalist by some it seems that it is in the internalist spirit. Your knowledge is not completely blind, you can at least give an answer to the question of how you know what you know – you know because you are good at forming beliefs about this particular matter.

### **2.4.2 Disjunctivism**

Another, more unusual, suggestion which takes into account internalist and externalist intuitions is disjunctivism. One of the fundamental motivations behind internalism is the idea that subjects who are exactly alike mentally must be alike justificationaly. However, despite being internalists in an important sense, disjunctivists, such as John McDowell (1998), deny that this is the case. McDowell can nonetheless be classed as an internalist because he thinks that one's justification is constituted by facts that one can know by reflection alone, although, again, not in the strong sense suggested by some internalists. McDowell believes that one's justification is constituted by a particular type of reflectively accessible state; *factive reasons*. A *factive reason* for a proposition *p* could be that I *see that p* or that I *hear that p*. These are *factive reasons* because such a reason to believe a proposition entails that the proposition is true. I can only *see that p* if *p*. Alan Millar describes seeing that *p* in terms of having 'cognitive contact' (Millar 2007:186) with the facts. When you see that *p* there is a connection between you and the fact that *p*.

If we compare the two subjects *S*, in the actual world, and *S\**, her recently envatted counterpart, then we can see that according to McDowell's framework the two subjects are not alike with regard to justification, even though they may appear to be. *S*'s justification for her belief that *p* – for example, the proposition that there is a cup of tea on the table – is that she *sees that p*. However, the justification that *S\** has for that same proposition can only be that she *appears* to see that *p*. She cannot see that *p* because she is envatted. She has no eyes with which to see and the experience that she is having is not of an actual cup but is an electrically induced hallucination of a cup. In the good case, the experience that the subject is having 'takes in' the world but this doesn't happen in the bad case. Although both experiences may appear the same to the subjects, only the one in the good case has cognitive contact with the facts.

McDowell criticises thinking in terms of the 'highest common factor' (McDowell 1998: 386-9). This way of thinking assumes that the reasons that support our beliefs in the good case and the bad case must be the same. Because the experience of the two subjects is identical, it is tempting to think that both subjects are equally justified. Therefore, if our reasons fall short in the bad case, then they must also fall short in the good case. McDowell thinks that the highest common factor assumption is false. The reasons that one has in the good case connect to the world but in the bad case they do not, therefore they are not the same reasons.

The disjunctivist approach appears to have one advantage over the other types of internalism namely that the disjunctivist may not need to add an extra, externalist anti-Gettier condition to his account of knowledge. If, in Gettier cases the subject does not take in the world in the right way then she will not have knowledge. It also has the advantage that it takes account of some internalist and some externalist intuitions. It places an emphasis on reflective access but at the same time insists there must be an appropriate connection to the facts.

## **2.5 Conclusions**

In this chapter, I have argued that internalism about justification is plausible but when it comes to knowledge, externalism is more plausible. However, this does not mean that the importance of reflection should be downplayed. We place a lot of importance on justification, as evidenced by the dissatisfaction with externalist answers to scepticism. Internalist justification also has an important role to play with regard to knowledge claims. Thus, any good theory of knowledge will try to take into account both internalist and externalist intuitions.

---

## CHAPTER THREE

### A Naturalist Approach

---

Hilary Kornblith argues that we should take a naturalist approach to epistemology and that this leads to a strongly externalist theory of knowledge. As we saw in Chapter One Kornblith argues that knowledge is a natural kind and that as such we should not investigate it conceptually but in the same way that we investigate other natural kinds: empirically. Kornblith's argument that knowledge is a natural kind turns on the fact that knowledge plays an important causal and explanatory role in cognitive ethology – the science of animal cognition. Only natural kinds can play such a role and so knowledge must be a natural kind.

Having put forward his argument that knowledge is a natural kind Kornblith goes about developing a theory of knowledge. Using examples from cognitive ethology he concludes that “The concept of knowledge which is of interest here...requires reliably true belief.” (Kornblith 2002: 58). He then goes on to argue that human knowledge and the knowledge of non-human animals is of the same kind. Neither reflective access to our reasons for belief nor any particular social interactions are necessary for knowledge, he argues. Since these peculiarly human activities are not necessary for knowledge there is no need to draw a distinction, as Sosa does, between human and animal knowledge. Thus, any theory of knowledge correctly developed from cognitive ethological data should apply equally to human and animal knowledge.

I take issue with Kornblith on two counts. Firstly, it is not clear that his argument that knowledge is a natural kind is correct. Secondly, even if Kornblith is right that animal and human knowledge are of the same kind that does not mean that we should therefore discount the importance of reflection and social factors to human knowledge. That is not to say that either is essential to knowledge but even if they are not *essential* they play such an important role that any investigation that focuses simply on the commonality between human and non-human knowledge and ignores its differences will not be a complete theory of knowledge.

### **3.1 Knowledge as a Natural Kind**

In this section we shall examine Kornblith's argument that knowledge is a natural kind. If knowledge is a natural kind then investigating the concept of knowledge will not tell us anything interesting about the actual phenomenon. Something can only play an important causal and explanatory role in a natural science if it is a natural kind and Kornblith argues that knowledge plays this role in cognitive ethology. So, knowledge must be a natural kind. As a natural kind conceptual investigation is an inappropriate method of investigation. Instead we should look at the phenomenon of knowledge which we can discover through cognitive ethology.

#### **3.1.1 What is a natural kind?**

Natural kinds are roughly groupings or orderings that occur independently of human beings. Standard examples of natural kinds include chemical elements, chemical compounds, species and elemental particles. There are various things that are commonly considered characteristic of natural kinds. Firstly, members of a natural kind should have at least some properties in common. All instances of gold have the atomic

number 79, for example. It may be, however, that not all the members share any one property; a tiger that is born without a tail or who loses a limb is still a tiger. Natural kinds should also permit inductive inferences and participate in laws of nature. Natural kinds are therefore at the heart of the natural sciences (Bird and Tobin 2009).

There is much debate about the exact nature of natural kinds which I shall not enter into here. Instead I shall assume for the purpose of this discussion that the account of natural kinds that Kornblith adopts is the correct one. Following Richard Boyle, Kornblith claims that natural kinds are

...homeostatically clustered properties, properties that are mutually supporting and reinforcing in the face of external change. (Kornblith 2002: 61)

There are certain 'contingently clustered' (Boyd 1988:197) properties in the world, properties that co-occur in a number of important cases. Tigers usually have four legs, stripes and sharp teeth. These properties are clustered together. However, no single one of these properties is necessary or sufficient to be a tiger. A tiger that loses a limb is still a tiger even though it has lost one of the paradigmatic properties of being a tiger.

In the case of a natural kind, claims Boyd, this clustering must be due to homeostasis. That is, the presence of properties within a particular cluster favour the presence of other properties in that cluster. This could be due to intrinsic factors, such as a particular gene or atomic number, or to extrinsic factors such as a particular environmental niche. This means that any outlying cases are unlikely to last making the cluster self-regulating and stable. In the case of tigers, there is similarity between tigers' genes that explains why they tend to have four legs, stripes and sharp teeth. Tigers that are born with three legs or who lose their teeth will not, in most circumstances, survive and so the cluster is maintained.

On occasion, however, mutations occur and environments change in such a way that a new property may prove beneficial, or at least neutral, with regard to survival. This may lead to variation within a kind. If, due to some change in their environment, it became advantageous for tigers not to have stripes then over time stripeless tigers would become more common (or possibly take over completely, depending on how disadvantageous it was to have stripes). Natural kinds, therefore, can change over time, a property that at one point is possessed by most members of a kind could later become rare and vice versa.

According to Boyd, it is also essential that the homeostatic clustering is causally important. That is, there are important effects produced by a combination of the properties and the underlying mechanisms. The tiger survives well in its environment because it has camouflage and sharp teeth and these are a result of its underlying genetic structure. Natural kinds, under Boyd's definition, are therefore stable enough to be part of natural laws and to support inductive inferences.

### **3.1.2 Knowledge as a natural kind**

It is common in the cognitive ethology literature to use intentional idioms to describe animal behaviour. Kornblith cites various examples where the behaviour of animals and birds is described with intentional terminology. For example, in discussion of the behaviour of ravens in Mongolia, Bernd Heinrich (Heinrich 1999: 133, cited in Kornblith 2002: 31) describes ravens as 'distracting' an osprey. John Alcock refers to animals as "hiding from", 'spotting', 'invading' and 'repelling' predators' (Alcock 1975, cited in Kornblith 2002:30). However, as Kornblith is aware, the fact that people speak in intentional terms does not prove that the animals have such intentional states, it just shows that people talk as if they do. People quite frequently talk about inanimate



objects as if they have intentional states even though we do not think that they have intentional states. Whilst I might complain that the printer doesn't want to print my document that is not evidence that I think the printer has desires. We need more evidence, therefore, in order to reach the conclusion that the ethologists are actually ascribing intentional states to their subjects.

It is impossible to find descriptions of animal behaviour described in non-intentional terms, however. Kornblith (2002: 32-7) argues that this is because there need not be any similarity between different bodily movements that come under the same intentional description. Even something as simple as the basic fight or flight reactions cannot easily be reduced to bodily movements alone. An animal may move its limbs quickly for many different reasons; to run away from a predator or rival, to run towards a food source or mate, or just for play. They may also move differently even when moving for the same reason; an animal could flee from a predator by running or walking, for example. The only way to understand a whole group of dissimilar actions as 'flight', therefore, is to refer to the reasons for the behaviour. The ravens in Heinrich's raven case work in pairs to gather eggs. One distracts a nesting osprey whilst the other steals its eggs. The raven doesn't necessarily use the same bodily movements every time it distracts a raven. It is only when we understand these diverse movements in intentional terms that we will fully understand the behaviour.

Kornblith (2002: 37-42) argues that we can go further than simply ascribing intentional states to animals; we also have the basics for a belief-desire psychology. Animals have certain biologically driven needs, they need to interact with their environment, and have information-bearing internal states about their environment. Informational states about the environment and ones that are responsive to it can be described as beliefs of a sort and biologically driven needs are akin desires. Although plants also react to the environment we do not need to attribute beliefs and desires to

them. Plant behaviour can be explained in non-intentional, lower-level terms – reactions to stimuli – but animal behaviour cannot. To fully understand animal behaviour we need to be able to group bodily actions under intentional descriptions.

We should not worry that ascribing intentional states to animals will constitute anthropomorphism (Kornblith 43-8). Of course we should be cautious not to ascribe too much to animals but we should be equally cautious about ascribing too little. In general the literature has almost entirely focused on the former. Even when too much is ascribed to an animal, the debunking explanation does not necessarily deny intentionality completely.

Kornblith gives the example of tests done on young chimps (Povinelli and Eddy 1996, cited in Kornblith 2002: 48-52). Their behaviour appeared to show that they were ascribing intentional states to others around them. For example, they would ask for food from a researcher that was facing them but not from one who was looking away, and they would follow the gaze of a researcher. This seems to imply that they had some understanding of whether or not the researcher could see them. However, further testing showed this initial hypothesis to be false. The chimps acted in the same way to a researcher with a bucket on her head and one with a bucket beside her head, for example. It seems, therefore, that the chimps are not aware of whether or not the researcher could see them as was previously thought. This does not show, however, that psychological explanations of chimp behaviour in general are incorrect, but rather that the particular hypothesis in this case was false. It certainly does not mean that all psychological explanations of their behaviour should be abandoned.

For many cognitive ethologists, once intentional terminology in general has been deemed acceptable, use of the term 'knowledge' is relatively uncontroversial, claims Kornblith (Kornblith 2002: 55). However, there are some that look in more detail at the issue of knowledge specifically. One example that Kornblith discusses is that of the

work of Louis Herman and Palmer Morrel-Samuels's (1990/1996) research on dolphins. Kornblith claims to find in their work the idea that knowledge attributions are derivative from knowledge-acquiring abilities. When explaining the behaviour of individual animals we explain behaviour in terms of beliefs and desires. According to Kornblith 'a distinction between belief and knowledge is simply irrelevant here' (Kornblith 2002: 57). However, he claims, talk of knowledge becomes relevant when looking at the species level.

What Kornblith seems to be saying is that when trying to understand the behaviour of an individual animal we need not make reference to knowledge, only to beliefs and desires. When we want to explain the success of a particular species or type of animal then we need to bring in talk of knowledge. A particular species is successful in its environment because its cognitive capacities are adapted to that environment. Its success is explained by its ability to gain knowledge of its environment. We can explain why the plover leaves its nest in an attention-grabbing fashion on the basis of beliefs and desires – it believes that a predator is nearby and it wants to protect its eggs – but the plover would have acted in this way whether or not its beliefs were true. Knowledge is not important at this level, claims Kornblith. However, when it comes to explaining why plovers in general have such capacities we need to refer to knowledge. Plovers are successful in their environment because they are good at detecting predators and good at distracting them, so that their eggs do not get eaten, ultimately leading to more plovers. Plovers do not just need to be able to form true beliefs in this regard but they need to do so reliably. If they were not good at detecting predators, and doing so reliably, then they would not have been successful within their environment.

However, it is not at all clear how Kornblith comes to these conclusions on the basis of the research that he cites. Kornblith claims that that 'knowledge, on this view, first enters our theoretical picture at the level of understanding of the species, rather

than the individual' (Kornblith 2002:57). However, Herman and Morrel-Samuels (1990/1996: 289) talk solely in terms of knowledge at the individual level. At no point in their article do they mention belief. They give a list of some of the 'the basic knowledge requisites for the adult dolphin' and claim that 'extensive knowledge of the world may be required for effective functioning in that world and much of the requisite knowledge is gained through the exercise of receptive skills' (289). This is completely at odds with what Kornblith claims. It is true that they talk about 'knowledge-acquiring abilities' and 'receptive competences', which presumably equate to reliable processes of belief formation but these competences are described as 'supporting knowledge acquisition'. It is not clear that they see knowledge attributions as deriving from the abilities. It is not clear, therefore, where Kornblith derives his idea that knowledge becomes important only at the species level.

Even if Kornblith's interpretation of Herman and Morrel-Samuels were correct it is not clear why one would *need* to talk of knowledge at the species level, anyway. Why not just say that the dolphins are successful because they consistently form true beliefs? It is not clear that talk of knowledge is any more necessary at this level than at the level of the individual. Kornblith claims that:

...the presence of cognitive capacities across individuals...is itself something that requires explanation; explanation of the presence of such cognitive capacities requires that we advert to knowledge. (Kornblith 2002: 58)

This seems to be the opposite of how epistemologists normally understand knowledge. Usually we think of certain cognitive capacities as explaining how we have knowledge but Kornblith seems to be suggesting that we only bring in knowledge to explain the capacities. If we are to reverse the usual direction of explanation then Kornblith needs to do more to show that his understanding is the correct one. It is true

that we need an explanation for the cognitive capacities that help a species to survive, but it is not clear that this explanation must make reference to knowledge.

### **3.1.3 Kornblith and the value of knowledge**

It is often claimed that people value knowledge and mere true belief differently, that knowledge is valued more highly than mere true belief. We shall look at this issue in more detail in chapter five, but it is an issue that may be problematic for Kornblith. Even though the results of our actions would be the same whether we act on knowledge or a true belief, we still seek out knowledge if we can. It has also been suggested that knowledge closes enquiry. Once we come to know whether or not  $p$  we can stop investigation into the matter. For human beings, at least, there seem to be important differences at the individual level with regard to how we treat knowledge and true belief. It is not clear, on Kornblith's theory, why we value knowledge and true belief differently.

Human beings care very much whether they and other people know or merely believe. It is different when it comes to non-reflective, non-human animal knowers, however. It doesn't seem that it would make a great deal of difference whether or not a cat knows that its food is in its bowl or whether or not it just believes it. Either way it will get to the food. It seems that we can just as well describe the behaviour in terms of beliefs and desires as we can in terms of knowledge. As we saw above, Kornblith thinks that we only need to describe it in terms of beliefs and desires. That is not to say that we can't, or don't, describe the behaviour in terms of knowledge but rather that it doesn't seem that very much hangs on the distinction between knowledge and true belief when it comes to animal knowledge.

When it comes to the behaviour of human beings, however, a lot more seems to hang on this distinction. Knowledge plays a role in our intellectual lives in a way that it doesn't for non-reflective knowers. We seek out knowledge and justification for our beliefs and do not consider a matter settled until we have good reason to think that we know the answer. We want to know rather than merely believe and this seems to be the case even if there is no practical difference between the two<sup>10</sup>.

Kornblith does give an explanation for epistemic normativity in general (2002: Ch 5.). Knowledge is conducive to fitness and so it is valuable because it helps animals to reach their goals. Humans are no different with regard to this; knowledge in general is instrumentally valuable, it helps us to survive and to achieve our goals. Even though knowledge is no better at getting you to your goals than true belief, it may be that we value knowledge more, because in general it is good to have knowledge. If you form your beliefs unreliably then you will often end up with false beliefs, something that will not generally aid your survival.

However, this would not answer the question as to why we prefer knowledge to merely true belief. It is true that in general it is better to have knowledge than belief but on any particular occasion it would make no difference whether that true belief was reliably formed or not, provided that it is true. What is valuable, it seems, is the *ability* to reliably produce true beliefs. It is this ability that is conducive to fitness, after all. This would not answer the question of why knowledge is more valuable than true belief. This does not, however, give any explanation of why we value knowledge over true belief even if there is no practical difference between the two. Kornblith's response to this could be that this value difference is not dependent on knowledge itself but the concept that we have of knowledge. His investigation is explicitly

---

<sup>10</sup> I shall discuss this question further in Chapter 5.

focussed on the natural phenomenon of knowledge not our human concept of knowledge (Kornblith 2002:3), and so the value difference is not relevant. We may value gold more than silver, but that does nothing to tell us what gold or silver actually are. The value of gold will vary from place to place but its atomic number remains 79. However, if there are important facts about knowledge that Kornblith cannot account for then they must surely be problematic for his theory.

A lot of what Kornblith claims appears to turn on his claim that knowledge is a natural kind. However, it's not clear that he needs to go that far. For the central claim that knowledge should not be investigated by conceptual analysis, all he needs to claim is that there are *some* truths about knowledge that are not conceptual truths. If it were the case that every truth specifying the nature of knowledge was a conceptual truth then it would be appropriate to investigate it by conceptual analysis. It is, however, inappropriate to investigate something about which there are conceptual and non-conceptual truths in a way that only takes account of the conceptual ones. If this is the case then we would need to investigate both the conceptual and the non-conceptual aspects to get the full picture.

### **3.1.4 Bermúdez's criticism of Kornblith**

José Bermúdez (2006) argues that it is not clear that even if knowledge plays an important causal explanatory role in cognitive ethology that this proves that it is a natural kind. Cognitive ethologists often characterise themselves as applying common-sense belief-desire psychology to the behaviour of animals. If this is what they are trying to do, then there doesn't seem to be any more reason to think that this would be part of a mature science than that folk psychology would be. Bermúdez suggests that it depends on which area of science you look to as a guide. Whilst modified versions of

belief, desire and knowledge are likely to be found in a mature computational cognitive science, there are unlikely to be any such things in mature cognitive neuroscience.

Kornblith is happy to accept that the phenomenon knowledge may very well differ from our stereotype of it, after all, it is not the concept that he is interested in. However, it is not clear that we have a good way of identifying the phenomenon apart from our stereotype of it. In the case of ‘water’ or ‘gold’ there are canonical instances which fix the reference of each term but it is not clear that there are similar canonical examples when it comes to ‘knowledge’. Many of the disagreements in epistemology, for example, between internalists and externalists, seem to come down to the two sides taking different examples to be canonical. If we cannot agree on some canonical instances then it will be very difficult for us to even work out which phenomenon we should be investigating (Bermúdez 2006: 306-7).

Kornblith’s answer to this problem is to suggest a new set of canonical instances of knowledge – those that we find in the literature of cognitive ethology. If we take these cases as canonical then we can investigate them to see what knowledge consists of, in the same way that we can investigate what water is by examining clear-cut samples of water. In order for this move to work, however, Kornblith needs to show that the knowledge that is discussed in cognitive ethology is the same kind of knowledge that epistemologists are interested in. This, claims Bermúdez, is the most challenging part of Kornblith’s project.

It is not clear, therefore, that Kornblith has done enough to show that knowledge is a natural kind. However, even if his arguments are successful he has only shown that knowledge *as described in cognitive ethology* is a natural kind. Epistemologists are generally most concerned with human knowledge, and various theories claim to show that there is a difference between animal and human knowledge. So, in order for Kornblith’s argument to be fully successful he needs to show that there is no relevant distinction



between human and animal knowledge. We shall examine his arguments against such a distinction in the next section.

### **3.2 Only One Kind of Knowledge**

There are two types of theories that claim that human knowledge is different from the knowledge of non-human animals, claims Kornblith. Theories of the first type assume that meta-cognition at the social level is necessary for knowledge (Kornblith 2002: Ch 3) The second type is theories that claim that meta-cognition at the individual level is necessary for knowledge (Kornblith 2002: Ch 4). However, Kornblith wants to deny that there is a good case for theories of either of these types. He aims to show that there is only one type of knowledge in order to show that the conclusions that can be drawn from cognitive ethology apply as well to human beings as to non-human animals.

#### **3.2.1 Knowledge and Social Practices**

One way of drawing the distinction between human knowledge and non-human animal knowledge is to claim that knowledge requires the ability to participate in certain types of social practice, and since animals are incapable of engaging in these practices it is not possible for them to have knowledge. Donald Davidson (1984) claims that to have belief one must be language user. Robert Brandom (1994; 1998) goes even further and claims that in order to be capable of belief you have to be able to engage in the practice of giving and asking for reasons

Brandom (1994:214) claims that in order to believe, you need to be able to engage with the social practice of giving and asking for reasons. If you cannot do this then you

will not be able to understand the claim being made and so cannot know it. Kornblith's first comment about this is that different people participate in these practices to greatly varying extents. Some people, who Kornblith terms 'highly Socratic' (Kornblith 2002: 76), engage in this practice frequently whereas others, like Kornblith's 'laconic Yankees' (80), will rarely engage in it. Presumably, argues Kornblith, if one does not engage in the practice at all then one is not a genuine believer. It also seems that Brandom needs an answer to the question of how far one needs to get along the 'Socratic scale' (Kornblith 2002: 77) to count as a believer. Does one need to engage with the social practice on only one occasion or regularly?

Brandom's view of belief is counter-intuitive. It is not only that we would probably want to class people as believers who are completely silent with regard to reasons, but also that belief seems to be prior to the practice of giving and asking for reasons (Kornblith 2002: 78-9). The explanations that one gives for why someone partakes in the practice more or less appeal precisely to beliefs and desires. For example, we could explain someone's lack of participation in the practice of asking questions by saying that she believes it to be boorish. Giving and asking for reasons seems to be a product of belief, rather than a necessary condition for having beliefs.

In addition, it seems that we cannot make sense of the behaviour of infants or animals unless we ascribe beliefs to them. Brandom acknowledges that non-linguistic beings can have mental representations but not beliefs. However, if these representations help the creature to negotiate and respond to its environment then it is not clear, claims Kornblith, how they can be distinguished from beliefs. However, even if it does make sense to talk of belief and desire in animals that is not to say that animal belief and desire is the same thing that we have as reflective creatures. For example, as Alan Millar (2004: xi-x) suggests, one cannot rule out a 'high conception' of belief, such as Brandom's on the basis of apparent propositional attitudes in animals and infants.

Although animals may be capable of belief in some sense, it may be of a different kind from the belief of creatures such as ourselves who operate ‘within the space of reasons’. Even if we can only explain animal behaviour in intentional terms that does not mean that the states are of the same kind as those that human beings have.

The practice of giving and asking for reasons does seem to be deeply embedded in our culture. If you are a source of some information and know that the matter is important to the enquirer, then very often we will offer up the reasons without even being asked for them. If my friend asks what time the train is because she urgently needs to get to Stirling, I will not only reply that the train is a 10.33 but tell them that I regularly get that train and assure them of my certainty. If, on the other hand, I am not sure, then I will hedge my answer with ‘I think...’ or ‘I’m not sure, but...’, I will not make a claim to knowledge – at least with regards to something important – unless I can back it up with some reason for thinking that it is so. It is very hard to imagine someone who never at all engaged in the practice of giving and asking for reasons. Such a practice may not be necessary for knowledge but it is so deeply involved in our knowledge practices that it needs to be accounted for.

Davidson’s position doesn’t go quite as far as Brandom’s but does require that a creature be a language user in order to have beliefs (1984: 164). This leads him, like Brandom, to deny that animals can have beliefs. He acknowledges, however, that in the normal run of things we do ascribe beliefs to animals and so he presents several arguments against the idea of animals having beliefs. Firstly, beliefs are underdetermined with regard to non-linguistic behaviour. For example, we cannot tell which beliefs lead a subject to choose an apple rather than a pear on the basis of non-linguistic behaviour. It could be because it’s more tasty, more expensive, or simply because it is on the left, but we cannot settle this question with non-linguistic behaviour. However, there is a problem with this claim (Kornblith 2002:85-7). It seems

that any theoretical claim will be underdetermined in this way. We can only ever have finite data but our theoretical claims will have infinitely many implications. Nor will adding linguistic data help to solve this problem. We will still only have finite data for infinite implications. Underdetermination in itself is not a reason to think that language is necessary for belief.

Davidson (1984: 163) also claims that there is work that can be done by language that cannot be done by non-linguistic behaviour. We may describe the dog as knowing that its master is home but we cannot tell whether it knows that the president of the bank is home or that Mr Smith is home, where these are one and the same person. Kornblith (86) gives this claim short shrift, saying that it is just obvious that the dog does not know that the president of the bank is home since dogs have neither concepts of bank or president but if needed we could construct tests to show that this is the case. And, as to whether he knows that Mr Smith is home, we can create tests to see if dogs can recognise people by name. These issues, claims Kornblith, can be settled relatively easily without needing linguistic data.

Davidson (170) also argues that one cannot have belief without having the *concept* of belief, because in order to believe you need to understand the possibility of being mistaken. However, according to Kornblith (86-7), this claim could be given a weak or a strong reading and according to a weak reading animals can understand the possibility of being mistaken. A dog can come to understand that it's mistaken when it sees the ball in its master's hand rather than flying across the field. This reading is the only plausible one, he claims, but it will not do the job that Davidson needs it to do. Thus, Kornblith rejects both Davidson and Brandom's theories.

A more moderate theory in the social vein is that of Michael Williams (2000). Williams does not want to deny that non-human animals have beliefs, but he does want to claim that they cannot have knowledge. In order to have knowledge you need to be

inducted into a linguistic community and to be able to keep track of epistemic entitlements and commitments. This community provides the normative standard and practice against which beliefs are justified. He distinguishes between grounded belief and responsible belief. A grounded belief is a reliably formed true belief. A responsible belief on the other hand requires meeting the normative standards set by the community (2000: 608). Human infants and non-human animals cannot be held responsible for their beliefs because they cannot be members of the linguistic community.

Kornblith (2002: 91) believes that Williams equivocates on the meaning of 'justification' and so claims that more is needed for knowledge than actually is required. He reconstructs Williams' argument thus:

1. Knowledge requires justified, true belief.
  2. Justification cannot exist apart from membership in a community with certain social practices.
- Therefore,
3. Knowledge cannot exist apart from membership in a community with certain social practices. (Kornblith 2002: 91)

'Justification' in the first premise is most plausibly read as meaning 'grounded', claims Kornblith, whereas in the second premise it is most plausibly read as meaning 'responsible'. It is this equivocation that leads to Williams' claim that we cannot have knowledge apart from as members of a linguistic community.

Kornblith seems to be very close to begging the question here. Because he believes that nothing more than reliable true belief is required for knowledge he reads the first premise as meaning this. He claims that this is the most plausible way of reading the first premise but it is far from clear that this a commonly held intuition. There is a very plausible reading of 'justified' in the first premise, which is only possible if you can give

reasons for your belief. Certainly, if someone cannot provide reasons then we will be doubtful of their claim to know. If we think of paradigmatic examples of human knowledge, then they will almost certainly meet this requirement. Because Kornblith wants to take animal knowledge as paradigmatic then it makes it seem more plausible that the 'justified' means something like 'well-grounded' or 'reliably formed'. Depending on what examples you choose, either reading could be plausible.

Perhaps, rather than claiming that animals cannot have knowledge, Williams would be better off claiming, as Sosa does, that there are two distinct types of knowledge, one that requires only grounded belief – something that adult humans, infant humans and non-human animals can all achieve – and then another type of knowledge that requires grounding plus the ability to give reasons for your belief – something that does exclude infants and non-human animals..

Although Kornblith disagrees with Williams, he notes that something important can be drawn from the discussion (Kornblith 2002: 91-2). Ordinary talk does not seem to draw a distinction between animals and humans with regard to knowledge but it does with regard to justification, he claims. We would not be concerned as to whether or not non-human animals are justified or not in their beliefs but we may ask this question about an adult human. We are generally happy, however, to ascribe knowledge to non-human animals and infants as easily as we do to adult humans. However, if any kind of knowledge equates to 'justified, true belief' then the fact that there are two very different understandings of 'justified' might show us that, although we are happy to ascribe knowledge to both humans and non-humans, that does not mean that we are talking about the same thing in both cases. I may describe both gold and fool's gold as 'gold' but that does not mean that they are the same thing.

Williams believes that epistemic responsibility explains epistemic normativity. The community holds people responsible for their beliefs, they must reach a certain

standard. However, Kornblith thinks that such a community is not necessary to create normativity. He argues (2002:92-3) that rather than creating normativity, the demands of the community reflect normativity that already exists. It seems that there are very good reasons not to drink and drive, quite apart from the social practices that frown upon it. Likewise, it seems that there are perfectly good reasons for wanting to hold true beliefs quite apart from the demands of the linguistic community.

In addition, Kornblith argues that not all epistemic social practice is in fact of value (2002: 95-102). Giving and asking for reasons will not always lead you to the truth. In part it depends on the community that you are a part of. Brandom and Williams both see the practice of giving and asking for reasons as having a 'default and challenge structure' (Kornblith 2002: 96). This means that within a community there are certain things that are accepted without question and others that need justification. In a highly religious community it may be certain scriptural claims that go unchallenged, which could lead to believing a great many untruths about the world. In other societies information is so heavily controlled that accepting default sources will lead you to false beliefs. Such social practices, claims Kornblith, can lead enquiry in unproductive ways and will not, in fact, be beneficial to the people taking part in it.

Kornblith denies that involvement in social practice is necessary for knowledge and so for him one route to claiming that human and non-human knowledge are distinct is closed. However, it's not clear that he has succeeded in this. Whilst there seems to be a *prima facie* case against views such as Brandom's and Davidson's – that such theories exclude animals and infants from having belief – such theories could, in principle, make sense of this by distinguishing between two kinds of belief as Williams does. Kornblith's arguments against Williams also seem to fail. Kornblith claims that it is more plausible to think of knowledge requiring justification in the sense of epistemic grounding or reliability than it is to require epistemic responsibility for knowledge but

he offers no reason why this should be the case. He has not, therefore, shown that there is no distinction to be made between human and non-human knowledge.

### **3.2.2 Knowledge and Reflection**

*a) Is reflection necessary for human knowledge?*

Kornblith also denies that reflection is necessary for knowledge. Firstly, he argues that introspection works very differently from how we think it does, and is actually rather unreliable (Kornblith 2006: Ch.4; 2008; forthcoming). His other arguments are aimed at coherentist and foundationalist accounts of knowledge separately but can be understood as variations on the same argument: requiring reflection for knowledge sets the bar so impossibly high that no one can ever have knowledge. Assuming that we do not want to be led into inevitable scepticism then we should not make introspection a requirement for knowledge.

Kornblith (2002: Ch. 4) begins with Descartes (1641/1996). For Descartes introspection has two roles to play, it can reveal defects in our epistemic position and it can also be used to fix those defects. Using introspection, Descartes realises that he has held false beliefs in the past and that consequently his current beliefs may also be false. The way to fix this, he suggests, is to first withhold all belief and then, through an introspectively accessible act of will, only believe those things that are clearly and distinctly true. This, Kornblith points out, would be an exceedingly slow process and so he suggests an alternative for the 'Impatient Cartesian' (Kornblith 2002: 108). What the Impatient Cartesian should do, he suggests, is to find out which processes of belief formation are faulty and dismiss any beliefs that are formed on that basis. Any beliefs formed by effective belief forming processes can be kept.



It is the program of the 'Impatient Cartesian' that contemporary internalists are pursuing, claims Kornblith. In particular he focuses on the theories of Roderick Chisholm (1989) and Laurence BonJour (1985). Chisholm and BonJour are both internalists but of differing kinds. Chisholm has a foundationalist theory and BonJour a coherentist theory<sup>11</sup>. Chisholm claims you can find out whether a belief is justified by going through a process of Socratic questioning. This process first involves examining your reasons for believing that *p*. If you find that your reason for holding that *p* is that *q* then you need to examine your reason for holding that *q*. You should repeat this process until you come to a proposition that is directly evident, which means that it is recognisable as true through introspection alone. BonJour, on the other hand, thinks that in order to have knowledge of *p*, you need to check that *p* coheres with your other beliefs. This coherence is available to introspection.

However, claims Kornblith, introspection is not the reliable guide to our processes of reasoning that philosophers such as BonJour and Chisholm take it to be. Firstly, he takes two examples from the psychology literature that show that things that we do not think have an effect on our reasoning actually do. The first is described as the 'position effect' (Kornblith 2002: 112). In tests where people were asked to choose the best item from amongst four identical items there was a clear and significant bias to items on the far left. When questioned about this bias, the subjects were sure that the positioning had no effect on their choices and thought the experimenter very odd for even asking such a question. It seems here that the subjects were biased toward the leftmost items but when asked to introspect they were only confirmed in their belief that they were not.

---

<sup>11</sup> Kornblith focuses on BonJour's coherentist theory in *The Structure of Empirical Knowledge*. BonJour has since rejected coherentism in favour of a foundationalist theory. For more on this see BonJour and Sosa (2003) *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, Oxford, Blackwell Publishing.

The other effect that Kornblith discusses is the ‘anchoring effect’ (Kornblith 2002: 113). In this experiment the subjects were given a number and asked whether it was too high or too low for the percentage of African members of the UN. When they were given the number 10 (chosen by a spin of a roulette wheel) the mean estimate as to the real number was 25, and when the number 65 was chosen as a base then the mean estimate was 45. It seems here that people’s estimates were significantly affected by random number that they were given in the first place. The subjects were not aware, however, of the influence that this had. Nor, claims Kornblith, is there any reason to think that asking the subjects to reflect on this would make a difference to the answers that they gave.

There are also cases that seem to show that people are not affected by factors that we assume would affect our belief formation. For example, in one study the experimenters got subjects to read a short story about a drowning baby and asked them to pick out the passages responsible for the emotional impact (Kornblith 2002: 114). They picked the same passages that the experimenters antecedently thought were responsible. However, when the story was given to a second group of subjects with the passages removed, the subjects rated it with the same emotional impact. The passages that were assumed to cause it to be judged emotionally affecting were evidently not the cause of this judgement after all.

A similar effect can happen in perception (Kornblith 2002:115). When on a plane it appears that we can judge the orientation of the plane by sight, we seem to be able to see that the plane is going upwards. However, it seems that this cannot be the source of our judgement, as this happens even when the plane is flying at night and there is nothing to judge its orientation against. The real source of our perception that the plane is ascending or descending is our kinaesthetic experience; we feel the plane ascending but interpret it as a visual experience. The problem is that you will not be

able to answer the question ‘How do you know that the plane is moving in an upwards direction?’ correctly. A natural response to the question would be ‘Because I can see that we are moving upwards’, which is false. However, as Kornblith acknowledges, they are not mistaken in their belief that they are reliable at detecting the movement of the plane. If you have a minimal account of what it takes to be reflectively justified, then this may in fact be enough.

A final example of where introspection can go wrong is that our reasoning suffers from confirmation bias (Kornblith 2002: 116-20). There are two forms that confirmation bias can take. The first is that when testing a hypothesis people only confirming instances of the hypothesis and don’t look at ones that would prove it false. The second is when disconfirming evidence is discovered it is not taken as seriously as the confirming evidence. Such bias leaves us in the position that if we assume that we are reliable (as we undoubtedly do in many cases) then we are likely only to be confirmed in this view. It seems that we also tend to reconstruct explanations for our behaviour based on what we think a rational person *would have* done rather than on the basis of what actually happened. Kornblith suggests that if this is the case then introspection is not the place to look for improving our cognitive functioning. Cognitive psychology is much more likely to give us practical, unbiased advice on how to improve our cognition.

Peter Wason (1960, cited in Kornblith 2002: 117) conducted a study into confirmation bias where subjects were given a sequence consisting of three numbers that they were told conformed to a rule. They were asked to try and discover the rule by suggesting their own sequences of numbers. The experimenters would tell them after each suggestion whether their sequence conformed to the rule or not. Most of the subjects said that they were testing a particular hypothesis and most of them went about it by ‘examining only confirming instances of it’ (Kornblith 2006: 117). Once

they had a large number of confirming instances the subjects claimed that they had discovered the rule. 'Strangely,' suggests Kornblith, 'when subjects were told that they had not discovered the rule, in more than half the cases the next sequence tested was an instance of the very rule that they had just been told was incorrect.' (Kornblith 2006: 117) Kornblith takes it that this shows that subjects only examine confirming instances and disregard disconfirming information.

However, as Bonjour (2006: 325) argues, it is not clear that this case of confirmation bias is actually very problematic for internalist theories of knowledge. Firstly, proposing random sequences of numbers is unlikely to tell you anything about the rule at all. It is highly unlikely that a subject will happen upon a confirming instance of the rule by accident. In the case suggested it seems that looking for confirming instances will be more productive than suggesting sequences at random. This, however, is not the strategy that the authors think that people should be pursuing. What is missing from people's strategy in these cases is that they do not test to see whether there could be a more general rule that would include all cases of their rule. For example, people do not test to see whether the rule is 'add the same number to the previous number' rather than 'add 2 to the previous number' (Bonjour 2006: 325). This does not seem like a case of bias but rather seems to be 'a more general mistake of failing to consider the full range of possible hypotheses' (Bonjour 2006: 326). It is certainly not clear that this demonstrates confirmation bias.

Kornblith seems to overstate his case to some extent. There may be cases in which our introspection goes awry, but that does not show that it is not normally reliable. He has shown that there are some everyday examples in which introspection misleads but he is far from showing that this is the norm. Perhaps we should not be as trusting of introspection as some internalists are but no one claims that it is infallible. There are good reasons for us to have developed introspection that is usually reliable; it is useful

to be able to check our beliefs to see if they were formed well, since in that way we will be more likely to follow the reliably formed beliefs. Kornblith has shown that we can be affected by unexpected factors in making judgements and that we should be vigilant against errors. There are many occasions on which my vision or my hearing has misled me but that does not mean that they are not to be trusted in general.

*b) Is non-reflective human knowledge of the same kind as that of non-reflective animals?*

Even if Kornblith had shown that neither social practice nor reflection was required for adult human knowledge, he still has not shown that animal and human knowledge are of the same kind. Even if he had shown that no kind of meta-cognition was necessary for knowledge, he would still have to do something more to show that human knowledge is the same kind as animal knowledge. If we think about Kornblith's position in terms of Sosa's 'animal' and 'reflective' knowledge, Kornblith is suggesting that the 'animal' knowledge cases are those that we should take to be canonical, rather than cases where reflection is involved. This, claims Bermúdez (2006: 308-9), is based on two assumptions. Firstly, that reflective knowledge simply consists of animal knowledge plus some second-order reflection on that knowledge and, secondly, the 'subtraction assumption' (309). This assumption is that if we take away the reflection from reflective knowledge then we get animal knowledge, of the type that is investigated by cognitive ethologists.

However, notes Bermúdez, the term 'animal' knowledge is doing 'double duty' (309). Animal knowledge means both, very simply, the kind of knowledge that animals have and also to contrast with reflective knowledge, so non-reflective human knowledge is also 'animal' knowledge. Bermúdez goes on to suggest that if there are significant differences between the cognitive processing of non-reflective animals and human beings, then these will not necessarily be the same.

According to Bermúdez, there are two norms that govern belief revision. Firstly, at the non-reflective level beliefs are subject to the ‘norm of coherence’ (310). That is, if there is some conflict between beliefs, or between belief and sensory evidence then this causes dissonance, which is removed by belief revision. Bermúdez gives the example of a rat who comes to believe, through experience, that food will appear when he presses the lever. If the system is then changed so that food appears at random then this will create cognitive dissonance which is resolved by the rat losing the belief that lever-pressing and food appearance are related. At the reflective level, however, our beliefs are also governed by the ‘norm of truth’ (310). In addition to whether our beliefs are coherent we are also aware of evidential concerns. For example, a belief may be rejected just because the evidence available does not support it.

Reflective knowledge is only available to human beings, suggests Bermúdez, because it is only language using creatures that are capable of such reflection. His suggested reason for this is that we can only think about thoughts if they are ‘clothed in language’ (312). This in itself is not necessarily problematic for Kornblith as he views metacognition of this kind as unnecessary for knowledge, it is simply an ability that we have on top of our ability to gather knowledge, it is just an ability to reflect on the knowledge that we have which is, fundamentally, of the same type as that which animals have. However, Bermúdez suggests that even the non-reflective knowledge that humans have must be of a different kind from that of non-linguistic creatures.

If reflective knowledge is simply animal knowledge that has been reflected on then the relations between the beliefs and the evidence available must be present in the ‘animal’ knowledge. However, Bermúdez argues that this cannot be the case. Logic, he argues, requires language, In order to entertain a thought such as “A or B” a creature must be capable of understanding that a certain relation holds between the two thoughts “A” and “B” – at the very least that they are not both false. To be capable of

this the creature must be able to think about the truth values of the thoughts. And, so, ‘no creature that was not capable of thinking about thoughts could have any understanding of truth-functional compound thoughts’ (Bermúdez 2006: 214). This is not just true of logical reasoning that depends on relations between thoughts but also of reasoning that depends on the internal structure of thoughts such as the move from  $Fa$  to  $\exists xFx$ . Thus, in order to be able to enter into such logical reasoning you need to be capable of thinking about thoughts, and for this you need to be a language-using creature, argues Bermúdez. He does not deny that animals are capable of reasoning but rather that it must be of a fundamentally different kind to that which human beings engage in.

I shall not here debate the merits of Bermúdez’s theory. However, the interesting possibility is raised that human and non-human cognition may be fundamentally different. If this were the case then there would be good reason to think that reflective human knowledge cannot just be animal knowledge plus reflection but that even at a non-reflective level human knowledge is of a different kind. Rather than the one distinction that we find in Sosa it seems that there are two possible distinctions. On the one hand, there is the distinction between animal and human knowledge and on the other between reflective and non-reflective knowledge. This leads to four potential types of knowledge: non-reflective animal knowledge, reflective animal knowledge, non-reflective human knowledge and reflective human knowledge. Thus human knowledge may exist in two kinds but these may both be different again from any kind of animal knowledge.

The implication that this has for Kornblith is that if he wants to make cognitive ethology the basis of our theory of knowledge then he needs to do more than show that reflection is not necessary for knowledge. He also needs to show that there is no important distinction to be drawn between animal and non-reflective human

knowledge. Kornblith dismisses Bermúdez's claim arguing that logic is not in fact dependent on language. A creature can be sensitive to logical structure without having complex thoughts regarding truth values. 'The medievals,' he claims

...referred to the disjunctive syllogism as "the syllogism of the dog," and they referred to it in this way because, they claimed, if a dog is chasing another animal down a path and it comes to a fork in the path, it will smell one of the two paths to see if the animal has proceeded down that fork; if the result is negative, it will take the other route without bothering to smell it first. Disjunctive syllogism in action! (Kornblith 2006: 341)

If a dog does indeed behave in this way then it certainly looks to be responsive to the constraints of logic. As there are only two options if the animal did not go down path A then it must have gone down path B. It is not clear if this is enough to respond to Bermúdez, after all, he does not deny that animals are capable of reasoning, just that it must be different from human reason. However, the problem still remains that we have no reason to presume that animal and human knowledge are of the same kind, even if both can be non-reflective in nature.

Kornblith is suggesting a radical move in what is taken to be the raw data for epistemology. Traditionally, human knowledge has been the main focus of investigation. Although narrowly focussing on human knowledge may be unreasonably anthropocentric, it must still play an important part. In order for Kornblith's change in focus to be accepted it seems that he must do more to show that there is no difference in kind between animal and human knowledge. The burden of proof lies with him, and at this point he has not yet succeeded in showing that there is just one kind of knowledge.

It is not even the case that Kornblith need insist that there are *no* differences between human and non-human animal knowledge. Two things can be both of the same natural kind and of different natural kinds. For example, 'tiger', 'whale' and



'mammal' are all natural kinds; a whale and a tiger are therefore of the same natural kind – mammal – whilst simultaneously being of different natural kinds – whale and tiger, respectively. There could be space to allow that animal and human knowledge are, in fact, different, whilst still being of the same natural kind 'knowledge'. This is not something that Kornblith would advocate but it is a possible route for someone who wants to hold on to the distinction between human and non-human knowledge and claim that knowledge is a natural kind.

Also, it is not necessarily the case that the difference between reflective and non-reflective knowledge need be one of kind. Rather, it could be the case that reflective knowledge is just a case of 1<sup>st</sup> order knowledge that  $p$  that is accompanied by 2<sup>nd</sup> order knowledge of how one knows that  $p$ . This seems to simply be a case of having more knowledge rather than having a different *kind* of knowledge. It is unclear which of these options would be preferable and at this point I will not explore the issue further but we shall return to it in Chapter 5. For now, it is simply important to note that it is possible, and reasonable, to draw a distinction between reflective and non-reflective human knowledge, whatever that distinction may, in fact, amount to.

### **3.3 The Consequences of Kornblith's Perspective**

Kornblith argues for several theses. He argues that we should not investigate knowledge conceptually, but rather look to empirical science to give us an answer to the question of what knowledge is. From this he argues that the correct understanding of knowledge is an externalist one and that there is no distinction to be drawn between human and non-human knowledge. This leads to Kornblith downplaying the importance of reflective access in human knowledge.

Even if we accept that an externalist account of knowledge is preferable to an internalist account, this does not mean that reflectively accessible justification is not important. One of the consequences of taking the kind of naturalist perspective that Kornblith advocates is that it emphasises the similarities between animal and human knowledge but ignores their differences. Human beings have reflective capacities that other animals lack and this is bound to have an effect on human knowledge. For example, human knowledge can be undermined simply by uncertainty about one's evidence, a problem that non-reflective creatures do not have to deal with. Knowledge is also of utmost importance in our interactions with other people, and for this reflectively accessible justification is often required. It is expected, for example, that one does not claim to have knowledge unless one has reflectively accessible grounds to back it up. Again, this is a problem that non-reflective animals do not have to face.

None of this is to say that reflection is necessary for knowledge but rather to say that any perspective that downplays these differences will not be giving a full account of knowledge. The role that reflection plays in human knowledge is too important to be ignored. The similarities between animal and human knowledge are also important but we should not let these be the only things that we base our theory of knowledge on. In taking this view, Kornblith's account obscures the important differences that there are in the way that reflective and non-reflective creatures interact with knowledge. To be a full account of knowledge we would need to add an account of these differences to Kornblith's account.

### **3.4 Conclusions**

In order to legitimately use cognitive ethology as the basis for his theory of knowledge, Kornblith needs to show that there is only one kind of knowledge which encompasses

both animal and human knowledge. However, his attempts to dismiss both social theories and reflection as being necessary to knowledge fail. Although reflection may not be completely reliable, something that surely no one would deny, he has not shown that it is so unreliable that we should not trust it. Even if he had succeeded, showing that reflection is not necessary for knowledge would not be enough to show that human knowledge is of the same kind as animal knowledge. It may, therefore, be reasonable to draw some kind of distinction between human and non-human knowledge.

Kornblith has not succeeded in demonstrating that it is *not* legitimate to draw a distinction between human and animal knowledge. However, that is not to say that we should draw such a distinction. If we can keep our theory of knowledge simpler by allowing only one kind of knowledge then that would be certainly advantageous. I believe that we do not need to distinguish between two *kinds* of knowledge, as such, but rather we can draw distinctions elsewhere that will accommodate the important distinctions that there are between human and animal knowledge. There are distinctive advantages and disadvantages to having the capacity for reflection as human beings do. Reflecting on our evidence can reinforce or undermine a belief, for example. Socially we are also often expected to cite our reasons for a particular belief or piece of knowledge. These challenges, and benefits, to being reflective, along with the social environment that we inhabit mean that we reflection and reflectively accessible reasons have a central role to play in human knowledge. This is what we shall explore further in the following chapter.

---

## CHAPTER FOUR

### A Genealogical Approach

---

In contrast to Kornblith, Edward Craig (1990) focuses very much on what knowledge is for human beings. Rather than looking at how the phenomenon of knowledge developed in an evolutionary context, Craig is interested in how and why the concept of knowledge may have developed, an approach that has been described as ‘genealogical’ (Kusch 2009; Pritchard *forthcoming a*). He is not interested in traditional conceptual analysis, however. Rather, he advocates ‘practical explication’ (Craig 1990:8) of our concepts, which involves constructing a new concept of knowledge, making changes to our intuitive concept if necessary. Whether we adapt the concept or not, we should not take any set of conditions we uncover as applicable to all instances of knowledge but rather as defining the prototype for knowledge. Variation from this prototype is perfectly acceptable.

Craig argues that our concept of knowledge originates in our need to tag good informants. We need to be able to identify good and bad informants so that we can get the best information and so a protoknowledge concept would develop in response to that need. The protoknowledge concept has since undergone a process of objectivisation leading to the concept that we are investigating today. That is, it has gone from being about who is a good informant with regard to my information needs here and now to a concept that covers potential informants for potential enquirers. A

knower must be potentially identifiable as such, be capable of communicating the relevant information and be as likely to be right as an enquirer would need him to be.

The focus on enquiry is an important move on Craig's part. When we are investigating knowledge we need to think about it in the context of our everyday lives not just in an abstract theoretical way. However, Craig neglects something which is a vital part of enquiry for both enquirers and informants: reflectively accessible justification. In the final part of this chapter I shall expand on Craig's theory by examining the important role that this plays in our interactions concerning knowledge. I shall also suggest a way in which the apparently contradictory theories developed by Kornblith and Craig might be brought together.

#### **4.1 Craig's Methodology**

Like Kornblith, Craig is dissatisfied with the traditional approach to epistemology of trying to provide 'an explicit intension to fit the intuitive extension' (Craig 1990: 1). There is, he claims, good reason to think that this approach is flawed, not least that this project seems to have been unsuccessful thus far. We have intuitions about the intension as well as the extension of 'knowledge'. For example, one may dismiss an example as not being a case of knowledge on the basis that the belief in question was formed in a lucky fashion. We can understand sceptical problems as arising from this kind of conflict of intuitions, he claims. The conditions for application of the concept of knowledge seem to mark out a much smaller number of cases than our extensional intuitions would indicate (Craig 1990:1). That is, the intuitive extension of knowledge is very broad, including all the things that we generally take ourselves to know. The

intension that seems to be implied by the sceptical argument, however, seems to show that we know almost none of the things that we ordinarily take ourselves to know.

Craig suggests a different starting point for our investigation into knowledge from the traditional approach. Instead of beginning with the ordinary usage of the concept of knowledge he suggests that we should begin with an 'ordinary situation' (1990:2). The word 'knowledge', it seems, has an equivalent in most, if not all, languages. If this is the case then there is good reason to think that it may fulfil an important need or function in the lives of human beings (Craig 1990: 2-3). The method that he suggests is to start with a hypothesis about what role knowledge plays for us and suggest what a concept that filled that role would be like. What Craig is suggesting is, in essence, a state of nature analysis of knowledge; how and why might such a concept have developed in a primitive situation? The investigation is still looking for conditions that govern usage but not actual usage, rather rules that would govern a concept that fitted the hypothetical role that we suggest. Of course, in order for it to still be a concept of *knowledge* neither the intension nor the extension must differ greatly from our everyday usage of the concept.

Craig describes his project as a 'practical explication of knowledge' (Craig 1990: 8). The idea of explication, which he traces to Carnap, is to

construct a new version of [the concept] satisfying certain standards, with the proviso that to count as a new version of that concept it [has] to emerge with many of its principle features intact. (Craig 1990:8)

The hypothetical concept that we construct must stay close to our intuitive understanding of knowledge. If it strays too far from this then it is the hypothesis that is at fault not our intuitions. We are not in the business of inventing totally new concepts but rather of attempting to illuminate the concepts that we already have.

However, our intuitions may be ‘indeterminate or elastic’ (Craig 1990:2) and such an investigation may give us ideas on ways in which they should be stretched, and why and it may also help to uncover the source of such problems.

Searching for a list of necessary and sufficient conditions that cover every case that we take to be knowledge is a lost cause. If we can come up with a list of conditions that take this form then we should instead take them as defining the prototypical case of knowledge (Craig 1990:5). Deviation should be allowed from this prototype because the concept developed in a practical situation. An enquirer needs to have a strategy for finding out about the world and such a strategy should be aimed at normal cases. We should not expect, therefore, that all cases that we are inclined to describe as knowledge will meet all the criteria.

Why should we adopt Craig’s unorthodox approach to epistemology? One reason that Craig suggests is that the traditional method of philosophical investigation seems to have failed thus far to produce a satisfactory account of knowledge, which at least suggests an alternative approach is needed. Also, to succeed in defining knowledge would not be the end of epistemological enquiry. There would still be an open question as to why such a concept developed. The fact that the concept of knowledge is so widespread amongst human communities – Craig claims that anecdotally there is no language that does not have a ‘comfortable and colloquial equivalent’ (Craig 1990: 2) to the English term ‘knowledge’ – suggests that it fulfils some important role in human life. What Craig is suggesting is that we take this as the starting point for our investigation rather than something to be investigated at a later date.

Craig’s main defence seems to be that of the ‘proof of the pudding’. That is, if this methodology ends with

conditions of application showing a close fit to the intuitive extension of 'know', and does fit well with a variety of facts about the 'phenomenology' of the concept, then those who hold it to be mere coincidence may be requested to make a case for their attitude. (Craig 1990: 3-4).

Craig also suggests that in adopting this approach we can both strengthen and widen the basis of the theory of knowledge. Instead of looking at only intuitive ideas about the extension of the term 'knowledge' we can also take into account facts about the success and failure of various analyses, controversies, reactions to scepticism and so on. This gives us a much broader base of data to work from. This also strengthens the basis for our theory because these data seem to be more akin to an 'undocored 'given'' (Craig 1990:7) than what is normally taken into account in analysing knowledge. The result of which, hopefully, will be that we can develop a less theory-laden and theory-driven account of knowledge.

A further defence of Craig's methodology can be found in Kusch (2009). Kusch expands upon Craig's idea that the concept of knowledge 'still bear[s] certain marks of its origin' (Craig 1990: 95). That is, there are signs that we can detect in our current concept of knowledge that indicate how it might have come about. Firstly, Kusch cites linguistic evidence that the concept of knowledge has roots in the practice of testifying. For a start there is the common use of testimonial metaphor in our talk of knowledge and learning. For example, 'The tyre tracks *tell* us that the assassin arrived by car after it rained.', 'I couldn't *believe* my eyes.', 'Jones is very well *informed* about this insect: he has studied it for years.' (Kusch 2009: 81). Linguistically, it seems, we often treat non-humans as testifiers. We may also say that someone is well informed even when there has been no informant involved. Secondly, there is the linguistic fact that that 'know' can take interrogative constructions as complements whereas believe, for example, cannot. E.g. 'I know who she is.' vs. 'I believe who she is.'; 'She knows where to go.'



vs. 'She believes where to go.' Testimony and knowledge appear to be linked here in a way that testimony and belief are not. Kusch also argues that attributions of knowledge have the function of honouring the knower and that markers of this can be seen in the concept of knowledge that we have now.

## **4.2 Knowledge as Tagging Good Informants**

### **4.2.1 What is enquiry?**

For human beings there are two ways in which we can gather information about the world. We can get information directly from the environment by using our senses or we can gather information from other human beings. Of course, human beings are not the only animals capable of sharing information; monkeys can warn the rest of their troop that a tiger is approaching by making a specific call and honey bees can communicate complex directions to good food sites through the movements of a 'waggle dance'. What is unique to humans is that we can engage in the process of enquiry using people as informants. We can ask and answer questions that are intended to assist enquiry; we rely on testimony with regard to places we have never been or experiences we have never had; we can enquire of experts in areas that we do not know much about; and so on.

By enquiry I mean nothing more than the search for an answer to a question. This could range from the simple 'Is it raining outside?' to complex investigations into the genetic make-up of elephants or the chemical make-up of the atmosphere on Mars. Enquiry is part of our daily lives. We want to know many things: at what time the bus leaves, where our keys are, whether the meeting is today or tomorrow, how much the concert ticket will cost and so on. Different enquiries will demand different methods.

The question of whether or not it is raining outside can be easily resolved by looking out of the window. On some other matters, however, the best method will be to ask someone else. If I do not know what time the bus leaves then one simple way of finding out is to ask someone who knows. Some enquiries will require a mixture of both. There are, therefore, two roles that one can play with regard to an enquiry. One can be the enquirer, who is seeking the answer to some question. Or you could be an informant, who can give an answer to the question. In many cases the enquirer and the informant will be one and the same person such as when you ask yourself ‘Now where did I put my keys?’

Enquiry and the transmission of information are vital parts of human social interaction. They are also distinctly human abilities. Given that we have such capabilities, it seems that there arises a need to be able to distinguish between those who testify well and those who don’t. Furthermore, given that we are presumably interested in *good* testimony there is a need to encourage accurate testimony and discourage bad testimony. In the next section we shall examine Craig’s argument that knowledge is a concept that evolved to tag good informants.

#### **4.2.2 What is needed to be a good informant?**

*a) True belief*

Human beings, like all animals, need to be able to find out information about their environment in order to survive. As such we have certain ‘on-board’ sources of information – eyes, ears and so on – that help us to find out about the world around us. However, we are not always in the best position to find out about a certain issue and so it would be advantageous to be able to access the resources of other human

beings. As Craig suggests, Fred, who is up the tree, may well be in a better position to spot the approaching tiger than the person on the ground (Craig 1990: 11). It seems, therefore, that we have good reason to want to be able to assess and mark people as good or bad people to get information from.

Getting information from an informant involves more than is required to get information directly from the environment. We need to be able to assess the informant's competence and decide whether they are reliable, trustworthy and so on. This seems to be a different process to the one in which we assess the information of our senses. Although in some cases you may need to check that your eyes or ears are working properly or that the environment is normal you do not need to worry about whether the apple in front of you is only pretending to be red, or whether it is a reliable apple. Craig's hypothesis is that the ancestor of knowledge would have developed to assist this process. If we can tag our potential informants as good or bad then it will save us having to go through these checks every time.

What does it take for someone to be a good informant? A first requirement is that they will tell us the truth as to whether  $p$ . This will primarily involve the informant holding a true belief as to whether  $p$ . That is, 'either  $p$  and he believes that  $p$  or not- $p$  and he believes that not- $p$ ' (Craig 1990:12). This comes very close to the traditional truth requirement for knowledge but couched in terms of a disjunctive condition. This is because Craig is thinking in terms of enquiry. When we make an enquiry into a certain matter it is an open question as to whether- $p$ . We will only have successfully completed this enquiry if we have a decisive answer either way. Traditional theory of knowledge, on the other hand, is couched in terms of individual propositions and whether they are known or not.

In addition to having the correct belief as to whether *p*, a good informant must also be able to bring an enquirer to the correct belief (Craig 1990: 13). A subject who will not or cannot tell us the truth will not be a good informant even if they do in fact have a true belief about the matter at hand. If you were to choose a compulsive liar as your informant then you would not be brought to a correct belief about the matter at hand, even if he has a true belief. Must a good informant be confident about the matter in question, then? Obviously, if he has a belief then he is more likely to state it with confidence. If he appears to be hesitant about his answer then you may not believe him.

However, an informant's confidence in a proposition is not the only reason that we might trust an informant's testimony. If I know independently that Mr Truetemp's utterances about the temperature are reliable then I can be brought to a true belief about the temperature even if he himself is unsure about his beliefs. Provided that he states what he feels the temperature is, it is unimportant whether he has a belief about it or not. It looks, therefore, as if belief on the part of the informant will not be necessary for them to be a good informant (Craig 1990: 14-5).

Craig draws a distinction between informants and sources of information, however, which is important here (Craig 1990: §V). A source of information is a state of affairs that tells us about the world, such as the reading on a thermometer. An informant, on the other hand, is an agent who responds to our question as to whether-*p*. The concept of knowledge, claims Craig, is tied to informants not sources of information (Craig 1990:34). We do not usually claim that the thermometer knows what the temperature is or that a tree knows how old it is, even though we can come to know these things by examining the thermometer or the tree. Getting correct information from a source of information is simply about being able to interpret the evidence in front of you.

Human beings can be both sources of information and informants. For example, we can discover whether John has been knocked unconscious by testing his reactions, or lack thereof, without having him tell us that he has been knocked out. Here we use John as a source of information. He would count as an informant only if we were engaging with him as an agent. The distinction between informants and sources of information helps to make it clear why we may want to add a belief condition to the conditions for protoknowledge. An enquirer may be able to find out whether  $p$  from someone who doesn't believe that  $p$ , but it's not clear that they are using them as an informant.

There are advantages to using informants rather than just sources of information (Craig 1990:36). Firstly they are convenient. You need no special skills, beyond the ability to communicate, to gather information from an informant. Sources of information, however, may require special skills to interpret the information that you find. An experienced tracker may be able to tell you which animals have passed by here recently by looking at the tracks on the forest floor but an ordinary hiker would not.

Secondly, informants can often be more helpful than a mere source of information. They are able to understand the purpose behind your enquiry and so can often give you additional information that will be of assistance to you (Craig 1990: 36). For example, if you ask where the nearest Starbucks is an informant is capable of looking at what motivates your request – the desire for a large cup of coffee – and will be able to tell you that although there are no Starbucks in the immediate vicinity, there is a very nice coffee shop just around the corner. An informant can empathise with the reasons for you asking a particular question and give you the information that is of most use to you in addition to, or even instead of, the information that you actually asked for.

Of course, there are also disadvantages to using informants. Informants may deliberately deceive in a way that sources of information are unlikely to. That is, they may give every indication of being a good informant with regard to a certain issue but still deliberately mislead you. Informants can also be unreliable through no fault of their own, we are not infallible after all. These facts, however, just underline the need for us to be able to tag which informants are reliable and which are not.

*b) Identifiability: Condition X*

So far we have seen that the conditions for being a good informant seem to come close to the traditional conditions for knowledge. Firstly, a good informant must be right as to whether p. Secondly, he must believe that p. Like traditional epistemologists, Craig doesn't think that this is sufficient. It is not enough that someone *would* be a good informant if we could never identify them as such. If nobody is aware of Mr Truetemp's ability, including Mr Truetemp himself, then he will not be a good informant as to the temperature. If there is no way of telling that he would be a good informant then we might not even think to ask him and would not be likely to believe him if he sounded unsure about the answer. A good informant, therefore, must be identifiable as a good informant. To tell if someone is reliable with regard to things that I already know is relatively easy, I ask them questions and test their answers against my own knowledge. However, it is in the nature of enquiry that the enquirer does not know the answer to the question that they are asking. If you knew, you wouldn't be asking. A good informant needs to have some feature, independent of the fact that they give correct answers, that identifies them as someone who is likely to be right as to whether p.

Craig examines various contemporary epistemological accounts to see if they can shed light on what this condition might be. Firstly, he looks at the modal account put forward by Robert Nozick (1981, cited in Craig 1990: III). As discussed in Chapter 2 of this thesis (72-3), Nozick claims that to be knowledge beliefs need to track the truth. S's true belief that p counts as knowledge if: (a) in the nearest possible world in world in which p is false, S does not believe that p, and (b) in nearby possible worlds in which p is true, S continues to believe that p.

Craig (1990: 21-23) thinks that there is something to be said for accuracy across a range of possible worlds but thinks that the conditions should be different from those that Nozick suggests. There is a good reason to want an informant to be accurate across a range of possible worlds because we cannot precisely identify which world we are in and we want our informant to be correct whichever one it happens to be. However, to do this job we would need to fix the range of worlds according to the interests of the enquirer. An enquirer should only be interested in 'open' possibilities; ones that he cannot rule out for himself. If you want to know the time of the next bus then it you do not need to take into account the world in which your informant is wearing a different coloured shirt. However, Nozick's conditionals do not exclude such worlds as all close possible worlds are relevant. The other problem with Nozick's conditionals for Craig's project is that we do not have a way of detecting accuracy across possible worlds and so it is of no use in detecting reliable informants. Although being accurate across possible worlds may coincide with an identifiable property, it cannot be what Craig is looking for.

Craig argues that a reliabilist condition will not do the job either. As we saw earlier, in §2.3.2, simple process reliabilism claims that knowledge is true belief formed by a reliable method. However, we are not equipped to tell whether or not someone has

formed their beliefs in a reliable way, or at least, it is not something that is immediately obvious to an enquirer. If you ask someone whether or not the train is due at 5pm, you generally have no way of telling whether that person formed their belief in a reliable manner. Perhaps they looked at the current timetable, or perhaps they looked at an old timetable which was very different to the current one. An enquirer can't easily identify someone as having formed their beliefs in a reliable way, at least not without a lot of investigation. Again, having reliably formed beliefs might coincide with the property that Craig is looking for but it cannot do the job that he'd like (Craig 1990: 30-1)

Instead of there being one property that makes informants detectable, Craig suggests that the condition should be understood to cover '*any* detectable property which has been found to correlate closely with holding true belief as to whether *p*' (Craig 1990: 25) and labels this *X*. Although this may correlate closely with beliefs being formed in a reliable way or to fulfilling various counterfactual conditionals these things are not enough. We need something that identifies a subject as an accurate source as to whether *p*. Craig is happy to allow any kind of feature to fill this role, arguing that being a taxi driver is a detectable property that correlates well with reliably giving correct answers about locations and so we can identify a taxi driver as a good informant in this regard simply because he is a taxi driver.

Although Craig claims that the connection between *X* and an informant being right about *p* must be reliable, this should be distinguished from the claims of reliabilism. According to reliabilism, the thing which must be reliable is the method by which the subject comes to believe that *p*. What Craig claims, on the other hand, is that *X* must reliably indicate that the subject will be a good informant about the matter in hand. That is, that in most cases where someone has property *X*, they will be a good informant as to whether *p*. This is irrespective of whether they formed their belief in a



reliable manner or not. Although these two things will often coincide, they are not the same thing.

There are various ways in which the 'reliable method' of reliabilism can be expanded. Does it need to be reliable in producing a true belief in exactly the situation that obtains; in situations like the one obtaining; in all likely circumstances; or, in all possible circumstances? A similar question arises with regard to the good informant; do we need them to be a good informant just about this particular question or about many (or all) related questions? It seems that there may be good reason to prefer the latter. For a start, one way of telling that someone will be a good informant as to whether- $p$  may be that they have been good informants about related matters in the past. For many propositions the truth value of  $p$  will not vary over time and so we will not ask exactly the same question repeatedly. Thus, we may have only an indirect indication of whether a particular person will be a good informant as to whether- $p$ .

This does seem to be a way in which we often judge whether someone will be a good informant or not. If you want information regarding the result of the latest Hibs-Hearts derby, for example, you will ask the football fan who has always been accurate before. This gives you a high chance of getting the right answer, even if you haven't asked about that particular game or even those teams before. If, on the other hand, someone had been highly unreliable in this respect previously then the enquirer has good reason to doubt that he can accurately answer the question on this occasion. Previous success, or lack thereof, is a good indicator to an enquirer as to whether or not a person is likely to be a good informant in this case.

*c) As likely to be right as needed*

So far, we have seen that a good informant should have a true belief as to whether  $p$  and that he should be recognisable to me as a good informant. Craig also wants to add

that a good informant should be as likely to be right as the enquiry needs and channels of communication should be open between the informant and enquirer. The second of these is pretty straightforward – if the enquirer cannot communicate with the informant then she will not be able to get the information that she requires. The first needs a bit more explanation. A good informant, as discussed above, needs to be correct across a range of possible worlds because we do not know which world we are in. However, he does not need to be correct in those worlds which are not ‘open’ possibilities. Nor does he need to be correct in worlds which are highly improbable.

How many worlds are ruled out will depend on various factors, such as: the urgency of the enquiry; the relative benefits and costs of being right versus being wrong; the enquirer’s personal attitude to risk; and, the background circumstances to the enquiry (Craig 1990: 86). If the train is in 5 minutes then an enquirer might accept a less reliable informant as to where the station is than if he has 30 minutes to find it. A less reliable informant will also do if there is very high potential benefit in getting the right answer but a low cost in getting the wrong answer. In addition, some people are just naturally more risk averse than others and so will be more exacting in their requirements for informants than those who are more inclined to take risks. Finally, these all need to be weighed up against the background circumstances. If a lot is riding on the answer to a particular question then the enquirer will be more careful over his choice of informants, however much of a risk-taker he might generally be.

#### **4.2.4 The process of objectivisation**

There are several cases which appear problematic for Craig’s theory (Craig 1990:17, 82). These are cases in which we would intuitively identify a subject as knowing that  $p$  but in which they would not make a good informant as to whether- $p$ . Firstly, there is

the case of a subject who has information relevant to an enquiry but will not or cannot reveal it. For example, the case of Luigi, who knows where the body is buried but will not tell. Secondly, is the case where a subject has no credibility as an informant. For example, Matilda has repeatedly 'cried wolf' and so no one believes her claim that the house is burning down even when it is true. A final case is where it is impossible to identify someone as a good informant on a particular matter. The case of Fred, the secretly studious milkman, for example. Fred was once a reserve on Mastermind but didn't tell anyone. His specialist subject was obscure and unrelated to dairy delivery. As such, he knows the answer to an obscure question but no one would know to ask him. If such cases are cases of knowledge then Craig needs to explain how we got from the concept of protoknowledge, which seems to exclude them, to the concept of knowledge, which includes them.

Craig suggests that protoknowledge underwent a process he describes as 'objectivisation' (1990: §X). Initially, creatures were only looking for things to satisfy their immediate desires in the here and now. However, there is pressure to move away from this subjective stance to something more objective. Craig takes the example of the concept of a chair. A 'protochair' is perhaps just anything in the vicinity that you can sit on right now but there are pressures to move away from this highly subjective viewpoint. You may be also interested in something to sit on in the future, or in things that you could sit on somewhere that you've never been before, or in things that other people, of a different size or shape to you, could sit on. In such cases you may need to rely on other people knowing what is suitable to sit on so that they can tell you.

So, we move away from a context-based, highly subjective concept to a concept that is abstracted away from any particular person, time or place. From something that is useful for me to sit on right now to the idea of something that is suitable for sitting

on *in general*. A chair may be of no use to anyone right now as it is stranded in the middle of the ocean or because it is a chair for a doll's house. However, these things are still chairs because they could be sat on in principle, although, in the second case, only by a very small person. The fact that right now they cannot be used as chairs does not mean that they are not chairs.

A similar thing has happened to protoknowledge, suggests Craig. In the first instance a good informant would be someone who has a true belief about *p* and meets the following criteria:

1. He is accessible to me here and now.
2. He should be recognisable by me as someone likely to be right about *p*.
3. He should be as likely to be right about *p* as my concerns require.
4. Channels of communication between him and me should be open. (Craig 1990: 85)

We can see that people may fail to be good informants even though they have a true belief by failing to meet one of these conditions. Some may fail to be a good informant simply by not being here when you need the information. Others, the secretive milkman and Matilda included, would fail the second test. You cannot possibly be aware that he would have the right answer to your question as he has never given anyone the slightest reason to think that he would. Due to her habit of telling tales, Matilda, will not be recognised as a good informant, either, and so she also does not fulfil this condition.

The strength of the third condition will vary with the situation. In cases where the matter is of great importance one may reject someone who would be perfectly adequate if the matter were trivial. In a life-or-death situation I should not risk getting the answer to my question from someone who seems even a little unlikely to get it

right. The final condition can fail to be met in various ways. Perhaps the enquirer and informant do not speak the same language, or, like Luigi, the potential informant refuses to reveal their information.

In any society there are likely to be many individuals facing the same general problem of how to settle various whether-p questions. If the members of that society are responsive to the needs of others then there is a pressure towards objectivisation, claims Craig (1990: 87-8). We may need to recommend informants or to ask for recommendations and so we need to start to think outside of our own immediate needs. You are no longer thinking solely of informants who are of use to you here and now but rather of subjects who would make good informants for people in general, irrespective of circumstance.

There will sometimes be good informants that you are not equipped to recognise and so you will have to rely on others to give you this information. Likewise, you will be able to recognise informants that others would not be able to. In particular, you are generally in a better position than other people to identify whether you yourself would be a good informant on a particular matter. Thus, the requirement that someone be recognisable as an informant becomes much broader. It does not have to be something that a particular enquirer can detect but rather something that *someone* could detect. Our milkman may not seem a good informant to us but it would surely be possible for someone, himself, if no one else, to identify him as a good informant on the subject in question.

The requirement that the informant be as likely to be right as the enquirer needs him to be also has to change. Now, an informant must be as likely to be right as *anyone* could need him to be. In different enquiries and in different circumstances the level of reliability needed will vary and an objectivised concept has to take this into account.

This means that condition (3) needs to be strengthened to require someone with a very high likelihood of being right (Craig 1990: 91).

This potentially pushes the requirement up so high that no informant could possibly meet it, leading to scepticism (Craig 1990: 91-2, §XI). In lottery cases a person's belief that she will not win is almost certain to be correct. In a lottery of 100 tickets with only one winner, the belief that I have not won will be correct 99% of the time. We would not, however, take someone as a good informant as to the result of the lottery just on this basis. It seems that we could have a lottery with a very large number of tickets and this would still apply, which seems to imply that the accuracy of our informants should be pushed up until they are 100% reliable.

Craig's answer to this problem is that lottery cases are special cases. Usually we do not call on our informants to be 100% reliable. Our concept of knowledge is one which has developed in a real life situation and should be aimed on the kind of cases that are the norm. Although lottery cases are normal in that they happen in everyday cases they are unusual in that they appear to demand a much greater accuracy than normal. In usual, practical situations this level is very far from being required. We should not isolate our philosophical investigations from real life (c.f. Kaplan 1991, 2000, discussed in this thesis §1.3)

The final condition also changes with objectivisation. There may be cases in which you could not communicate with the informant but in which someone else could. Gareth can find out information from an Arabic speaker when I cannot, due to my ignorance of his language. Although Luigi will not tell you where the body is buried that does not mean that he would not tell Carlo. As we move to an objective concept, the conditions change allowing for cases where a subject may not be a good informant for me but may be a good informant in a more generalised sense.

The process of objectivisation also answers the ‘cart before the horse’ objection (Craig 1990: §X): it seems natural to say that someone is a good informant as to whether  $p$  only if they know that  $p$  in the first place. To dismiss this criticism Craig looks back to his comparison with chair. When investigating the concept we may first notice that people like to sit down and that the concept of a chair may have developed in response to this. Oh no, claims the detractor, there are things that are chairs that are of no use to people who want to sit down. There could be a chair in the middle of the Sahara or a chair on a blazing bonfire. These things are chairs nonetheless. You need to first tell us what a chair is, and then we will find out why they are good for sitting on.

This response seems adequate when applied to the concept of a chair and Craig thinks that we can respond in the same way with regard to knowledge. If the cases are indeed parallel, then it seems that a detractor would need to find additional reason to reject Craig’s prioritisation of the concept of a good informant. Through objectivisation, the concept does move away from its origins but the mark of those origins can still be seen (See discussion of Kusch 2009, above). The concept of knowledge is not identical with the concept of a good informant but rather is an objectivised version of a good informant.

However, it’s not clear that the objection cannot be pressed further here. Craig’s hypothesis is that we can gather information about the world through our senses but that it is also advantageous to be able to gather information from other people’s sensory equipment. So, we have a need to identify those people who can give us good information about our environment. Because of this, we develop a concept ‘knowledge’ that we apply to people if they are good at telling us the things that we need to know; someone knows whether  $p$  if they believe truly and can potentially inform others that  $p$ . This seems to be a strange way to put things as surely you have

the information available to yourself even if you can't pass it on. Reliably gathering information from your senses is surely prior to being able to pass that information on.

However, Craig would claim that we would have no use for the concept of knowledge except in a social environment. The concept only develops because we need to identify who can give us good information. The idea is that the concept of knowledge developed in a social context to fulfil a particular need that human beings have. This does not mean that it cannot now be applied to those who are in some way removed from social interaction. The objectivised concept of knowledge says that a person must in principle be identifiable as an informant if they are to count as a knower. A person can, however, identify themselves as an informant for their own purposes even if they were incapable of communicating with anyone else. Robinson Crusoe can know that the tide is coming in even though he cannot communicate this to anyone. He is a good informant with regard to his own enquiry.

### **4.3 Building on Craig: The Role of Reflective Access in Enquiry**

Craig's theory gives a good account of how the concept of knowledge may have originated. However, even though he concentrates on the importance of enquiry, he does not place enough emphasis on the role of reflection both for enquirers and informants. In this section we shall explore the importance of reflectively accessible evidence for both enquirers and informants and build upon Craig's theory to give a fuller understanding of the concept of knowledge in a social context.



### **4.3.1 The Importance of Reflective Access for Informants**

Craig claims that externalism fits better with his account of what it takes to be a good informant and therefore his theory of knowledge (1990: 68). A good informant does not need to have reflectively accessible reasons for his belief in order to pass on the information effectively. Provided that there is some reason for the enquirer to think that the informant is reliable, there doesn't seem to be any need for the informant to also be aware of it.

This is only obviously the case, however, if we look at things from the third person perspective i.e. that of the enquirer. Frequently, a person will have to assess whether or not he himself is a good informant, either in order to play a part in their own investigation or to declare themselves an informant for the purposes of someone else's enquiry. Internalist ideas are enticing in this context, claims Craig (1990: 64-5), but this is a mistake arising from the perspective that we have had to take towards our own beliefs. When assessing your own beliefs you have to play the role of both enquirer and informant and so you have to assess your ability as an informant. The only way that you have of assessing this ability is by looking inside yourself for markers that you are reliable as to the question of whether p.

However, this reflection is only necessary when examining oneself as a potential informant; it is not necessary for all informants to do this when giving information. The first person perspective is not primary, it is just that in certain circumstances we need to adopt this perspective. The football fan doesn't need to know himself that he is a good informant but we need to be able to identify him as reliable. We may know that he's always been reliable in the past or we may take his wearing of a football scarf and singing of football-related songs as an indication that he will be reliable with

regards to the outcome of today's game. He need not have any reason himself to think that he is a good informant. Craig's theory, therefore, is externalist in nature.

It is not clear, however, that the first person perspective is at all unusual. A very common way of conducting an enquiry is to ask people if they know the answer to your question. A responsible enquirer would not ask just anyone – if you need directions to a local tourist spot, it makes more sense to ask the ice-cream vendor than the lost-looking tourist. However, there is often no way of telling whether a person will know the answer to the specific question that you are asking other than by asking them whether they know. In fact we often couch enquiries in terms of 'Do you know...?' questions. The initial question does not state our enquiry directly but rather asks the potential informant whether they know the information that you seek. One does not ask 'How do you get to the castle?' but rather 'Do you know how to get to the castle?' or 'Can you tell me how to get to the castle?' Of course, the correct affirmative response is not just 'Yes'. It is normally understood that the question does two jobs; to enquire as to whether you know and also to ask for the information. Even if you have good reason to think that the person that you have asked will be a reliable informant – when you ask at the tourist information office, for example – you may still very well formulate the question in terms of the informant's knowledge. The phenomenology of testimony, in English at least, seems to be that we ask people to declare or at least confirm themselves as informants in addition to giving the information that is required.

To be a responsible informant you should have some idea of how you came to know. If you do not have some reason to think that you are correct then you should not put yourself forward as an informant. We can see this in the Gricean maxim that you should not assert knowledge unless you have adequate grounds for your knowledge (Grice 1975: 46). We may be able to use someone as an informant when

they do not have reflectively accessible justification but we would need to have some other reason to think that their answer would be reliable. From the 1<sup>st</sup> person perspective, however, we only have reflectively accessible justification available. In normal circumstances, if an informant is not sure how they came by a particular piece of information then they will not claim it as knowledge but rather express it in qualified terms. When faced with difficult quiz questions, for example, someone who does not have reflectively accessible evidence will couch things in terms of ‘I think the answer is...’ or ‘I’m not sure, but I’d say ...’ and so on. This would be the same whether you know the answer, have a true belief about it, or a false one for that matter. When we lack second order knowledge of how we came to have a certain piece of information we hedge what we say to make sure that the claim is not misunderstood as a claim to knowledge.

The fact that people on the whole are unwilling to make a knowledge claim unless they have some reflective backing for their belief can be taken as evidence that, as informants, they want their testimony to be accurate. We do not want to be understood as claiming knowledge unless we have reason to think that we have knowledge. Kusch claims that knowledge, in addition to having the role of tagging good informants, also has the role of bestowing honour on good informants and shame on bad ones. If an informant is tagged as a (proto-) knower then that is a form of praise. Reynolds (2002) argues that the primary purpose of classifying people as knowers is to improve testimony. To identify someone as a knower is to praise them, which, therefore, gives people incentive to testify accurately. If we have an interest in good testimony then it seems reasonable to suggest that a way of encouraging good testimony would emerge. The concept knowledge, therefore would play a dual role. From the point of view of

the enquirer, it would have the role of tagging informants and from the point of view of the informant it would have the role of encouraging them to testify well.

As the proto-concept of knowledge was developing presumably an awareness would develop amongst informants as to what was expected of them. A common method of enquiry is to ask people whether they know the answer to a particular question of whether or not *p*. In these cases we would expect our informants to go through this reflective process and declare themselves as good informants, or not, and may enquire into the reflective reasons that they have. In fact, this is often the case, especially if the answer to our question is of particular importance. We very often have to rely on people to declare themselves as informants and so it is odd that Craig dismisses internalism in the way that he does. It is true that there are other ways of us identifying reliable informants, they may have given us good information in the past or they could be identifiable by their job as one likely to give good information such as teachers and people who work in a tourist information office. However, in many cases the only way I can find out if someone is a good informant is by asking them. The first person perspective is much more important and central to testimony than Craig claims.

### **4.3.2 The Importance of Reflective Access for enquirers**

It is often stated that 'knowledge closes enquiry' (See, for example, Kvanvig 2003: 171, 2009: 344). Once we come to know the answer to the question that we are investigating we no longer need to continue enquiring. In this section I will argue that it is only reflective, 2<sup>nd</sup> order knowledge that can successfully and responsibly close enquiry; that is, knowledge accompanied by reflective justification.

There are two ways in which enquiry may be closed. In the first instance, the enquiry comes to an end because the enquirer just no longer feels the need to

investigate. They have come to a belief about the matter in question. In this case the enquiry is closed because the enquirer has alleviated the 'irritation of doubt' (Peirce 1877/1992: 114). We can also ask, however, what it takes for an enquiry to be closed responsibly and successfully. Firstly, for an enquiry to be successful, we should at least be brought to a true belief about the answer to the question posed. A false belief would not be a successful end to enquiry. Nor would it be successful if the enquirer came across the truth but did not come to believe it. If the enquirer has no belief about the answer to the question being asked then he would have no epistemic reason to stop his enquiry.

To successfully close enquiry, then, one needs to have a true belief about the answer to the question in hand. Is a true belief sufficient or do we need knowledge? To a large extent this will depend on what kind of knowledge is in play. If an enquirer comes to have brute externalist knowledge that *p* then it seems that he is not yet at the point of being able to halt enquiry responsibly. He may know that *p* but this knowledge will be indistinguishable from a mere belief that *p*. There is no reason for him to think that his belief is true beyond the fact that he believes it and so there is no reason for him to cease his enquiry.

For example, if Mr Truetemp wants to know what the current air temperature is then he can come to a reliably formed true belief just by thinking about it as, unbeknownst to him, he has a chip in his brain which accurately records the temperature. However, the fact that he has formed a true belief in a reliable manner does not give him reason to stop his investigation because from his point of view his belief could as easily be false as true. He does not have reason to cease his investigation until he has checked his belief. Once he has checked his belief against a reliable thermometer, he has reason to stop his enquiry. It seems, therefore that reflectively

accessible justification is needed in order to close enquiry. It is not enough to have knowledge without reflective backing because you would have no reason to stop enquiry at this point.

In fact, reflectively accessible justification for your belief seems to be sufficient to close enquiry. Of course, in some cases an enquirer will possess such justification but have a false belief. For example, I may think that I have found out that the bus leaves out at 4.15pm by looking at the timetable. If the timetable is out of date, however, and the bus now leaves at 4.10pm then I may have closed my enquiry too early. To successfully and responsibly close enquiry, however, you need to have knowledge that is accompanied by reflectively accessible justification. The level of justification may vary depending on the enquiry. For example, if it is a very important matter we may want stronger justification than on something trivial, or if it is urgent we may accept a lower level of justification than we would if we had all the time in the world. However, in all cases an enquirer needs some understanding of how he or she came by his or her belief or some reason to believe that it is true. In most enquiries it will be clear to the enquirer how he or she came by his or her belief – ‘I saw that the keys were on the table.’; ‘I read that the bus is due at 4.15.’; ‘John told me that the party is tomorrow night.’ and so on. If the enquirer does not have even this level of backing for their belief then they it is clear that they cannot responsibly close their enquiry.

Having closed enquiry we want it to remain closed so that we do not have to start investigating the matter again. A satisfactory end to enquiry then, is stable knowledge. A piece of purely externalist knowledge does not have this stability. If an enquirer does not have any reflectively accessible reason to believe the proposition then it seems that any counter-evidence should shake his belief. A responsible enquirer would reopen his or her enquiry if this were the case. So, it is preferable to have a true belief that is

backed up by reflectively accessible evidence for the belief as an enquirer is less likely to lose his belief once he has gained it.

In order to close enquiry we need stable belief which can only come about through having internalist justification. We also want this belief to be non-accidental in nature – which externalist knowledge gives us. So in order to successfully close enquiry we need to have knowledge which is backed-up by reflectively accessible justification. I shall examine the idea of stable knowledge further in Chapter 5.

#### **4.4 Contrast with Kornblith's Approach**

Both Craig and Kornblith take unorthodox approaches to epistemology. Unlike Kornblith, however, Craig does not want to abandon conceptual investigation completely. Although his method is unusual he is still investigating the same thing as the traditional epistemologists – the concept of knowledge. In fact, he dismisses a Kornblithian-type theory saying that he doesn't think that knowledge is a 'given phenomenon' and our concept developed in order to allow us to speak about this 'common and important stuff' (1990: 3). Craig does not put forward an argument for this position claiming only that the 'proof of the pudding' (1990: 3) will be whether his methodology does result in giving conditions for the correct application of 'knowledge' that fit closely with our intuitive understanding of what knowledge is. If it does, he suggests, then the burden of proof will be on those who suggest that this is a mere coincidence to show it to be so.

I shall not attempt a direct response to Kornblith's position on Craig's behalf but rather suggest that the two theories are not at odds to the extent that they initially appear. Despite both ostensibly being interested in the question of what knowledge is

the two philosophers are focusing on very different aspects of knowledge. Craig is primarily interested in how the concept of knowledge as a social phenomenon developed – why do we need the concept of knowledge and what role does it play for us? On the other hand, Kornblith is interested in why the natural phenomenon of knowledge developed – why do animals, including human beings, need such a thing as knowledge? What evolutionary benefit does it have? These two different sets of interests are focusing on the same thing – knowledge – but there is no reason to think that they are straightforwardly contradicting each other.

One way to reconcile the two positions might be to claim that there are two different things that fall under the description of ‘knowledge’. On the one hand is a phenomenon called knowledge which evolved because it assists in the survival of those animals that have it. In addition to that is the concept of knowledge as it is used in human social interaction. Kornblith claims that he has no interest in the concept of knowledge and is only investigating the natural phenomenon. Craig, on the other hand is not interested in the natural phenomenon. He is not involved in the project of ‘Evolutionary Epistemology’ which ‘looks at our cognitive faculties as adaptive responses to changing circumstances and changing needs for information.’ (Craig 1990: 10).

Craig focuses on the role that knowledge plays in a human social context whereas Kornblith is primarily concerned with knowledge in the context of non-human animals in the natural world, although he thinks that this account extends to humans as well. As such, it is far from clear that they are looking at the same thing. The two investigations could, perhaps, even be seen as continuous. First, the phenomenon knowledge develops – animals need reliably-formed, true beliefs to survive. Then, once human beings have developed the capacity for reflection and complex communication, a new



need arises – to be able to sort good from bad informants – necessitating the need for a concept to classify people. As we saw above, Craig claims that having reliably formed true beliefs may correlate well with a third condition for knowledge but that it cannot do the job that human beings require – the identification of good informants. We could perhaps take this to show that there may be something that is common to both human and animal knowledge – reliable, true belief – but given the particular needs of human beings in a social environment that requires us to sort and tag good and bad informants we need to have a concept that reflects those needs. If these two ways of thinking are reconcilable in this way then perhaps taking insights from both will give us a better understanding of knowledge. I shall discuss this further in the following chapter.

#### **4.5 Conclusion**

Craig's methodology – to examine the possible origins of the concept of knowledge – is a useful one and gives us a different perspective on the old question of what knowledge is. We can see a strong link between knowledge and the process of enquiry. However, Craig's theory does not put enough emphasis on the importance of reflectively accessible justification in the process of enquiry. In this chapter, I have sketched out a way of building on Craig's insight which restores this importance. I have also suggested a way in which the apparently contradictory methodologies of Kornblith and Craig could be reconciled. The difference between them is primarily one of focus. In Chapter 5 I shall go on to examine how the thoughts developed in this chapter can tell us something about why we value knowledge.

---

## CHAPTER FIVE

### Value Problems

---

Several value problems have come to light in recent epistemology. The first is in fact an old problem dating from and named after Plato's *Meno*. It simply asks 'Why is knowledge more valuable than mere true belief?' The two states appear to have the same practical value but we seem to prefer knowledge. The *Meno* problem uncovers a problem with certain types of theory such as reliabilism. It seems that such theories cannot give an answer to the *Meno* problem because they are based on a view whereby truth is the only epistemic good. If this is the case then any condition on knowledge that is truth-conducive will not add value to mere true belief. Two further value problems have also been identified; the secondary and tertiary value problems. The secondary value problem asks why knowledge is more valuable than any proper subset of its parts and the tertiary value problem asks why knowledge is distinctively valuable, why it has a different *kind* of value to anything that falls short of knowledge.

This chapter will examine the challenge posed by these different value problems and suggest an answer. I will argue that knowledge is not in fact more valuable than mere true belief but that knowledge accompanied by justification is more valuable. I will argue that there is a difference between mere true belief and internally justified knowledge; if our knowledge is accompanied by justification then it is more stable and also better suited to being used to pass on knowledge. Most of our knowledge is accompanied by this kind of

value and so we can explain why we seem to think that knowledge is always more valuable than mere true belief even if this is not actually the case. Making this value distinction between knowledge and internalistically justified knowledge also has the added bonus of being able to explain why an externalist answer to scepticism is unsatisfactory.

## **5.1 Value Problems**

There are various value problems that crop up in the contemporary literature and that we shall examine here. The Meno problem asks simply, why is knowledge more valuable than mere true belief. The swamping problem is a particular challenge that certain theories face in answering the Meno problem. Those that understand truth as being the only thing of value in the epistemic domain cannot explain why knowledge is more valuable than mere true belief. The secondary value problem asks why knowledge is more valuable than any proper subset of its parts, and the tertiary value problem asks why knowledge has a distinctive value.

### **5.1.1 The Meno, or primary value problem**

Why do we value knowledge? An initial and relatively straightforward answer is that it is very useful to have knowledge. On an everyday basis my knowledge helps me to fulfil my desires. If I want chocolate then my knowledge that the corner shop sells chocolate, and is open right now, will help me fulfil my desire for chocolate. I would find it a lot harder to fulfil my desire for chocolate if I was in a strange place and didn't know where the nearest chocolate vendor was.

The fact that knowledge is useful answers our initial question about knowledge but one may think that there are further questions that need to be answered. Is knowledge

distinctively valuable, for example? That is, do we value knowledge more than other epistemic states, such as true belief? If so, why is this the case? In the *Meno*, Socrates claims that we value knowledge more than mere true opinion. However, this seems to be problematic. If I have a true belief about the way to get Larissa and I act upon that belief, then I will arrive in Larissa just as well as if I had knowledge of the way to Larissa. There seems to be no practical difference between the two.

We are therefore faced with a *prima facie* problem. Knowledge is more valuable than true belief but there seems to be no practical difference between the two. It is not clear, therefore where the value difference lies. We seem to value knowledge more than mere true belief. Given the choice between the two most people would be inclined to choose knowledge every time. If I want to get to Larissa, for example, I would rather know the way than just truly believe it. However, there seems to be no practical difference between the two. If I have a true belief about how to get to Larissa then it seems that I can get there just as well as if I *knew* how to get there. There seems to be a conflict between two intuitions about knowledge: 1) knowledge is more valuable than mere true belief, and 2) knowledge and true belief have the same practical value.

The most straightforward answers to the question would involve denying one of the intuitions behind the problem; either knowledge is not more valuable than mere true belief or they do not have the same practical value. A third option would be to maintain that both (1) and (2) are true but to claim that the value difference is not to do with practical value. For example, knowledge may have an intrinsic value that mere true belief does not.

### **5.1.2 Swamping Problem**

Theories such as reliabilism seem to have a particular problem in answering the Meno problem; the swamping problem. It is perhaps best explained by reference to an analogy (originally due to Zagzebski 2003:14-15): Jill likes good coffee and she has a good cup of coffee in her hand. As it happens this coffee was produced by a reliable coffee-machine. In most cases, the coffee produced will be high quality. Jack also likes good coffee and has an equally good cup of coffee in his hand. This coffee, however, was produced by an unreliable machine which often produces terrible coffee but on this occasion just happens to have produced a good cup. Is Jill's coffee worth more than Jack's just because it is reliably produced? The answer to this seems to be no. If both cups of coffee are equally tasty, then it makes no difference how they were produced. Both Jack and Jill are primarily interested in having a good cup of coffee and that is exactly what they both have, the reliability of the production methods is irrelevant. That is not to say that reliability is not a good quality in a coffee machine. If Jack and Jill are interested in regularly having good cups of coffee, then it makes sense for them to buy the most reliable machine their budget can stretch to. However, the value that reliability confers onto the coffee machine doesn't transfer to the coffee that is produced by that machine. Given that Jill has a good cup of coffee in her hand, it doesn't matter that it was produced reliably. The value that it has in being the product of a reliable process is 'swamped' by the value of the good coffee itself.

The case with reliably-produced true belief is analogous. When it comes to knowledge we are primarily interested in getting to the truth. Once you have a truth in your possession it is irrelevant whether it was formed reliably or not. The value of reliability is swamped by the value of the truth. Again, it is valuable to have reliable belief-forming mechanisms because we are generally interested in getting to the truth, but this value is not conferred onto the beliefs themselves. If what we are interested in is

forming true beliefs then once we have true beliefs in our possession it doesn't matter how we came by them.

Although the swamping problem was originally aimed at reliabilism, it is in fact problematic for any view that has a truth-conducive condition on knowledge. One way to understand the swamping problem is to see it arising from a particular view of what is valuable from the epistemic perspective. This view, described variously as 'epistemic value monism' (Zagzebski 2004), 'epistemic truth monism' (Sosa 2007:70-2) and 'epistemic value T-monism' (Pritchard *forthcoming a*: 14, and *passim*) claims that truth is the primary, only, thing that is valuable in the epistemic domain. Other valuable properties such as justification or coherence are valuable because they are instrumentally valuable with regard to getting to the truth. A belief which is justified or coheres with one's other beliefs, or both, is more likely to be true. These things are only valuable in so far as they are truth-conducive. Beliefs that are justified or coherent are more likely to be true.

Epistemic value T-monism leads to a problem for certain types of theory of knowledge. If knowledge consists in a true belief plus a third truth conducive condition, such as reliability, then it appears that we cannot explain why knowledge is more valuable than true belief. A true belief does not seem to gain any value from having been formed by a reliable process because the value of reliability is swamped by the value of truth. As we saw earlier, a good cup of coffee gains nothing from being produced by a reliable coffee machine, it is just as good however it came about<sup>12</sup>.

It is, of course, valuable to have reliably produced beliefs in general. We want to have cognitive systems that reliably produce truths, otherwise we would not have much luck getting around our environments and getting what we want. However, this value is not conferred on the belief because if the belief is true then it doesn't matter how it got to be

---

<sup>12</sup> In fact, if you know that the machine is unreliable then you might value it more because you were expecting the coffee to be bad.

that way, it is the truth that is valuable. These truth-conducive processes are valuable because they help us to get to the truth, their value derives from their relation to the thing that is of primary value – the truth.

### 5.1.3 Secondary Value Problem

The primary value problem and the related swamping problem are not the only value problems that knowledge faces. Jonathan Kvanvig (2003) also discerns a secondary value problem (See also, Pritchard 2007). If we want to explain the distinctive value that knowledge has then we need to explain not just why it is more valuable than true belief but also why it is more valuable than any other proper subset of its parts. Knowledge is almost never analysed as simply true belief<sup>13</sup>. Rather, it is usually analysed as consisting of true belief plus at least one other condition, for example, justified true belief, reliably formed true belief or non-lucky true belief. In some cases such as subjectively justified true belief a fourth condition is needed in order to avoid Gettier examples. If this theory of knowledge was correct and we could show that justification adds value to true belief then we would have an answer to the Meno problem. This would not, however, explain the *distinctive* value of knowledge as knowledge is not just justified true belief. In addition to the primary value problem Kvanvig discerns a secondary value problem. If we want to explain why it is *knowledge* in particular that we value then we need to explain not only why knowledge is more valuable than mere true belief but also than any other subset of its parts. My answer to this, however, is that knowledge does not, in fact, have a distinctive value. What we take to be the distinctive value of knowledge is in fact the combination of true belief and justification. In this section I shall argue first that it is not

---

<sup>13</sup> One notable exception is Crispin Sartwell (1991; 1992). Sartwell argues that a) there is no pre-theoretic commitment to the distinction between knowledge and true belief, and b) that knowledge couldn't be anything but true belief. However, Sartwell's position has been almost universally rejected, primarily due to its counter-intuitive nature and also on the basis of counter-examples. See, for example, Lycan (1993).

clear that we need an answer to the secondary value problem even if knowledge is distinctively valuable and secondly, why I think that knowledge is not distinctively valuable but that justification is the important factor in the value distinction.

The secondary value problem asks why knowledge is more valuable than any proper subset of its parts. For example, why is knowledge more valuable than justified true belief, or than non-Gettiered true belief? This is perhaps not as straightforwardly a problem as the value difference between mere true belief and knowledge. However, Kvanvig claims, this is a problem if we want to claim that it is knowledge in particular that we value. If knowledge is no more valuable than justified true belief, for example, and knowledge is more than justified true belief – perhaps it is non-Gettiered, justified true belief, for example – then we have not explained the distinctive value that knowledge has. If it is knowledge that we value then we need to be able to explain why it is knowledge and not some lesser state that is of importance to us.

#### 5.1.4 Tertiary Value Problem

Pritchard (2007: §2; 2008: §2; 2009:4, *forthcoming a*) has also discerned a tertiary value problem. In addition to asking why knowledge is more valuable than any subset of its parts, we can also ask why it is knowledge in particular that we value. That is, does knowledge have a distinctive value?<sup>14</sup> There are good reasons to think that knowledge is distinctively valuable. Firstly, we need an explanation for why it is *knowledge* that has been the focus of epistemological investigation rather than some other state. The other reason is that we seem to treat knowledge differently from other states, as something that is ‘precious’. For example, we berate students who only think of their degree in terms of

---

<sup>14</sup> Kvanvig (2003) appears to lean towards the view that knowledge is distinctively valuable (e.g. xiv, xv-xvi, 116) but he has since claimed that any suggestion of this was for ‘stylistic effect only’ (Kvanvig 2009: 343).



potential future monetary gain. The gaining of knowledge should be at least as important, if not more so, than the career advantage a degree may give you. This implies that knowledge has a value independently of purely practical considerations. The importance of the tertiary problem has been underplayed, claims Pritchard (*forthcoming* a: 8), as if we can answer this problem then we also have the answer to the primary and secondary value problems.

## **5.2 Responses to the value problems**

### **5.2.1 Virtue epistemology**

Virtue epistemology looks as though it may be able to solve these value problems. Achievements seem to be finally valuable and so that can explain why knowledge is finally valuable. However, the going is not quite as good for virtue epistemology as it first seems. As Pritchard argues, the idea that knowledge is an achievement doesn't seem to capture what knowledge is; cognitive achievement is neither necessary nor sufficient for knowledge. Knowledge may generally go hand-in-hand with cognitive achievement but that cannot account for the value of knowledge.

Virtue epistemology of a certain kind claims that a subject has knowledge if they form a true belief because of an epistemic ability (e.g. Greco 1999, 2003, 2007; Sosa 1991, 2007; Zagzebski 1996, 1999). Sosa, for example, claims that in order to be knowledge a belief must be apt. That is, your belief must be true because it was formed through the exercise of a cognitive ability. According to this account, knowledge is best understood as a cognitive achievement. If you hit the target because you are a skilled marksman and it is properly creditable to your ability then you have achieved something. In the same way, if you come to a true belief that there is a target in front of you because of your excellent visual skills, then the achievement is creditable to you.

This understanding of knowledge can give us an answer to the Meno problem. We seem to value things that are the product of achievement more than those that are the product of luck. We would value a shot that hits the bulls eye differently if it were due to the archer's skill than if it were the product of freak winds. Hitting the target because of her skill is more valuable than hitting it through luck. If the same is true of all achievements then we can explain how knowledge is more valuable than true belief; knowledge is more valuable because it is an achievement in a way that mere true belief is not. Achievements are finally valuable. That is they are valuable for their own sake but because of how they are formed. Pritchard (*forthcoming a*: 30) suggests that it should be no surprise that we value achievements as they 'constitute the exercise of one's agency on the world'. Our lives are better for including these achievements even compared to a life where our goals are regularly met but without any achievement on our part.

The virtue epistemologist also appears to be able to answer the tertiary value problem. If knowledge is an achievement then it has a different kind of value from those states that are not achievements. A mere true belief is not an achievement, nor is just any reliably-formed or safe true belief. It is only if the success is creditable to the agent does it count as an achievement. For example, an agent 'Temp' forms his belief about the temperature according to the reading on a thermometer on the wall. However, the thermometer is actually broken and fluctuates within a certain range. Nevertheless, Temp always forms a true belief about the temperature because whenever he looks at the temperature a hidden second person changes the temperature of the room to match the reading. His belief is reliable because it will always be right<sup>15</sup>. However, it is not a cognitive achievement. This, therefore, can explain the distinctive value of knowledge. It

---

<sup>15</sup> This example is due to Pritchard (*forthcoming a*: 68)

has a different kind of value to states that fall short of knowledge because it is an achievement and achievements are finally valuable.

Pritchard (*forthcoming a*) argues that although initially plausible virtue epistemology cannot give an adequate answer to the value problems. He argues that if knowledge could be equated with cognitive achievement then it would answer the problems but in fact cognitive achievement is neither necessary nor sufficient for knowledge.

Firstly, one could have achievement without knowledge. Although virtue epistemology rules out the kind of luck involved in Gettier cases, it does not rule out the kind of environmental luck that is at issue in barn façade cases. Henry is looking at an actual barn and forms a true belief because of his cognitive ability yet due to the preponderance of fake barns in the area, he does not have knowledge<sup>16</sup>. This seems to show that one can have a belief that is an achievement in the relevant sense yet still fails to be knowledge. The virtue epistemologists' criteria do not rule out environmental luck.

Secondly, it seems that one can have knowledge without cognitive achievement. In particular, it is not clear that there is any cognitive achievement in certain types of testimonial case. For example, in Jennifer Lackey's case (Lackey 2007: 352), a tourist, Morris, arrives in Chicago and wants to find out how to get to the Sears tower<sup>17</sup>. In order to do so he asks the first person he sees for directions. The person he asks is local and gives him good directions to the tower. Intuitively we would say that Morris comes to know in this case. However, it doesn't seem that the truth of the belief can be accredited to Morris's cognitive abilities. Whilst it is undeniable that he deserves some credit – he doesn't ask just anyone for directions but chooses someone who looks likely to know – the fact that he is correct is not primarily creditable to his abilities. Rather, it is the person

---

<sup>16</sup> Sosa (2007) appears to bite the bullet on this issue and accept that Henry does know in barn façade cases. He lays out an example – the kaleidoscope case (Sosa 2007: 31-33) – that seems to be structurally similar and at a later point acknowledges that barn façade cases may fall into the same category (Sosa 2007: 96n1).

<sup>17</sup> For further discussion of this case see also Pritchard (*forthcoming*: §2.6).

that he asks who has to be awarded the main part of the credit. It is their cognitive ability that has paid off with a true belief, not Morris's.

Thus, it seems that cognitive achievement is neither necessary nor sufficient for knowledge and if this is the case then we have not found an adequate solution to the value problems. However, not all is lost. If knowledge generally goes in hand with cognitive achievement, then, we can have an answer to the swamping problem (Pritchard *forthcoming a*: 64). In order to answer the swamping problem, all that is needed is that knowledge is *sometimes* finally valuable. This shows that epistemic value T-monism is false, undermining the basis for the problem. That knowledge is sometimes finally valuable, because it is sometimes an achievement, shows that truth is not the only thing of value in the epistemic domain.

### 5.2.2 Pritchard's Anti-luck Virtue Epistemology

Pritchard (*forthcoming a*) argues for an anti-luck version of virtue epistemology. That is, a theory that combines a virtue-theoretic ability condition with an anti-luck safety condition. Knowledge is, according to this account,

safe belief that arises out of the reliable cognitive traits that make up one's cognitive character, such that one's cognitive success is to a significant degree creditable to one's cognitive character. (Pritchard *forthcoming a*: 51)

The ability condition that is used here has an advantage over that used in Sosa's account as it does not require that a belief be primarily creditable to the agent but only significantly so. This avoids problems such as the Jenny case discussed above. In fact, the combination of the two conditions means that this theory is well equipped to deal with many of the problem cases that are discussed in contemporary epistemology. It also has a significant advantage when it comes to answering the value problems.

Pritchard argues that whilst it is not possible to give an answer to the tertiary value problem, it is possible to give a diagnostic answer as to why such an answer is not possible (Pritchard *forthcoming* a: 62-4). As we saw above, cognitive achievements are distinctively valuable but not all knowledge counts as a cognitive achievement, as we saw in the Jenny case. However, argues Pritchard, the paradigm cases of knowledge are those in which there is a cognitive achievement of some sort and so are finally valuable. There is, he argues, a type of understanding that usually accompanies knowledge, which is finally valuable. Although knowledge is not, in itself, finally valuable we can see how we might come to think that it is, as paradigm cases of knowledge have such value in virtue of the accompanying understanding.

Pritchard uses a Craig-style genealogical approach to explain why knowledge has the bipartite structure that he claims it does (*forthcoming* a: 60-1). There is, he claims, an ambiguity in the idea of a reliable informant. In one sense an informant is reliable if she has a reliable cognitive ability and in another sense she is reliable if she is someone that an enquirer can rely on. One can be reliable in the first sense without being reliable in the second. For example, your informant might be a good barn-detector but you think (falsely) that she is in barn façade county. She is reliable in that she is forming true beliefs but you cannot rely on her because of you cannot defeat the (misleading) defeater (Pritchard *forthcoming* a: 61). An agent could also be reliable in that you can rely on her but not be reliable. For this Pritchard gives the example of an agent with poor maths skills who is unknowingly being helped by a benign wizard such that he gets all the answers right (*forthcoming* a:61). Pritchard argues that it is because of this ambiguity in what it means to be a reliable informant that knowledge has the structure it does. This Craig-style account can also help to give an answer to the primary and secondary value problems. If knowledge developed to fulfil a particular need then we can see that it is an epistemic standing that has particular value to us. Knowledge is not finally valuable but it

has more instrumental value than those standings that fall short of knowledge (Pritchard *forthcoming a*: 63).

Pritchard (*forthcoming a*:64) also suggests that his theory can give an answer of sorts to the swamping problem. One could either take the position that knowledge is not in fact more epistemically valuable than mere true belief but it is more practically valuable. Or one could take the view that since knowledge is often, although not always, a cognitive achievement and therefore finally valuable, we have an epistemic pluralist account.

### **5.3 The Importance of Internal Justification**

Although Pritchard's account is effective in answering the value problem, it does not do enough to take into account the value that justification adds to knowledge. In this section I shall outline an alternative account of the apparent value of knowledge, which demonstrates that whilst knowledge that is not accompanied by reflectively accessible justification is not necessarily more valuable than mere true belief, knowledge that is justified in this way is more valuable.

It doesn't seem that there is any way to hold on to both the intuition that knowledge is more valuable than true belief and also hold an externalist theory of knowledge. This, I believe is because the intuition that knowledge is more valuable than true belief is tied to a more internalist understanding of knowledge. The case for externalism seems pretty clear and no form of externalism excludes the possibility of internalist knowledge as well as externalist knowledge or the possibility that knowledge is sometimes accompanied by reflectively accessible justification. In the previous chapter, we discussed the importance of internalist justification when it comes to social interaction and the transmission of knowledge. Here we shall see why internalist justification makes a difference on an individual level as well.

It doesn't look like an externalist account of knowledge can answer the value problems but if we add justification into the picture then we can solve them. Firstly, the primary value problem. This just asks why knowledge is more valuable than true belief. From a practical perspective it looks like knowledge will not be any more valuable than mere true belief; you will get to Larissa just as effectively if you follow a mere true belief as if you follow something you know. When we think in terms of internally justified knowledge, however, we can see a real practical difference.

Let's compare several cases. Firstly, the classic road to Larissa case.

**Case 1:** Janet is trying to get to Larissa and is a stranger to the area. She has a mere true belief about the way to Larissa; let's say it was formed in an unsafe manner and it is not internally justified either. When she comes to a fork in the road she believes, truly, that she should take the left hand fork. She follows her belief and gets to Larissa.

As we saw above, this appears to show that knowledge is no more practically valuable than mere true belief. You will arrive at your destination either way. However, this is not as obviously the case when conflicting evidence is introduced into the picture:

**Case 2:** John also has a mere true belief about how to get to Larissa and is a stranger to the area. When he comes to the fork in the road, he believes truly that he should take the left-hand fork. However, there is a signpost that says 'Larissa' pointing down the right-hand fork. The signpost appears to be genuine and doesn't seem to have been tampered with in anyway. In fact the signpost has been put there by the people of a neighbouring town who are fed up with Larissa getting the trade of every passing epistemologist. What should John's response to the sign post be? He has a true belief that he should follow the left-hand fork, but the evidence that he is presented with seems to imply that he should go down the right-hand fork. It seems likely that he would change his route to go down the right-hand path, leading him wrong.

The presence of the signpost will, it seems, undermine John's belief. When faced with evidence that contradicts our belief then the rational thing to do is to reassess our belief.

It seems that the only sensible conclusion that John can come to is to take the right-hand fork. The only evidence that he is aware of is the sign post, as he has no internally accessible reason for his true belief. Although there is objective reason for John to take the left-hand fork – not only is it the correct way to go but he has a true belief that this is the case – he has no subjective reason to act on this belief. The only thing that he has to go on is the false evidence of the sign. As reflective creatures, we do not simply follow our beliefs blindly. We try to act in a way that will help us to fulfil our desires and the way that we do this is by assessing our internally accessible evidence. We not necessarily aware of the objective reasons, we are only aware of the subjective reasons.

How then, does the subject that has knowledge fair?

**Case 3:** Peter is in an apparently better situation than either Janet or John. He knows, in a purely externalist way, that the road to Larissa is down the left-hand fork. He formed this belief in a reliable way, however he has no internal justification for his belief. What does Peter do when faced with the misleading sign? It seems that in spite of knowing that he needs to follow the left-hand road, he is in no better position than John. He has no reason to think that his belief is true, never mind that it is knowledge. It seems that he too would be likely to take the wrong path.

As we can see having knowledge is of no benefit in this kind of case, the problem being that he has no more evidence for his belief than John. Knowledge is of no more use than true belief when it comes to these kinds of situation. However, if we add internal justification, then we have a state that is more valuable than mere true belief.

**Case 4:** Jane is a local of the Larissa area and has taken the road many times. Not only does she know that the way to Larissa lies down the left-hand fork but this knowledge is accompanied by strong internal justification. When she reaches the fork in the road she takes note of the road sign and surmises not that the road to Larissa has moved but that someone must have put the sign there for some other reason. She is not put off her course and successfully gets to Larissa.



Because Jane has internal reassurance that her belief is correct, she is able to dismiss the false evidence of the sign, something that neither John nor Peter could do. In this way, having justification is valuable and knowledge that is accompanied by justification is more valuable than knowledge *simpliciter*.

This harks back to the original answer to the value problem given by Plato. He claims that knowledge is more valuable than mere true belief because knowledge is ‘tethered’ to the truth. Like the statues of Daedalus, which were so life-like that they had to be tied down to prevent them from running away, belief which is not tethered is likely to slip away. If our belief is not in some way tethered, then we might lose it, in the face of countervailing evidence, for example. Mere true belief has nothing to tie it down and so if we are faced with something that appears to contradict it then we should at least withhold judgement, if not reject the belief outright. Internal justification adds stability and durability to a belief and this is something that is in general valuable.

As Craig (1990: 7) notes, however, stability does not always add value to belief. For example, I may need to know what time it is right now so that I don’t miss my train. However, it will be of no value to me tomorrow or in three weeks time to know what time it was right now. We do not need to claim, however, that we should keep every belief that we form for an indefinite period. Obviously, some beliefs will be of use in the future whereas others will only be of use in the short-term. The important point about stability that is, for as long as they are useful, it is beneficial for beliefs to be stable. If the belief is of use then you do not want it to disappear with the smallest challenge. It may not be of any use for me to know what time it is now at some point in the future. If it is, however, then it will be beneficial if that belief is stable. If I am asked to account for my whereabouts on a certain occasion, then it is better if I have some justification; ‘I know that it was 7.15pm because ‘The Archers’ had just finished’ is more useful than ‘I think it

was around 7.15', even if you have a true belief in the second case. So, although we do not need every belief to last indefinitely, stability adds value for as long as the belief is useful. In some cases that may just be a few minutes and in others it may be years, but however long the period, it is better that a belief be stable than not.

Unlike Plato, I do not want to draw the line between knowledge and mere true belief. Rather, the relevant difference is between internally justified knowledge and brute externalist knowledge. I think that the case for externalism about knowledge is good (See Ch. 2) but I think that this kind of knowledge cannot answer the value problems set out above. Instead we need to add justification to knowledge in order to do this. Justification will also add value to true belief and so justified true belief will be more valuable than brute externalist knowledge in some cases. For example:

**Case 5:** Topsy has a justified true belief that the way to Larissa is down the left-hand fork in the road. However, because her belief was formed in an unreliable way, this belief does not amount to knowledge. When faced with the misleading road sign Topsy is in a similar position to Jane. She can discount the evidence of the sign because of the internally accessible evidence that she has.

As we can see from this case, it is the justification that is doing all of the work. The reason for this is simple, having internal justification for your belief is the only indication that you as the subject have that your belief is true. When it comes to making decisions about what to believe or how to act in a given situation then that is all that you have to go on.

We do not necessarily need to draw a distinction between two types of knowledge as Sosa does (e.g. Sosa 1997: 422; 2007: Ch2, Ch5).<sup>18</sup> That is, there is no necessity to claim that knowledge accompanied by reflective justification is a different *kind* of knowledge. It

---

<sup>18</sup> See Chapter 2 for further discussion of this.

is enough to claim that on the one hand there is knowledge *simpliciter*, knowledge of a brute externalist kind. On the other hand, there is knowledge *plus*, that is, knowledge plus internalist justification. It is still, at heart, the same kind of knowledge but it is accompanied by extra justification. It's not clear that very much hangs on whether there are two different kinds of knowledge or not but either way it is not necessary to claim that the difference is one of kind.

The reason that internally accessible justification is important is that we are creatures that are capable of reflection. A cat, who is not capable of this, does not face the same kind of questions that a human being does on a daily basis. Not only can misleading evidence undermine a true belief, the mere lack of evidence can do the same damage. If I cannot remember why I think that the correct road lies to the left then I will be left in a quandary. In such a case, once the seeds of doubt have been sown – which can be as simple as yourself or someone else asking ‘Are you sure?’ – you often become less sure, on some occasions abandoning your belief entirely. Just reflecting can sometimes undermine a true belief. Because we are capable of such deliberations, the role of reflectively accessible evidence is key. Even if we can have knowledge in the absence of such reflection, and I think that we can, such knowledge is of less value than knowledge that is accompanied by internal justification.

If we take internally justified knowledge then we can give an answer to each of the value problems. Firstly, as we saw in the examples above, there does seem to be a real, practical difference between mere true belief and justified knowledge (or justified true belief). The justification adds stability, which in itself is valuable. It is beneficial if we can use our knowledge in the future as well as at the immediate moment and so stability adds value. Justification also adds value in that it makes our knowledge more useful in terms of sharing our knowledge. As discussed in the previous chapters, there is a norm of assertion that says that we should not assert something unless we have some way of

backing it up. Thus, if you have knowledge but no justification for proposition *p* then you should not assert that *p*. Justification plays an important role in the way that we socially transmit information and so adds value in this way as well.

This approach can also give a response to the swamping problem. Justification is valuable not just for its connection to truth; that is, epistemic value T-monism is false. It is not solely important to gain a true belief but also to have a justified belief. Although our main aim may be to find out the truth, that is not the only thing that is valuable even within the epistemic domain. True belief that is fleeting may not be fit for purpose. Often in commencing an investigation we hope to gain lasting knowledge. Justification is not valuable simply because of any connection to the truth that it may have but also because it stabilises our belief, making it long-lasting and because of the role that justification plays in our social interactions. We need justification not just because it is an indicator that a belief is true but also because it can be used to back up our claims when acting as an informant.

The mistake in the value problem is not epistemic value T-monism but rather that it isolates a single belief from the wider context. That is, it is not that there is something other than truth that is valuable in the epistemic domain but rather that we value truth for more than just the fact that we hold it in our hands right now. We don't just want truth right here and now but true belief that will persist and not be easily lost. We also value truth that can be passed on easily. For both of these things we need internalist justification. When it comes to a cup of coffee, I may only be interested coffee that tastes good right here and now. If I am looking for some ground coffee, or coffee beans, however, I may have extra criteria. These may all be derived from the value that I place on the taste of the coffee but the swamping problem does not take into account this wider context. I am not interested in ground coffee that will only be good today before losing its flavour, as I may want coffee on another occasion. I also want coffee that I can

share with friends, that fits within the norms of my coffee-drinking society. Coffee that would not be acceptable socially – that has passed through the digestive system of a civet, perhaps – is not fit for purpose. So too with knowledge, if it is not long-lasting or socially useful then it is no more valuable than true belief. Knowledge, even justified knowledge, is not necessarily more valuable than true belief if you take a snapshot, but this is not a good way to assess whether one thing is more valuable than another. To do this we need to take in the wider context of knowledge and if we do this we will discover that internally justified true belief is more valuable than mere true belief, or even than brute externalist knowledge.

We can also respond to the secondary problem. Knowledge is not more valuable than any proper subset of its parts. Knowledge and reliably-produced true belief are equally valuable. However, justified knowledge is more valuable than any subset of the parts of knowledge because justification adds value. That is to say, it is not knowledge *simpliciter* that is more valuable than any of the sub-parts of knowledge but knowledge *plus*.

Of course, if this understanding of knowledge is correct then we need to explain why we have the intuition that it is *knowledge* that is valuable rather than that it is knowledge plus justification that is valuable. What I think is going on here is that we have conflicting intuitions about knowledge. This is why there often seems to be an impasse between internalists and externalists. I think it is plausible that we have both internalist and externalist intuitions and the different sides on this debate emphasise different sets of intuitions. Whilst we seem to have some intuitions that support an externalist account of knowledge – the importance of belief being connected to the world; ascribing knowledge to infants and animals; the importance of forming beliefs reliably; and so on – this isn't the whole story. Reflectively accessible justification plays a central role in the way in which we think about knowledge, talk about knowledge and use knowledge.

These internalist intuitions could be what lie behind our intuition that knowledge is more valuable than mere true belief. After all, knowledge is, in the normal run of things, accompanied by such justification even if there are cases where we lack it. For most pieces of knowledge we can give some account of how we came to know it or evidence that backs it up. It is a rare case in which we can give no explanation of how we came by a particular piece of knowledge or why one might think that it is true<sup>19</sup>. The fact that most of our knowledge is accompanied by justification may obscure the fact that it is not knowledge itself that is valuable but only knowledge that is also internally justified.

The tertiary problem still remains but it is not clear that this is a particularly problematic issue. The initial intuitions that we started with was simply that knowledge is valuable and from this we moved to the claim that it is more valuable than true belief. There is nothing in these intuitions that means we need to show that knowledge has a different kind of value to true belief. All it claims is that knowledge has more value, and this could be of the same or another kind. It is not clear that we need demand any more. Whilst it seems to be true that we value knowledge over other epistemic states, there is no reason to think that it is not just at the top of a sliding scale. If I can't have knowledge about a particular matter, then justified, true belief is a good second and failing even that, true belief is acceptable. If I wanted to know which road goes to Larissa and I am

---

<sup>19</sup> There are cases in which we apparently cannot bring forth such evidence. Millar refers to such knowledge as 'detached standing knowledge' (Millar *forthcoming*: 202-7). This is knowledge where we have lost touch with the evidence that we once had. Good examples of this are cases of general knowledge, such as my knowledge that the capital of Malaysia is Kuala Lumpur. Such cases definitely count as knowledge, but they don't obviously count as internally justified knowledge. Millar spells out his answer to the problem in terms of an ability to recognise what one remembers (as opposed to what one appears to remember). In this it is like the perceptual recognitional abilities that he claims play a central role in knowledge. We could take something like Millar's line and claim that in such cases you do have reflectively accessible justification for the proposition; the justification is that you know that you are reliable on such matters, it is the type of thing that you know. For example, I know that my belief that Accra is the capital of Ghana is questionable as my grasp of African geography is shaky at best. That the capital of France is Paris, however, is the kind of thing that I know, having a better grasp on geography closer to home. I have justification for my knowledge that Paris is the capital of France, even if I cannot cite any one source that lead me to believe it. I know that I have a reasonable grasp of European capital cities and so am justified in believing that I am correct in this case.

prevented in some way from having knowledge of this fact, then a mere true belief will do. Just because we value knowledge *more* does not mean that we value it differently.

#### **5.4 An Explanation of Metaepistemological Scepticism**

If we understand the value of knowledge in the way sketched out above, that is, not as value of knowledge but rather as value of justified knowledge, then we have an explanation for why the externalist answer to scepticism is unsatisfactory. Knowledge accompanied by justification is the only valuable type of knowledge and this is precisely what we cannot have in answer to scepticism.

The sceptical argument appears to show that we cannot have knowledge because we cannot rule out the possibility that we are not BIVs. However, the externalist counters that it is not necessary for knowledge that we rule out sceptical hypotheses. As long as our beliefs are formed in the right kind of way, as long as we are connected to the facts in the right way, we can have knowledge. So, my belief that I have two hands counts as knowledge because it is connected to the fact that I have two hands in the right way. The standard sceptical argument is based on the closure principle:

**Closure Principle for Knowledge:** If *s* knows that *p*, and *s* knows that *p* entails *q*, then *s* knows that *q*.

So if I know that I have two hands then I must know that I am not a BIV. The traditional sceptical argument claims that, since you cannot know that you are not a BIV, you cannot know that you have two hands. The externalist denies that this is the case by claiming that in order to know that you have hands, you do not need to rule out the possibility that you are a BIV. Instead you just need to form your belief in the right way, whether that be reliably, safely or whatever. The externalist can maintain that the closure

principal holds because you can also know that you are not a BIV, provided that your belief is again formed in the right way.

The reason that the sceptical argument fails, according to the externalist, is that it is wrong to say that you do not know that you are not a BIV. You can know this but you know it in an externalist way. It does not matter that you cannot rule out this possibility. Provided that it is the case that you are not a BIV and your belief that this is the case was formed in the right way then you do know that you are not a BIV. It seems that the internalist, on the other hand, will either have to accept scepticism or deny the closure principle.

Knowledge plus internalist justification is basically equivalent to internalist knowledge. The only difference is that we are not claiming that the internalist justification is necessary for *knowledge* but rather that justification is a beneficial extra that usually accompanies knowledge. However, we can still ask whether internalist justification is closed. It seems that if it is then we will have to accept the restricted sceptical conclusion that we cannot have any reflectively justified knowledge. If this is indeed the only knowledge that is more valuable than mere true belief, then for it not to be possible would be a big problem. So, it seems preferable to take the other option; that is, to deny closure for internalist justification. This would mean that one could have internally justified knowledge of everyday propositions without having to do the impossible and rule out sceptical scenarios (For a similar argument, see Pritchard 2005: Ch3).

It doesn't seem that this will be particularly problematic for the view sketched above, as we do not need justification to be closed for either of the purposes that we discussed. Even if justification is not closed, it will still be able to stabilize our beliefs. Justification has the benefit of stabilising our beliefs, making them stronger in the face of conflicting



evidence. For this purpose we need to be able to respond to real and likely possibilities such as the possibility that the sign points in the wrong direction.

The other reason that justification is valuable is because it is needed to back up our knowledge claims. And this is unaffected by the denial of closure. It can, however, go some way to explaining why the externalist answer to scepticism seems unsatisfactory. Although we can have brute externalist knowledge of the denials of sceptical hypotheses, we cannot have the more valuable type of knowledge, that which is accompanied by internalist justification. If our brute externalist knowledge is no more valuable than mere true belief then, if this is all that we can have with regard to the denials of sceptical hypotheses, something will appear to be lacking.

## **5.5 Conclusions**

In this chapter I have examined the various value problems and sketched out an answer to them. The important factor is the value of reflectively accessible justification. This gives knowledge a value beyond that of mere true belief. We can understand knowledge in an externalist way but we must also make space to account for the importance of reflection. In doing so we can come to understand why we value knowledge, of a particular, common type and also understand why the externalist answer to scepticism is not satisfactory. Internalist justification is of the utmost importance even if it is not necessary for knowledge.

---

## Conclusion

---

This thesis had two aims: firstly, to defend a particular methodology in epistemology and secondly to defend the importance of reflective access. The methodology that I advocate is to be as comprehensive as possible rather than being very focussed on one particular model. When we take a very focussed perspective it can leave us blinkered to other important aspects of the thing that we are investigating. Within this methodology there will be space for traditional methodology. Although this has faced criticism from naturalists and experimentalists, I believe that it still has a fundamentally important role to play. The concepts that we investigate in philosophy, including the concept of knowledge, are inherently complex and this is perhaps why we are interested in them. Conceptual enquiry, will, therefore be an important starting point to any investigation and will have to play a large part in any philosophical investigation.

Conceptual enquiry is not, however, the only valid methodology in epistemology. Insights into knowledge can be gained from non-traditional sources as well. To this end, I examined two non-traditional approaches to epistemology, the naturalism of Hilary Kornblith and the genealogical approach of Edward Craig. Each of these methodologies gives us different insights into knowledge. Kornblith, with his focus on knowledge in the natural world, demonstrates the important continuity between human and non-human animal knowledge. It also shows us how scientific data can be useful in giving a different perspective on the subjects of philosophical study. However,

## *Conclusion*

Kornblith's focus on this continuity obscures the distinctively human aspects of knowledge

In contrast, Craig focuses almost entirely on the human aspects of knowledge and ignores the continuity with the animal world. He is primarily interested in looking at how the concept of knowledge developed, unlike Kornblith who is explicitly interested in the natural phenomenon of knowledge. Craig's approach tells us a lot about how the concept of knowledge that we have now might have developed in a state-of- nature, giving us insight into the role that knowledge plays. Information exchange is a vital part of our interaction with other human beings and it should come as no surprise that we have a concept that developed from the need to assess our informants.

Although both Craig and Kornblith have important things to say about knowledge the fact that they take the particular perspectives that they do means that neither of them can develop a comprehensive theory of knowledge. We need to take into account the possible evolutionary circumstances that led to knowledge existing in nature, but we also need to understand the pressures that have led to the cultural evolution of our concept. These two perspectives are not necessarily mutually exclusive, however. In a sense, the cultural evolution that Craig discusses can be seen as building on the natural evolution that Kornblith focuses on. Obviously, it is not as simple as to say that Craig's theory takes up where Kornblith's leaves off, but we can perhaps see them as being interested in such different aspects of knowledge that they are not as at odds as they first seem.

Although Craig's focus on the role of enquiry to human beings is good, I think that he does not go far enough in emphasising the role of reflectively accessible justification. This brings me to the second aim of the thesis; to bring back a focus on reflection that seems to have been lost somewhat with the advent of externalist

## *Conclusion*

accounts of knowledge. I argue that reflection has an important role to play, even if it is not strictly necessary for knowledge. It is vital when it comes to making knowledge claims and in the transmission of information between agents. There are cases where knowledge is not accompanied by this justification, as evidenced by various thought experiments. However, these cases are rare. The norm is to have knowledge accompanied by reflective justification.

This fact is what underlies our intuition that knowledge is more valuable than mere true belief. Generally, it is more valuable because it is accompanied by justification. This justification makes our beliefs more stable in that they are better equipped to stand up to conflicting evidence, this stability adds value to our knowledge. However, knowledge that is not accompanied by this justification is not more valuable. It is these two different types of knowledge that underlie our apparently conflicting intuitions. Allowing that there is knowledge *simpliciter* and knowledge plus can account for our intuitions about value. In addition, it can explain the unsatisfactory nature of externalist accounts of scepticism. We cannot have the more valuable type of knowledge of the denial of sceptical hypotheses and so we feel as if we have not got the full answer that we were looking for. Instead we are left with an adequate but much less valuable type of knowledge.

The methodology that I advocate I believe leads naturally to this distinction between different types of knowledge. If we are to fully understand knowledge then we need to take into account the different aspects of knowledge. As it is a complex concept, attached to a complex phenomenon, it should be no surprise that it necessitates a complex theory to explain it. I do not claim, however, to have fully explained all aspects of knowledge. There are bound to be other perspectives that

### *Conclusion*

could also be brought to bear on the questions of what knowledge is and why it is valuable. However, such investigation will have to be left for a later date.

---

## **APPENDIX**

### **Real-life vs. Philosophical Chicken Sexing**

---

The chicken sexing case, which is generally presented as evidence for externalism is sometimes quoted as a real-life example and sometimes as a fictional or hypothetical case. Although the roots of the example are in the real-world, the case is actually quite different. It is not clear that the real-life case supports externalism to the extent that the hypothetical case does. This, of course, is not necessarily of great philosophical importance, as there are plenty of thought experiments that are unlikely to ever actually occur – brains-in-vats, swampmen and Mary the colour scientist, to name but three. However, as this example is sometimes presented as one taken from real life it seems that there is good reason to try and sort the fact from the fiction.

#### **A.1 A Brief History of Chicken-sexing in Philosophy**

The first mention of chicken sexers in a philosophical context appears to be in Gasking's paper 'Avowals' (1963). In this paper Gasking is talking not about grounds for belief but rather grounds for statements. However, the example used is very similar to the example used later in the literature by epistemologists. Gasking presents the example as if it is a real-life case but doesn't explicitly say that it is. He claims that the chicken sexer is trained to identify chicks by looking at photographs of chicks and making a judgement about whether they are male or female. He is then told whether he

is right or wrong. After a period of time the trainee chicken sexer is increasingly correct in his judgements. Gasking's claim is that the chicken sexer does not know what it is that distinguishes male from female even though he can reliably sort the chicks by sex.

The example came to be used by epistemologists not as an example of statements without grounds but rather as a case where the chicken sexer has knowledge but lacks reflectively accessible grounds. There are also two major developments in how the case is described. The first can be seen in Richard Foley's book *The Theory of Epistemic Rationality* (1987).

This chicken sexer has no inductive evidence that that his past beliefs about the sex of the chicks usually have been correct and suppose he has no other reason to think his beliefs have been correct. He has never been told, for instance, whether he has been right... he simply picks up a chick and waits for the belief... to arise in him (Foley 1987: 168)

Foley's chicken sexer is not just in the position that he does not know how he differentiates between the sexes but also he does not know whether or not he has the ability to differentiate. Foley sounds sceptical as to whether such people do in fact exist although claims that 'some people testify that they do' (Foley 1987: 168).

The second development can be seen in Robert Brandom's paper 'Insights and Blindspots of Reliabilism' (1998). Like Gasking, he writes as if he sees it as a real-life case but does not explicitly say that it is. Brandom sets up the example such that it is not only the case that chicken sexers do not know how they do it but that they are actually mistaken about how they do it. Specifically, the chicken sexers think that they are using some visual cue to distinguish the chicks but in fact they discriminate on the basis of smell.

In fact, as I hear the story, it has been established that although these experts uniformly believe that they make the discrimination visually, research has shown that the cues their discriminations actually depend on are olfactory. (Brandom 1998: 575)

This development in the story means that even the limited grounds that the chicken sexer may have had – “I know it’s male because it looks male” – is undermined, making sure that there are no internally accessible grounds available to the sexer. Brandom is using the example as a case “in which someone in fact reliably responds to some sort of stimulus, without having any idea of the mechanism in play” (Brandom 1998: 575).

Both in this form and in the earlier varieties the chicken sexer example is used widely in contemporary epistemology. For example, in *Virtues of the Mind* (1996) Linda Zagzebski discusses the chicken sexer describing it as an example that highlights the difference between the intuitions of internalists and externalists; an internalist would deny that the chicken sexer knows, where the externalist is quite happy to claim that he knows. Although Zagzebski admits to scepticism about the veracity of the example (Zagzebski 1996: 301), she is happy to let it go for the sake of argument. Chicken sexing is also cited by Goldman (1975:114); Kornblith (1982:245); Pritchard (2005: 43-4) to name but a few.

## **A.2 Real World Chicken sexers**

Where does all this talk of chicken-sexing come from in the first place? Are there actual chicken sexers or are they the creation of a philosopher’s imagination? Chicken sexers do indeed exist but are rather different to the ones in the philosophical examples. For the purpose of producing eggs farmers need a large amount of female chicks – cockerels being of little use in this regard – however, in most breeds of chicken male



and female chicks are identical for the first few weeks of life.<sup>20</sup> For reasons of economics, it is better for an egg producer not to rear male chicks that will be of no commercial use to him and so it is important to have a way of distinguishing between the sexes of the chicks.

It wasn't until the 1920's in Japan that such a method was developed. In 1925 Kiyoshi Masui and Juro Hashimoto discovered a method of determining the sex (Masui and Hashimoto 1933). Although there are no external distinguishing marks there is a genital eminence or 'bead' in the chicks' cloaca - the cavity into which the intestinal, urinary and genital tracts open. By everting the vent the bead can be observed. This process is not particularly pleasant and is described by Biederman and Shiffrar in the following way

The chick is held in the left hand (for a right-handed person) and the fecal contents are squirted into a container to clear the cloaca. Gentle but firm pressure from the two thumbs and right forefinger are exerted to spread the ventral surface of the cloaca upwards to expose the eminence, called the "bead". (Biederman and Shiffrar 1987: 641)<sup>21</sup>

To begin with, the discovery of these beads was of little use. The eminences are about the size of a pinhead and sexing decisions must be made quickly because keeping the cloaca open for any length of time puts the chick at risk of injury or even death. Because of this using the beads as a sexing tool was initially impossible. People had trouble finding the eminence in the short amount of time required. However, a farmer by the name of Manabu Kojima spent three years examining the genital eminences of over 6000 baby chicks finally becoming extremely accurate in his determinations. From

---

<sup>20</sup> There are certain cross-breeds, such as Black Stars and Red Stars, that produce different colours in male and female chicks from hatching making it easy to identify the sexes.

<sup>21</sup> For more on the methodology of chicken sexing see Blount (1934) and Kiyoshi (1933)

this, a training method was developed and chicken sexing academies began to spring up. A trained sexer will look at each chick for less than a second before making their judgement and can be up to 98% accurate in their judgements whilst sexing up to 1000 chicks an hour.

### **A.3 The similarities and differences**

As we can see the difference between the philosophical example and the actual case is quite large. In the first place the chicken sexers do know how they tell the difference between chicks; they look for the bead and classify the chicks accordingly. Also, they know that they are reliable at sexing, high levels of reliability are required by employers, records are kept and sexers are paid according to their accuracy (Biederman and Shiffrar 1987 :643). Finally, I can find no research that shows that smell plays any part in the sexing of chicks. In the Biederman and Shiffrar experiment photographs were used rather than live chicks and the professional chick-sexers were still highly accurate (although slightly less accurate than when sexing live chicks, possibly due to the quality of the pictures).

The interesting philosophical point that can be taken from the real-life chicken sexing example is not one about externalist knowledge but rather about skilled action. It is an extreme example of a case where someone is very good at something to the extent that they could no longer really explain how they do what they do. The chicken-sexers do have justified beliefs because they know that they are skilled at what they do. However, what they can't do is to point out exactly how they do what they do as quickly as they do it. There is no question about the general method by which they do it – they are looking at the genital eminences – but they are doing it so quickly it is surprising that they are able to take it in at all. What we have in real-life chickens sexing

*Appendix: Real-life vs. Philosophical Chicken Sexing*

cases is not a case of brute externalist knowledge but it is a case of ability knowledge that is hard to explain. The chicken sexers have developed such good recognitional capacities that they are sorting the chicks seemingly by immediate recognition, seemingly without thinking about it and the fact that we are capable of learning to do such a thing is interesting enough in itself.

---

## Bibliography

---

- Alcock, J. (1975), *Animal Behaviour: An Evolutionary Approach*, Sunderland, MA: Sinauer Associates.
- Audi, R. (1998), *Epistemology: A Contemporary Introduction to the Theory of Knowledge*, London: Routledge.
- Austin, J. L. (1946), 'Other Minds', *Aristotelian Society Supplementary Vol.* 20: 148-187. [Page references are to the reprint in J. L. Austin, *Philosophical Papers*, 3<sup>rd</sup> ed., J. O. Urmston and G.J. Warnock, eds., Oxford, Oxford University Press, 1979, pp. 76-116.]
- Ayer, A.J. (1957, reprinted 1976), *The Problem of Knowledge*, Harmondsworth: Penguin Books.
- (1973), *The Central Questions of Philosophy*, London: Weidenfeld and Nicholson.
- Bealer, G. (1993), 'The Incoherence of Empiricism', in S.J. Wagner and R. Warner, eds., *Naturalism*, Notre Dame, IN: University of Notre Dame Press.
- Beckoff, M. and Jamieson, D, eds., (1990), *Interpretation and Explanation in the Study of Animal Behaviour*, i, *Interpretation, Intentionality, and Communication*, Boulder, CO: Westview Press. [Page references refer to the reprint as *Readings in Animal Cognition*, Cambridge, MA: Massachusetts Institute of Technology Press, 1996]
- Biederman, I. and Shiffrar, M.M. (1987), 'Sexing Day Old Chicks: A Case Study and Expert Systems Analysis of a Difficult Perceptual-Learning Task', *Journal of Experimental Psychology: Learning, Memory and Cognition* 13(4): 640-645.

*Bibliography*

- Bermúdez, J. L. (2006), 'Knowledge, Naturalism and Cognitive Ethology: Kornblith's Knowledge and Its Place in Nature', *Philosophical Studies* 127: 299-316.
- Bernecker, S. and Dretske, F. (2000), *Knowledge: Readings in Contemporary Epistemology*, Oxford: Oxford University Press.
- Bird, A and Tobin, E. (2009), 'Natural Kinds', *The Stanford Encyclopedia of Philosophy* (Spring 2009 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2009/entries/natural-kinds/>.
- Blount, W.P. (1934), *Sexing Day-old Chicks*, London: Poultry World Ltd.
- BonJour, L. (1980), 'Externalist Theories of Empirical Knowledge', *Midwest Studies in Philosophy* 5: 53-73.
- (1985) *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard University Press.
- (2006) 'Kornblith on Knowledge and Epistemology', *Philosophical Studies* 127: 319-335.
- and Sosa, E. (2003) *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, Oxford: Blackwell.
- Boyd, R. (1988), 'How to be a moral realist', in G. Sayer-McCord, ed., *Essays on Moral Realism*, Ithaca, NY: Cornell University Press, pp.181-228.
- Brandom, R. B. (1994), *Making it Explicit Reasoning, Representing, and Discursive Commitment*, Cambridge, MA: Harvard University Press.
- (1998), 'Insights and Blindspots of Reliabilism', *The Monist* 81(3): 371-392.
- Butler, R.J., ed., (1962), *Analytical Philosophy*, Oxford: Basil Blackwell.
- Chisholm, R., (1988) 'The Indispensability of Internal Justification', *Synthese* 74(3): 285-96.
- (1989), *Theory of Knowledge* (3<sup>rd</sup> Edition), London: Prentice Hall.
- Conee, E. and Feldman, R. (1997) 'The Generality Problem for Reliabilism', *Philosophical Studies* 89: 1-29.

*Bibliography*

- (2001), 'Internalism Defended', in H. Kornblith, ed., *Epistemology: Internalism and Externalism*, pp. 231-60.
- Craig, E. (1990), *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Clarendon Press.
- Davidson, D. (1984), *Inquiries into Truth and Interpretation*, Oxford: Oxford University Press.
- DePaul, M. and Zagzebski, L. (2003), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford: Oxford University Press.
- Descartes, R., trans. J. Cottingham, (1641/1996), *Meditations on First Philosophy*, Cambridge: Cambridge University Press.
- Dretske, F. (1970), 'Epistemic Operators', *Journal of Philosophy* 67(24): 1007-23.
- (1971), 'Conclusive Reasons', *The Australasian Journal of Philosophy* 49(1):1-22.
- Foley, R. (1987), *The Theory of Epistemic Rationality*, Cambridge, MA: Harvard University Press.
- Gasking, D. (1962), 'Avowals' in R.J. Butler, ed., *Analytical Philosophy I*, Oxford: Basil Blackwell, pp.154-169.
- Gettier, E. (1963), 'Is Justified True Belief Knowledge?', *Analysis* 23(6):121-3.
- Ginet, C. (1975), *Knowledge. Perception and Memory*, Dordrecht: D. Reidel.
- Greco, J. (1999), 'Agent Reliabilism', *Philosophical Perspectives* 13: 273-296.
- (2003), 'Knowledge as Credit for True Belief', in M. DePaul and L. Zagzebski, *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford: Oxford University Press, pp. 111-34.
- (2005) 'Justification is not Internal' in M. Steup and E. Sosa, eds., *Contemporary Debates in Epistemology*, Oxford: Blackwell, pp.257-270.
- (2007), 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17: 57:69.

*Bibliography*

- Goldman, A. (1975), 'Innate Knowledge', in S.P. Stich, ed., *Innate Ideas*: Berkeley, CA: University of California Press.
- (1976), 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73(20): 771-791.
- (1980), 'The Internalist Conception of Justification', *Midwest Studies in Philosophy* 5(1): 27-52.
- (1986), *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.
- Grice, H.P. (1975), 'Logic and Conversation', *Syntax and Semantics* 3: 41-58.
- Haddock, A., Millar, A. and Pritchard, D., eds., (2009), *Epistemic Value*, Oxford: Oxford University Press.
- Haddock, A. and McPherson, F., eds., (2006), *New Essays on Disjunctivism*, Cambridge, MA: Harvard University Press.
- Heinrich, Bernd (1999), *Mind of the Raven: Investigations and Adventures with Wolf-Birds*, Harper Collins.
- Herman, L. and Morrel-Samuels, P. (1990), 'Knowledge Acquisition and Asymmetry Between Language Comprehension and Production: Dolphins and Apes as General Models for Animals', in M. Beckoff and D. Jamieson, eds., *Readings in Animal Cognition*, Cambridge, MA: Massachusetts Institute of Technology Press, pp. 289-306. [Page references refer to the 1996 reprint.]
- Kaplan, M. (1991), 'Epistemology on Holiday', *The Journal of Philosophy*, 88(3): 132-154.
- (2000), 'To What Must an Epistemology Be True?', *Philosophy and Phenomenological Research* 61(2): 279-304.
- Kauppinen, A. (2007), 'The Rise and Fall of Experimental Philosophy', *Philosophical Explorations* 10(2): 95-118.
- Kiyoshi, M. and Hashimoto, J. (trans. H. Okumura) (1933), *Sexing Baby Chicks*, Vancouver: Journal Printing Company.

*Bibliography*

- Knobe, J. and Nichols, S. , eds, (2008), *Experimental Philosophy*, Oxford: Oxford University Press.
- Kornblith, H. (1982) ‘The psychological turn’, *Australasian Journal of Philosophy*, 60(3): 238- 253.
- (2001), *Epistemology: Internalism and Externalism*, Oxford: Blackwell.
- (2002), *Knowledge and its Place in Nature*, Oxford: Oxford University Press.
- (2006), ‘Reply to Bermúdez and Bonjour’, *Philosophical Studies* 127: 337-349.
- (2008), ‘Knowledge Needs No Justification’, in Quentin Smith, ed., *Epistemology: New Essays*, Oxford: Oxford University Press, pp.5-23.
- Kusch, M. (2009), ‘Testimony and the Value of Knowledge’, in Haddock, A., Millar, A. and Pritchard, D., eds., *Epistemic Value*, Oxford: Oxford University Press, pp. 60-94.
- Kvanvig, J. (2003), *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- (2009), ‘Responses to Critics’, in A. Haddock, A. Millar and D. Pritchard, eds., *Epistemic Value*, Oxford: Oxford University Press.
- Lackey, J. (2007), ‘Why We Don’t Deserve Credit For Everything We Know’, *Synthese* 158: 345-61.
- Lewis, D. (1996), ‘Elusive knowledge’, *Australasian Journal of Philosophy* 74(4): 549-567.
- Lehrer, K. (1990), *Theory of Knowledge*, London: Routledge.
- and Cohen, S. (1983), ‘Justification, Truth and Coherence’, *Synthese* 55(2): 295-413.
- Luper, S., ed., (2003), *The Sceptics*, Aldershot: Ashgate.
- Lycan, W. (1993), ‘Sartwell’s Minimalist Account of Knowing’, *Philosophical Studies* 73(1): 1-3.



*Bibliography*

- McDowell, J. (1982), 'Criteria, Defeasibility and Knowledge', *Proceedings of the British Academy* 68: 455-79. [Page references are to the reprint in J. McDowell, *Meaning, Knowledge and Reality*, Cambridge, MA: Harvard University Press, 1998, pp.369-94.]
- (1998), *Meaning, Knowledge and Reality*, Cambridge, MA: Harvard University Press.
- Millar, A (2004), *Understanding People: Normativity and Rationalising Explanation*, Oxford: Oxford University Press.
- (2007), 'What the Disjunctivist is Right About', *Philosophy and Phenomenological Research* 74(1): 176-198.
- (*forthcoming*), 'Knowledge and Recognition', in D. Pritchard, A. Millar and A. Haddock, *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press, pp. 122-239.
- Neta, R. and Pritchard, D. (2007), 'McDowell and the New Evil Genius', *Philosophy and Phenomenological Research* 74(2): 381-96.
- Nichols, S., Stich, S. and Weinberg J.M. (2003), 'Meta-skepticism: Meditations in Ethno-Epistemology' in S. Luper, ed., *The Sceptics*, Aldershot: Ashgate Publishing, pp.227-247
- Nozick, R. (1981), *Philosophical Investigations*, Cambridge, MA: Harvard University Press.
- Peirce, C.S. (1877), 'The Fixation of Belief', *Popular Science Monthly* 12: 1-15, reprinted in C.S. Peirce, *The Essential Peirce*, N. Houser and C. J. W. Kloesel, eds., Bloomington, IN: Idaho University Press, 1992, pp. 109-123.
- Plantinga, A. (1997), 'Warrant and Accidentally True Belief', *Analysis* 57(2): 140-145.
- Prinz, J. (2004), *Gut Reactions*, Oxford: Oxford University Press.
- Pritchard, D. (2002), 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10(3): 283-307.
- (2005), *Epistemic Luck*, Oxford: Oxford University Press.

*Bibliography*

- (2006), ‘McDowellian Neo-Mooreanism’ in Haddock, A. and McPherson, F., eds., *New Essays on Disjunctivism*, Cambridge, MA: Harvard University Press.
- (2007), ‘Recent Work On Epistemic Value’, *American Philosophical Quarterly* 44(2): 85-110.
- (2008), ‘The Value of Knowledge’, *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta, ed.,  
URL=<<http://plato.stanford.edu/archives/fall2008/entries/knowledge-value/>>.
- (2009), ‘The Value of Knowledge’, *The Harvard Review of Philosophy* 16:2-19.
- (*forthcoming a*), ‘Knowledge and Understanding’, in D. Pritchard, A. Millar and A. Haddock, *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press, pp. 8-188.
- (*forthcoming b*) ‘Evidentialism, Internalism, Disjunctivism’, Trent Dougherty, ed., *Evidentialism and its Discontents*, Oxford, Oxford University Press. (Draft accessed online at <http://www.philosophy.ed.ac.uk/people/full-academic/documents/EvidentialDisjunctivism.pdf> 14/05/2010, cited with permission of the author)
- Pritchard, D., Millar, A., and Haddock, A. (*forthcoming*), *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Putnam, H. (1975), *Mind, Language and Reality: Philosophical Papers*, ii, Cambridge: Cambridge University Press.
- Povinelli, D and Eddy, T. (1996), ‘What Young Chimpanzees Know About Seeing’, *Monographs of the Society for Research in Child Development* 61(3).
- Reynolds, S. L. (2002), ‘Testimony, Knowledge and Epistemic Goals’, *Philosophical Studies*, 110(2): 129-161.
- Russell, B. (1948), *Human Knowledge: Its Scope and its Limits*, London: Allen and Unwin.
- Sartwell, C. (1991), ‘Knowledge is Merely True Belief’, *American Philosophical Quarterly* 28(2): 157-65.

*Bibliography*

- (1992), ‘Why Knowledge is Merely True Belief’, *The Journal of Philosophy* 89(4): 167-80.
- Sayer-McCord, G., ed., (1988), *Essays on Moral Realism*, Ithaca, NY: Cornell University Press.
- Schiffer, S. (1999), ‘Contextualist Solutions to Scepticism’, *Proceedings of the Aristotelian Society* 96: 317-333.
- Smith, M., ed., (2008), *Epistemology: New Essays*, Oxford: OUP.
- Sosa, E. (1991), *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- (1994), ‘Philosophical Scepticism and Epistemic Circularity’, *The Aristotelian Society*, Suppl. Vol. 68: 263-290.
- (1997), ‘Reflective Knowledge in the Best Circles’, *Journal of Philosophy* 94(8): 410-30.
- (1999), ‘How to Defeat Opposition to Moore’, *Philosophical Perspectives* 13: 141-153.
- (2007), *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Clarendon Press.
- Steup, M. and Sosa, E., eds., (2005), *Contemporary Debates in Epistemology*, Oxford: Blackwell.
- Stitch, S.P. (1975), *Innate Ideas*, Berkeley, CA: University of California Press.
- Stroud, B. (1984), *The Significance of Philosophical Scepticism*, Oxford: Oxford University Press.
- (1989), ‘Understanding Human Knowledge in General’ in M. Clay and K. Lehrer, eds., *Knowledge and Scepticism*, Boulder, CO: Westview Press. [Page references refer to the reprint in H. Kornblith, ed., *Epistemology: Internalism and Externalism*, Oxford: Blackwell, 2001, pp.126-46.]
- Spicer, F. (2008), ‘Are there any conceptual truths about knowledge?’ *Proceedings of the Aristotelian Society* 108(1): 43-60.

*Bibliography*

- Strawson, P.F. (1992), *Analysis and Metaphysics*, Oxford: Oxford University Press.
- Vogel, Jonathan (2000), 'Reliabilism Leveled', *Journal of Philosophy* 97(11): 602-623.
- Wagner, S.J. and Warner, R., eds., (1993), *Naturalism*, Notre Dame, IN: University of Notre Dame Press.
- Wason, P. (1960), 'On the Failure to Eliminate Hypotheses in a Conceptual Task', *The Quarterly Journal of Experimental Psychology* 12: 129- 40.
- Weatherson, B. (2003), 'What good are counterexamples?', *Philosophical Studies* 115(1): 1-31.
- Weinberg, J., Nichols, S., Stich, S. (2001), 'Normativity and Epistemic Intuitions', *Philosophical Topics* 29:429-460. [Page references are to the reprint in J. Knobe and S. Nichols, eds., *Experimental Philosophy*, Oxford: Oxford University Press, 2008, pp.17-46.]
- Williams, M. (2000), 'Dretske on Epistemic Entitlement', *Philosophy and Phenomenological Research* 60(3): 607-12.
- (2001), *Problems of Knowledge: A Critical Introduction to Epistemology*, Oxford: Oxford University Press.
- Wittgenstein, L. (1953), *Philosophical Investigations*, 2<sup>nd</sup> ed., Oxford: Blackwell.
- Zagzebski, L. (1996), *Virtues of the Mind: An Enquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.
- (1999), 'What is Knowledge?', in J. Greco and E. Sosa, *The Blackwell Guide to Epistemology*, Oxford: Blackwell, pp. 92-116.
- (2003), 'The Search for the Source of the Epistemic Good', *Metaphilosophy* 34:12-28.
- (2004), 'Epistemic Value Monism', in J. Greco *Ernest Sosa and His Critics*, Oxford: Blackwell.