PSYCHOLOGY AND THE GOALS OF EPISTEMOLOGY

Katherine Helen Puddifoot

Submitted for the degree of PhD

Department of Philosophy
University of Sheffield

May 2012

Abstract

This thesis defends the controversial view that results taken from psychology should be used by epistemologists developing their accounts, methodological naturalism, and describes how it should be implemented. I argue methodological naturalism should be implemented by epistemologists adopting an approach to epistemology where they study the goals that can be achieved through epistemic classifications, accepting an account of knowledge, rationality or justified belief only if it describes epistemic concepts in such a way that they can be used by humans to achieve important goals. I show how such an approach complements methodological naturalism by allowing results from psychology to come into the picture to reveal important goals to be achieved.

My project begins with a general defence of methodological naturalism. I show how results from psychology provide important insights relevant to epistemology that are not available in the armchair. I respond to a number of common criticisms of methodological naturalism, showing that they are based on misconceptions about the position. I evaluate a number of accounts of the methodology to use in epistemology consistent with methodological naturalism. I evaluate accounts that suggest that epistemologists should retain armchair methods such as conceptual analysis, simply adding a naturalistic component. I reject these on the basis that they cannot avoid some fundamental problems with the armchair methodology. I evaluate the claim that knowledge is a natural kind that can be revealed by consulting intuitions about obvious cases, providing a number of examples showing how implausible the claim is. I argue that knowledge is not a natural kind and instead it is a category required to answer to human needs. Finally, I show how epistemologists can give sufficient recognition to how epistemic concepts are answerable to human needs, avoiding the various criticisms of armchair methods, while adopting a methodology fully consistent with methodological naturalism.

Contents

Acknowledgements

I would like to thank a number of people who have helped me over the course of writing this thesis. To begin with I would like to thank my family for their constant support, especially my mum, Mary McHugh, who always gave my siblings and I the belief that if we wanted to achieve something then we could, no matter how the signs indicated otherwise. I would like to thank the Woods, for that extra support they have always so generously provided.

I thank my supervisor, Stephen Laurence, for his patience, encouragement and great attention to detail. I also thank Christopher Hookway, who acted as my secondary supervisor and provided helpful advice and comments throughout the process of writing the thesis. Thanks are also due to Bob Hale, Rosanna Keefe and Robert Stern who acted as Directors of Graduate Studies during my studies at Sheffield.

In writing this thesis I have benefitted from the comments and advice of many friends and colleagues in the Sheffield Philosophy Department and beyond, among them Bernardo Aguilera, James Andow, George Botterill, David Ekstrand, Paul Faulkner, Graeme Forbes, Carl Fox, Paul Giladi, Kate Harrington, Joseph Kisolo-Ssonko, Suilin Lavelle, Jamie McIver, Jonathan Payne, Cristina Roadevin, Robin Scaife, Naoki Usui, Inga Vermeulen, Jack Wadham and Stephen Wright. I thank them for their comments and advice on my work. I also want to thank all my friends in and around the Sheffield postgraduate community for their support. I want to give a special thanks to Kate Harrington and to Stephen Wright for that extra support, willingness to discuss my work, and for the proof reading skills that they provided. Thanks also go to Amy Gray for her proofreading.

Financial support for this thesis was provided by the Will and Ruby George Trust and by the British Society for the Philosophy of Science, for which I am extremely grateful.

Most of all I thank William for just being himself and being there. He has always been my very best friend, providing me with love, encouragement and support. I thank him for everything. Without him this thesis wouldn't be what it is.

Introduction

Interdisciplinary approaches to academic subjects are becoming increasingly popular. Projects that are interdisciplinary are thought by many to be more interesting, more worthy of funding and generally more fruitful than projects that use the findings and methods of a single discipline. Yet among many philosophers there remains the feeling that there are aspects of philosophy that should not be influenced by other disciplines. For example, many feel that epistemology in particular should not be interdisciplinary, that there are distinctively epistemological questions that cannot be answered by using the methods and findings of those from other disciplines. They aim to shelter epistemology, and epistemologists, from having to be informed by the sciences, for example, defending the idea that we can understand the nature of knowledge, rationality and justified belief in the armchair without considering the findings of people working in other disciplines. Particularly prevalent is the idea that normative questions, about how people ought to behave, such as those within the purview of epistemology can only be answered by philosophers using distinctively philosophical methods.

It is this idea about the division of labour within epistemology (or lack thereof) that is challenged by advocates of 'methodological naturalism' or 'naturalized epistemology' (henceforth, methodological naturalism). They defend the idea that epistemologists should be informed by scientific methods, findings and theories when they develop their accounts of knowledge, rationality or justified beliefs. They argue that facts about how we actually obtain and revise beliefs should inform these accounts about how people ought to obtain and revise their beliefs.

There are two ways to think about the motivation for adopting a naturalistic approach to epistemology, one positive and one negative[1]. Sometimes methodological naturalism is defended by focusing on its positive features. It is argued that naturalistic approaches to epistemology can provide realistic and achievable epistemic notions and concepts. While conceptual analysis might provide *some* account of what knowledge, rationality and justification could be, only a naturalistic approach to epistemology will guarantee an account of what they are that puts them within the reach of ordinary thinkers, with human cognitive limitations. A naturalistic approach to epistemology, so the argument goes, will be able to provide an accurate account of how people can successfully revise their beliefs. At other times, the merits of the naturalistic project are identified with the failures of alternative projects. Methodological naturalism is viewed primarily as an alternative to traditional philosophical approaches such as conceptual analysis and a priori theorising, which have come under increasing criticism in recent decades. If epistemological categories such as knowledge, rationality and justified belief are taken to be discoverable through the methods of science, epistemologists are not restricted to conceptual analysis and a priori reasoning.

The purpose of this thesis will be to evaluate the prospects for methodological naturalism and highlight the positive arguments that can be provided in favour of it. I will highlight the strengths of methodological naturalism and challenge those who argue that there should be a division of labour, with philosophers using distinctively philosophical methods to answer normative questions while people working in other disciplines address descriptive questions about how people actually behave. I highlight the advantages of developing an epistemology informed by empirical findings, scientific

[1] This method is especially evident in Phillip Kitcher's 'The Naturalists Return' but most methodological naturalists are motivated by both of these ingredients.

theories and other thoughtful descriptions of how people behave. But before we engage in this project there are a few preliminary points to clarify about methodological naturalism.

1. Methodological not Ontological Naturalism

The first point to clarify is about the exact nature of the methodology that I advocate in this thesis. There are various ideas that fall under the description 'naturalism'. As I have mentioned, the form of naturalism that is the main concern of my thesis is what is called 'methodological naturalism'. It involves claims about the methods that should be used in epistemology. According to methodological naturalism, methods from the sciences—empirical methods that study what nature is like—should influence the theories that philosophers develop[2]. It is important to distinguish this position from a weaker and less controversial position, ontological naturalism. Ontological naturalism is the commitment to the idea that reality is made up of only natural rather than supernatural entities. Defenders of ontological naturalism maintain that philosophers should not discuss magical or mysterious entities because reality is only constituted of non-magical and not-so-mysterious natural entities. One might be tempted to conflate methodological naturalism with ontological naturalism because the claim that the methods of the natural sciences should be used to study knowledge, rationality and justified belief suggests that these entities are natural—there seems to be a connection between the idea that knowledge, rationality and justified beliefs are natural and the idea that the methods of the natural sciences should be used to study them. However, there is an important difference between ontological naturalism and methodological naturalism.

[2] There is an important distinction underlying this choice of phrase between philosophers who make direct use of the methods of science and those who use scientific findings and theories to develop the philosophical position without actually engaging in scientific inquiry. The former can be counted as 'experimental philosophers' whereas the latter don not fall naturally into this category. Both, however, can be counted as naturalised epistemologists.

Like ontological naturalists, methodological naturalists are committed to the idea that reality is made up of natural entities, but they are also committed to a particular way of developing epistemological theories to understand that reality. They think that science should be used in epistemology, something that ontological naturalists need not accept[3].

An ontological naturalist is committed to the idea that the items under consideration in epistemology—beliefs, belief forming processes, and so on—are natural entities. In the current philosophical climate this is not really a controversial matter. Most contemporary philosophers would agree that the entities that are the object of consideration in epistemology are natural rather than supernatural[4]. They accept that the subject matters of epistemology are psychological states and/or processes, i.e. beliefs and/or belief forming processes, rather than supernatural entities[5]. What a methodological naturalist says about epistemology is more controversial. She says that the categories we use to evaluate the psychological states or processes under consideration in epistemological theories—categories such as 'knowledge', 'rationality' and 'justified belief'—can only be properly understood by taking the findings and theories of the sciences into consideration. In her view, it is only possible to get an accurate picture of knowledge, rationality and justified belief by considering facts about nature—about the psychological processes and conditions under which beliefs are obtained—through the study of science. This *is* a controversial claim. It is rejected, for example, by all philosophers who defend *a priori* analysis and/or conceptual analysis as

[3] Papineau (2009) has a nice summary of the distinction between ontological and methodological naturalism.
[4] I owe this point to Papineau (2009) once more.
[5] This point becomes important in some discussions of the merits of naturalised epistemology. Kitcher (1992) is sometimes interpreted as arguing for epistemic naturalism on the basis that epistemologists refer to psychological entities in their discussions. Bonjour (1994) and Feldman (1999) rightly point out that epistemologists can speak about psychological entities without being committed to the idea that scientific methods should be used in epistemological theorising. They can be ontological naturalists without being methodological naturalists.

the sole methods to be used in epistemology[6]. It is rejected by those epistemologists who do not study empirical findings, who remain in their armchairs, studying their concepts and trying to understand what they and other people happen to think knowledge, rationality and justified belief might be, without considering what the sciences say about the belief forming processes we use. These epistemologists might well be ontological naturalists, and accept that knowledge, rationality and justified belief are natural entities, but they think that the nature of those entities can be discovered without taking results from the sciences into consideration.

The relationship between ontological and methodological naturalism is thus asymmetrical. Methodological naturalists are ontological naturalists[7]. They are committed to the idea that epistemologists should use science to discover how people can obtain beliefs because the beliefs that people have when they have knowledge or rational belief are natural entities and can therefore be understood by natural science. They are committed to the idea that science is important because they think that we can only understand what the *natural* entities of knowledge, rationality and justified belief are through the sciences. Ontological naturalists, on the other hand, don't have to be methodological naturalists. They can accept that beliefs and belief forming processes, justified beliefs and knowledge, are natural entities while maintaining that the categories that the beliefs fall into such as 'knowledge' and 'justified belief' can be discovered without recourse to the sciences. For example, they can argue that knowledge is a

[6] Bealer (1998) provides a good example of a philosopher who insists that a priori theorising is the only way to investigate epistemic notions.

[7] It is possible to imagine a person is a methodological naturalist and also an ontological naturalist but if such a person existed it would seem that she was not adopting a reasonable position. A person who genuinely thought that the subject matter of epistemology was supernatural would not seem to be very wise if she spent her energies examining natural psychological phenomena, such as beliefs and belief forming processes. It would seem to be a vain hope that a person could understand the supernatural phenomena they associate with knowledge, rationality and justified belief by examining psychological phenomena.

natural phenomenon but that we can only uncover its nature by examining our concept of knowledge and working out when it is appropriate to apply the concept. The epistemological side of the study of knowledge could involve the study of our concepts, although those concepts are applied to natural, psychological entities.

My goal in this thesis is to establish that epistemologists should be methodological naturalists as well as ontological naturalists. I aim to show that it is right for epistemologists to draw from the sciences when they want to understand the nature of knowledge, rationality or justified belief. Not only should we accept that knowledge, rationality and justified belief are natural, we should aim to understand how beliefs deserve the positive status associated with being classified as any of these things by considering what sciences says.

2. Which sciences can inform epistemology?

Before launching into my defence of methodological naturalism there is another thing that it might be helpful to briefly consider to get a fuller picture of methodological naturalism; that is, which sorts of sciences can be relevant to epistemology. Methodological naturalists often focus upon psychology, and in this thesis I shall often use talk of the influence of results from psychology as shorthand for the influence of results from the sciences more generally. This is because psychology is the science that is most obviously pertinent to epistemology. If we think that epistemology should take into consideration how the mind works, then we are committed to the idea that disciplines like psychology, which directly study how the mind works, are relevant to epistemology. Neuroscience also has the potential to influence our epistemological theories because it too can tell us how the mind works. This is because it can inform us about how the brain computes and represents information. Once we understand the

styles of processing available to the human brain, we can work out what it would be for the brain to process information well[8].

The arguments in favour of using psychology in epistemology have the same sort of structure as the arguments in favour of using neuroscience. Both disciplines provide some sort of an understanding of how the mind works, and therefore provide some ideas about how people equipped with a human mind/brain can best go about forming beliefs. There are, however, other disciplines that could be relevant to epistemology not because they tell us how the mind works but because they are instrumental in the development of evolutionary theories. Archaeology and anthropology are two of the disciplines that influence evolutionary theories about cognition. Findings from archaeological and anthropological studies are used to establish how human cognition has developed, including through the process of natural selection[9]. Theories about how the brain evolved can influence theories about how well adapted the brain is to the environment. They can be used to develop an account of how people can attain the best beliefs possible by understanding how well-attuned a cognitive system is to the environment as the result of the process of natural selection. This is a more complicated way in which empirical studies can influence epistemological theories, but it is an interesting way, and we shall consider in more depth how evolutionary theories can influence epistemology in various sections in the thesis[10].

[8] Churchland (1987) argues, for example, that if connectionist theories about cognitive processing are correct this has implications for how people can best form beliefs. In her view we can understand whether connectionist theories are correct by considering the findings of neuroscience, so neuroscience will be highly influential on epistemology.

[9] The work of Boyd and Richerson and Joseph Heinrich provide examples of the use of methods from archaeology and anthropology to support evolutionary psychology.

[10] The work of Gerd Gigerenzer and his colleagues, e.g. Gigerenzer and Todd (1999), is a good example of an evolutionary approach to epistemic norms, and one that shall be discussed throughout this thesis.

3. Overview of the thesis

The reader should now have an understanding of the subject matter of this thesis. It is a position on the methods used in epistemology that suggests that the sciences should not be kept apart from epistemology, that they should be embraced by epistemologists, and the results from the sciences should influence the theories that they develop. Epistemologists should not only accept that knowledge is natural rather than supernatural, they should also accept that it can only be understood by considering the implications of results from the sciences, such as psychology, neuroscience and evolutionary theories like those presented by anthropologists and archaeologists. The purpose of the thesis is to defend this controversial claim and to describe the best way to implement a methodology to use in epistemology consistent with this claim.

The task of chapter one is to show that methodological naturalism provides the only way to develop an adequate epistemological theory. Methodological naturalists have argued for their position on the basis that progress can be and has been made by epistemologists who draw upon the results of science. But this argument has been challenged by those who defend armchair methods, who have argued that methodological naturalists overplay the benefits of considering results from science. By responding to some of the stronger arguments used to challenge the methodological position, I will present a defence of methodological naturalism.

I shall present my case for methodological naturalism by responding to some particularly powerful criticisms presented by Richard Feldman. Feldman argues that any progress that has been made by epistemologists in response to psychological results could have been made in the armchair. He claims that epistemologists should be able to develop theories that accommodate all possibilities rather than just actual facts, and that

armchair epistemologists are able to imagine all possibilities. I shall argue that Feldman's conclusion is based on a number of false assumptions: that epistemologists should aim to develop theories that reflect all possibilities about human cognition; that science has not revealed any possibilities that could not have been imagined in the armchair; and, finally, that even if science reveals possibilities more efficiently than armchair methods, the default position should be that people should use the armchair methods rather than the scientific methods that reveal the same possibilities. In chapter one, we shall see that epistemologists should be interested in facts about human cognition rather than possibilities, things that science is uniquely placed to provide access to, and that science can and has revealed facts relevant to epistemology that would have been unavailable in the armchair. We shall see that a non-naturalistic epistemology will be inadequate because it will fail to reflect scientifically established facts about human psychology.

After the defence of methodological naturalism in chapter one, the focus of chapter two will be on addressing a number of criticisms generally levelled against the position. Remember that methodological naturalism has been highly controversial, particularly because some philosophers are reluctant to accept that normative theories should not be influenced by the descriptions provided by the sciences. The aim of chapter two is to show that the criticisms generally levelled against methodological naturalism are unfounded and based on a number of misunderstandings about what it is to be committed to the position. I aim to persuade the reader that if she properly understands what it is to be a methodological naturalist she will see that it is not the unattractive prospect that it might initially seem to be. On the contrary, it is a very effective approach to epistemology.

In chapter three I shall begin to examine the prospects of various accounts of the methods to use in epistemology that are consistent with methodological naturalism currently on the table. Some philosophers have attempted to defend traditional methods, such as conceptual analysis and intuition-driven reflective equilibrium, by arguing that they can be integrated with results from psychology to develop a viable account of knowledge, rationality and justified belief consistent with methodological naturalism. They think that the results of conceptual analysis and/or intuition-driven reflective equilibrium can be adapted to reflect what science says about human cognitive capacities, thereby gaining all of the advantages of methodological naturalism. Chapter three will highlight a number of problems with this particular way of implementing methodological naturalism. In it I will argue that it is advantageous to aim for some degree of continuity between folk concepts and epistemological theories, but that there are many problems with the traditional methods of epistemology that cannot be overcome by approaches that attempt to integrate these methods.

In chapter four I evaluate Kornblith's claim that knowledge is a natural kind, that is, a category that exists in nature and independently of human interests, and that the aim of epistemology should be to understand this natural kind. I highlight a number of problems with this approach. Kornblith argues that intuitions can provide a guide to the nature of the natural kind of knowledge because they identify obvious cases of knowledge. I deny that intuitions should be expected to perform this function. Then I discuss the possibility also presented by Kornblith that a natural kind revealed by science will be identical with the category of knowledge. I focus on the natural kind discussed by Kornblith—true belief produced by cognitive equipment performing its proper function—and argue that it is highly implausible that it is either necessary or

sufficient for knowledge that a belief should belong to this natural kind. I argue that the reason why the natural kinds account is so unsatisfactory is that the category of knowledge is not natural and is required to answer to human needs. I conclude that epistemologists should aim to understand the goals that people can achieve through their epistemic classifications because whether it is right to classify any belief as knowledge, rationality or justified belief will depend upon the goals that we can achieve by giving something these labels. Because our epistemic concepts are required to answer to human needs, so that people can achieve important goals, knowledge should not be considered to be a natural kind.

Chapter five will develop the idea first outlined in chapter four that epistemologists should study the goals that can be achieved and the work that can be done by any classification that they might make of something as knowledge, rationality or justified belief. In this I am influenced by Craig (1990) and Weinberg (2006). I argue that epistemologists should adopt what I shall call a 'teleological approach to epistemology' and examine the goals that people can achieve through their epistemic classifications. They should only accept a theory if it describes our epistemic concepts in such a way that we can achieve important goals by applying them. I will show that by adopting such an approach to epistemology it is possible to avoid many of the problems with alternative methodologies used in epistemology.

I shall argue that a teleological approach to epistemology should be combined with methodological naturalism. A teleological approach can be made consistent with methodological naturalism because results from the scientific study of human psychology can be used to reveal valuable goals that can be achieved through our epistemic classifications. By considering results from psychology when developing an

idea of the goals that can be achieved through our epistemic classifications, epistemologists will be able to develop epistemological accounts which reflect all the valuable goals that can be achieved through our epistemic classifications and therefore genuinely answer to human needs, as they are required to. I shall argue that by taking such an approach to epistemology, epistemologists are able to avoid all the criticisms of current methods to use in epistemology, developing an account consistent with results from psychology as well as one answerable to human needs.

By the end of this thesis, then, the reader should be persuaded that methodological naturalism is a valuable project, and that it can only really be successful if it is developed in the manner just described, by combining a teleological approach to epistemology with a scientifically-informed understanding of human psychology that reveals goals that can be achieved by our epistemic classifications. As the thesis progresses, the flaws of the prominent current accounts of the methods to use in epistemology consistent with methodological naturalism will be exposed and we shall see how epistemologists can improve upon these methodologies by taking a particular combination of a teleological approach to epistemology and methodological naturalism, which I shall call radically naturalistic neo-pragmatism.

Chapter One: Why Epistemology Should Be Naturalized

Introduction

The aim of the current chapter will be to develop a positive account of how the study of results from psychology can and has led to progress in epistemology. My tactic will be to address a number of arguments that have been used to suggest that armchair epistemologists can make the same, or at least equally good, progress as methodological naturalists. I will challenge various arguments presented against the positive position of methodological naturalism by showing that results from psychology have provided a unique contribution to our understanding of knowledge, rationality and justified belief unavailable to non-naturalists. We shall see that only by leaving the armchair can epistemologists achieve the best possible epistemological theory.

1. A challenge to methodological naturalism

In my view, the importance of methodological naturalism can be understood by considering the progress that has been made by epistemologists who have used psychological results to develop their theories. The proof is in the pudding when it comes to epistemology; we can see that methodological naturalism is valuable because of the gains that can be made by epistemologists willing to engage with results from psychology. It has been suggested by Richard Feldman, however, that epistemologists will not benefit from engaging in methodological naturalism because the facts about human psychology that are relevant to epistemology are accessible in the armchair. In

this chapter I will highlight the benefits of methodological naturalism by responding to this challenge to the idea that it has made a valuable contribution[11].

Feldman (1999) supplies the following argument against the suggestion that methodological naturalism has been shown to lead to important progress in epistemology:

1. To show that their methodology is valuable, methodological naturalists must show that the consideration of psychological results has led to improvements to epistemological theories that could not have been made in the armchair.

2. Methodological naturalists have not shown this.

Conc. Methodological naturalism has not been shown to be valuable.

Feldman's argument comes out in his discussions of the work of Hilary Kornblith (1994) and Alvin Goldman (1993). Both of these authors are methodological naturalists, who have argued that improvements have been made to reliabilism in response to psychological results.

According to the simple version of reliabilism, a belief is justified if and only if it is produced by a reliable cognitive mechanism. Kornblith suggests that this simple version of reliabilism should be modified or abandoned because psychological results suggest that, on this version, reliabilism implies that inductive inferences cannot be justified. Psychological results suggest that inductive reasoning it is unreliable, so if justification comes with reliability, inductive inferences cannot be justified because they are not

[11] There are, of course, many other challenges to methodological naturalism. Some of these will be discussed in the next chapter. For current purposes I am focusing on Feldman's challenge because I think that once we see that his assumptions are incorrect it will be directly apparent that methodological naturalism is valuable.

reliable. Kornblith thinks that this is an unsatisfactory conclusion because he is committed to 'ballpark psychologism'; he thinks that people obtain beliefs roughly as they ought to and he acknowledges that people often use induction, so he concludes that all inductions cannot be unjustified. He suggests that reliabilism should be modified or abandoned because a successful epistemological analysis would not imply that all inductive inferences are unjustified.

Goldman presents a similar argument with respect to visual perception. He suggests that the simple version of methodological naturalism should be modified because it implies that all beliefs based on visual object recognition are unjustified. This is because psychological results suggest that visual object recognition can be unreliable under some conditions, although it is reliable under others. Goldman suggests that reliabilists should attempt to avoid the undesirable conclusion that all beliefs based on visual object recognition are unjustified by concluding that all beliefs produced by a single cognitive process do not have to be categorised together as reliable or unreliable, justified or unjustified. In Goldman's view, assessments of the reliability of a particular process should be made relative to specific environmental conditions. Some beliefs based on visual object recognition can then be categorised as justified, because they are the product of a cognitive mechanism used in sympathetic conditions, although the same cognitive mechanism would yield unjustified beliefs in other environments in which it is unreliable.

In Feldman's reply to these arguments he suggests that methodological naturalists have not established the value of their methodology simply by showing that epistemologists have made modifications to their analyses in response to psychological results such as those about induction and visual object recognition. He suggests that they can only

establish the value of their methodology by showing that modifications *that could not have been made in the armchair* have been made in response to results from psychology.

> If armchair epistemologists are allowed to assume that ballpark psychologism is true (or that scepticism is false), then surely they are also entitled to assume that we can (and do) sometimes generalize on the basis of too few examples, ignore irrelevant information, or let our emotions influence our inferences. So, not only is the fact that it is possible that some inductive beliefs are not justified available to us in our armchairs, so also is the fact some actual inductive beliefs are not justified. So, if the point of the induction example was to make us see that we should reject (R) [the simple version of reliabilism] in favor of something like (R1) [a modified version of reliabilism], then it fails to show that results from scientific psychology are particularly useful (Feldman, 1999: p. 175).

The implication of this passage is that the scientific study of psychology will only be proved to be useful by methodological naturalists if they can show that the relevant facts about human psychology are unavailable to armchair epistemologists. In Feldman's view, if a psychological fact is available in the armchair it is not worthwhile for epistemologists to use results from psychology to find out that fact. He illustrates this point by suggesting that scientific psychology hasn't been helpful for epistemologists developing reliabilism because they could have known about the unreliability of induction in the armchair. In section 2 I shall present a challenge to this assumption, arguing that it can be valuable to use results from psychology in epistemology even if it is possible to find out facts in the armchair.

Even if we were to accept Feldman's first premise, that results from psychology are only relevant to epistemology if they present facts unavailable in the armchair, however, there is another response methodological naturalists can make to Feldman: that is, they can show that many relevant facts about human cognition are simply not available in the armchair. Feldman's position depends upon the contention that all psychological facts

relevant to epistemology are available in the armchair. In section 3 I shall present some examples that challenge this assumption.

The methodological naturalist might think that she has succeeded in refuting Feldman's argument if she can show that there are facts about psychology relevant to epistemology that are available only to those who consider results from psychology. However, Feldman has one more response he could make use of; he maintains that epistemologists should be interested in possible facts about human psychology rather than empirically supported facts about psychology.

> In these cases, knowledge of contingent facts—the fact that people sometimes make bad inductive inferences, the fact that people are able to identify objects better in some circumstances than others—may influence one's epistemology. The way this knowledge influences epistemology is to make one aware of actual, and thus possible, examples that an analysis of justification must accommodate.
>
> However, these cases do not provide the materials needed for a defense [sic] of methodological naturalism. Neither argument shows that we really need results from empirical science in order to formulate versions of reliabilism that are good enough... Since the analyses are supposed to be adequate to all possible cases, these cases only need to be possible in order to show that (R) needs revision. Since it is easy to figure out from our armchairs that the examples like these are possible, it is easy to figure out from our armchairs that (R) needs revision. We don't need results from psychology to tell us this (ibid.: p. 173).

In this passage Feldman suggests that results from psychology are useful and important to epistemology because they are able to reveal possible ways that human psychology might be. He thinks that it is valuable to understand the possible ways human psychology might be because, in his view, a successful epistemological analysis should accommodate all possible cases. He continues by suggesting that 'it is easy to figure out from our armchairs that the examples like these are possible' (ibid.), showing that he thinks that armchair epistemologists have a good grasp of the space of possibilities of

human psychology. He concludes that results from psychology haven't contributed to epistemology because armchair epistemologists are really interested in the possible nature of human psychology, which, in his view, is accessible in the armchair.

My view is that the argument just presented is based on two false assumptions. The first is that epistemologists should be able to accommodate all possible ways that human psychology might be. If the methodological naturalist can show that epistemologists should be interested in *facts* about human psychology *rather than possible ways human psychology might be*, she will provide reason for thinking that results from psychology are important because, as it shall be argued in section 4, these provide the best understanding of human psychology. The second assumption is that epistemologists in their armchairs would be able to obtain a good understanding of the space of possibilities of human psychology, if possibilities were to be the subject of epistemologists' investigations. If the methodological naturalist can show that epistemologists in their armchairs are unlikely to imagine some of the possible ways that human psychology could be then this calls into question whether armchair epistemology is adequate, even if people accept that every possibility about human psychology should be taken into consideration. Section 4 will present the responses that methodological naturalists can make to both of these assumptions.

In sum, then, Feldman's position rests on the following assumptions:

Assumption 1. If armchair epistemologists can access facts about human psychology or imagine the space of possibilities of human psychology, then the appeal to psychological findings is not necessary and hence methodological naturalism isn't valuable.

Assumption 2. Armchair epistemologists can access the facts about human psychology relevant to epistemology from the armchair.

Assumption 3. Epistemologists are interested in empirically supported facts only because they highlight possible ways human psychology might be.

Assumption 4. Epistemologists are able to imagine all relevant possibilities about human psychology, that is, the space of possibilities of human psychology, from the armchair.

In this chapter I argue that all four of these assumptions are false and show that Feldman's challenge to the argument for methodological naturalism therefore fails. On the contrary, I will show that there are good reasons for thinking that progress can be made by epistemologists as a result of their being influenced by results from the scientific study of psychology.

2. Assumption 1

We have found, then, that Feldman argues that results from the scientific discipline of psychology are only useful to epistemologists if and only if they reveal possibilities about human psychology that epistemologists could not have imagined in the armchair. If armchair epistemologists can access facts about human psychology or imagine the space of possibilities of human psychology, then it is not necessary to appeal to psychological findings and hence methodological naturalism isn't valuable. In his view, what methodological naturalists must show is that similar results to those produced in response to psychological results could not have been achieved in the armchair. I would suggest, on the contrary, that methodological naturalists can provide a satisfactory defence of their position simply by showing that scientific results provide a particularly

effective way to understand the psychology of human cognition relevant to epistemology.

Recall that the aim of methodological naturalism is to suggest that results from psychology can be useful in epistemology. Feldman argues that such results can only be useful if they reveal facts about human psychology unavailable in the armchair. He does not deny that facts about psychology can be relevant to epistemology. For Feldman's position to be acceptable, then, it has to be consistent to maintain that it can be useful for epistemologists to be aware of facts about human psychology but that it is nonetheless not useful to adopt a particular method available to discover these facts. In other words, it has to be consistent for him to maintain that it is acceptable to desire a particular outcome but resist using a particular method available for achieving it.

Now there certainly are circumstances in which it is appropriate to reject a certain means to a particular end. One can easily imagine occasions in which it would be too risky or time consuming to use a particular means. Under such conditions it would be consistent to have a desire but not be willing to use a particular means to achieve it. But the case that we are discussing does not fit into this category. We are discussing a situation in which epistemologists have scientific results about human psychology at their disposal, while accepting that facts about human psychology are relevant to epistemology, but nonetheless maintain that they cannot benefit from considering the scientific results. It would not be particularly risky or time-consuming for epistemologists to consider empirically gathered results from psychology. On the contrary, it seems that the risk-averse epistemologist should be keen to borrow results from the sciences rather than making modifications to her theory based on her merely speculative conclusions reached in the armchair. And as we shall see in coming

sections, an epistemologist who is concerned with efficiency would benefit from focusing on empirically supported facts about human psychology rather spending time speculating about the various possible ways human psychology might be.

We have good reasons, then, for thinking that results from psychology can provide an efficient way to obtain an accurate picture of human psychology for epistemologists to consider when developing their theories. Consequently, Feldman is unwarranted in the conservatism he demonstrates when he accepts that facts about human psychology are relevant to epistemology but maintains that we should not use results from the scientific study of psychology to discover these facts. Even if Feldman is correct and facts about human psychology can be discovered from the armchair, epistemologists should be willing to consider results from psychology because they provide an efficient way to achieve the desired understanding. For epistemologists to resist using results from psychology would be to unnecessarily tie their hands, neglecting to use a useful resource, without good reason.

3. Assumption 2

In my view, then, methodological naturalism can be valuable because results from the scientific study of psychology provide an efficient way to understand human psychology, even if facts about human psychology are available in the armchair. But the methodological naturalist doesn't need to concede to Feldman that facts about human psychology are available in the armchair. We can reject Feldman's assumption 2; that armchair epistemologists can access the facts about human psychology relevant to epistemology from the armchair. It is worth noting that Feldman bases assumption 2 on his discussion of Kornblith and Goldman's arguments about induction and visual object recognition. Within the context of a discussion of these examples it seems convincing

that the facts about human psychology used to make modifications to epistemological analyses are accessible in the armchair. Armchair epistemologists are aware that inductive inferences sometimes go wrong and that people are better equipped to recognise objects in particular environments. But there are other facts about human psychology that are relevant to epistemology, which are not readily accessible in the armchair.

The claim that epistemologists are aware of the facts about human psychology in the armchair runs into trouble as soon as one considers the relevance of unconscious and automatic processes to epistemology. Psychologists in recent decades have gathered a significant body of evidence suggesting that beliefs that would have previously been assumed to be based on consciously accessible thought processes are actually based on unconscious and automatic processes. For example, psychological findings suggest that people are poor at accessing the basis for their beliefs through introspection and often confabulate reasons for their beliefs[12]. These findings present a challenge for epistemological internalists[13]. For internalists, knowledge requires one to be able to access, on reflection, reasons for a belief. So, for example, what makes one's perception of a tree a case of knowledge is the reflectively accessible visual awareness of the tree. At first glance, the evidence of unconscious and automatic processes suggest that many beliefs which people would usually call knowledge do not actually meet the criteria for knowledge presented by internalists because the believer does not have access to the

[12] For philosophical discussion of some of the psychological literature on confabulation see Carruthers, P. 'Introspection: Divided and Partly Eliminated', *Philosophy and Phenomenological Research* 80 (2010), 1, pp. 76-111.

[13] My discussion here has some parallels with arguments presented by Kornblith, 2002, chapter two based on evidence about failures of introspection. However, I think that the recent studies on unconscious and automatic processes such as implicit biases present an even more compelling case against internalism than those that Kornblith describes for the reasons presented below.

basis for the beliefs. The psychological results suggest that the internalist constraint that people should be able to access the basis for their beliefs could lead to a sceptical picture of what is required of knowledge according to which few beliefs belong in the category.

How can internalists respond to this worry? One option for them is to simply bite the bullet and maintain that beliefs based on unconscious and automatic processes are not knowledge, even if many of our beliefs are based on such processes. They can maintain that a person who knows something has to have access to *every* basis for the belief, and that when people aren't aware of every basis, for example because they are under the influence of unconscious and automatic processes, they fail to know. Another option available to the internalist is to maintain that people have knowledge as long as they have access to *some* essential basis for their beliefs (Pappas, 2008). They could accept that unconscious and automatic processes often influence the outcome of belief formation, and that people can't access every basis for their beliefs when these processes are influential, but maintain that people can nonetheless have knowledge in these cases, as long as they have partial access to the basis for their beliefs. Alternatively, the internalist could deny that any belief formed on the basis of a process that is inaccessible is justified but allow that such a belief is justifiable; if a person is able to marshal some reason for the belief it could become justified[14].

[14] Richard Feldman himself stresses the importance of distinguishing beliefs for which support can be given and beliefs which have been properly arrived at. Beliefs are 'justified' in the view of Conee and Feldman as long as reasons can be marshalled in support of the beliefs, whether or not the reasons given really identify how the belief is arrived at. If a belief is both well-supported and properly arrived at it gets the even better status of 'well-foundedness'. According to this view, then, a belief can be justified without being properly arrived at, as long as some reasons can be marshalled in support of it (Feldman and Conee, 2007).Thanks to Paul Faulkner for pointing out that this option is available to the internalist.

What is most important for the current discussions is that psychological findings are relevant to assessments of the plausibility of these options. The first option will only be viable if people are sometimes able to access every basis for some and perhaps many of their beliefs. Otherwise, internalists taking the option would present a sceptical account of knowledge according to which people never or at least rarely actually have knowledge. The concept of knowledge would, according to this account, never or rarely apply to anyone because people would not be able to meet the demand to have full access to every basis for their beliefs. Because results from psychology are able to show us how unlikely it is that people will be able to access every basis for their belief, they are directly relevant to any assessment of this position[15]. As results from psychology can tell us if people are able to access every basis for their beliefs the viability of the first option for internalism will depend upon what science says about human psychology.

Results from psychology are also relevant to the assessment of the second option because precise details about the nature of unconscious and automatic processes could decide whether it is acceptable for internalists to hold that beliefs influenced by unconscious and automatic processes can be classified as knowledge. An example might help to illustrate this point. There is a growing consensus among psychologists that implicit biases underlie our evaluations. These findings are particularly interesting because they are so unsettling. They suggest that people have implicit biases against

[15] Epistemologists have realised from the armchair that people cannot access the basis for all of their beliefs. Armchair epistemologists have recognised that people are not able to remember the reasons why they formed a belief in the distant past. They accept that a belief can be justified without requiring that we remember these past reasons for those beliefs and they do so based on this armchair observation. However psychological findings are still able to inform epistemologists about the extent to which we can be expected to access the basis for all our beliefs at the point at which we form a belief because they reveal the extent of non-conscious influences on our beliefs—processes that we cannot hope to understand in the armchair.

members of some groups and for other groups which lead them to make positive associations with members of some groups and negative associations with members of other groups without these associations tracking the truth[16] in any way. They suggest that people favour white people to black people, young people to old people, white children to black children, non-Arab Muslims to Arab Muslims, and the list goes on (Nosek et al, 2007). Psychological results suggest that people of all political persuasions, from liberals to conservatives, are subject to these biases. The results also suggest that the biases are reflected in people's behaviour, even when they are dealing with high stakes. For example, one study showed that implicit racial biases accurately predicted voting behaviour of US citizens in the 2008 election race between Barack Obama and John McCain (Greenwald et al., 2009).

These results from psychology suggest that some unconscious and automatic processes that influence people's judgements are extremely undesirable. Implicit biases which do not track the truth and have nothing like a rational basis are a paradigm case of mere prejudice, something that is the opposite of knowledge. It is therefore paramount that epistemologists should ensure that knowledge is not ascribed lightly when people are committed to beliefs as a result of the influence of unconscious and automatic processes. Consider, for example, a voter assessing Obama's policies prior to deciding who to vote for in the 2012 US elections. There is evidence suggesting that people's voting behaviours are related to implicit biases, so it also seems fair to assume that people's judgements about Obama's policies are also likely to be so affected. Imagine now that the voter judges that Obama's health care policies have been more damaging to the US economy than those that were outlined by the Republican Party. She has a

[16] I mean to be neutral in my discussions of truth tracking between the various theories associating epistemic notions with truth tracking. But I take it that a process that does not track the truth should not be viewed positively from an epistemic perspective.

number of reasons for thinking this, but she also has a number of equally compelling reasons for thinking the alternative, and what is decisive in her judgement is that she is influenced by an implicit bias against black people. This belief about Obama's health care policy should not be categorised as knowledge because it is only accepted as a result of the influence of a negative implicit bias which fails to track the truth. This means that if internalism is to provide an adequate account of the requirements for knowledge it should be able to exclude cases such as this from the category. However, if internalists take the second option outlined above, and accept that a person can have knowledge while having access to only some essential basis for her belief, the belief about Obama's health care policy could be categorised as knowledge as long as the person has access to some basis for the belief, even if what is decisive is some prejudice of which she is unaware (and which, moreover, she might find unsatisfactory as a basis for her belief if she were aware of it). The psychological results on implicit bias therefore present a particularly strong case against classifying beliefs as knowledge based on only partial access to the bases for the belief.

The third option is that internalists could maintain that any belief based on an inaccessible basis is unjustified but could become justified if the believer could marshal some reasons for the belief. On this account the belief about Obama's health care policy would be unjustified if based on an implicit bias but could become justified if the person who believes it can provide some good reasons for it. The internalist should be resistant to this move because it would allow that a belief which is based upon a process that is unconscious and fails to track the truth can become justified just because its believer can provide some reason for the belief, even though the belief is not based on the reason. They should insist that the justifying reasons for a belief should be causally related to

the process that produced the belief in order to avoid classifying those based on implicit biases as knowledge. Once again the results from the scientific study of psychology are relevant to the viability of the epistemological stance because they show the nature of the psychological processes that underlie many of our beliefs; that is, they show that they are unconscious and fail to track the truth. These facts about human psychology, revealed by results from science, suggest that epistemologists should make an active effort to ensure that they do not allow any reason that is provided to be a source of justification for a belief.

So, then, it seems that psychological results about unconscious and automatic processes can be relevant to epistemology because they present a challenge for internalists to deal with and can be used to suggest ways to modify the approach. Equally important for current purposes is the fact that epistemologists are only likely to address the challenges presented by the prevalence of unconscious and automatic processes if they are informed by the sciences since information about these processes are not available in the armchair.

The scientific study of human psychology is particularly important to epistemologists interested in unconscious and automatic processes because the processes are revealed by psychologists using methodologies that discover hidden aspects of human psychology that simply would not be understood in the armchair. For instance, implicit biases are revealed by using a method that uncovers the smallest difference in reaction speed between responses to members of particular groups. The Implicit Association Test (IAT), for example, asks people to respond to members of particular groups (e.g. female or male) and certain items (e.g. flower or insects) when they are shown on a computer screen (Greenwald et al., 1998). Researchers then analyse the data and if they find that

people are faster at categorising one group with positive stimuli then they conclude that the people more readily associate that group with positive features. This method aims to uncover implicit biases by discovering very small differences in response speed that cannot be detected outside the lab. It also aims to get people to show biases that they wouldn't be aware of so couldn't describe if asked explicitly about their preferences. Because people are asked to respond quickly, and do not know what they are being tested for, it is assumed that they respond using automatic processes rather than the controlled ones that they might use if asked for an explicit response (Ferguson, 2007).

Armchair epistemologists are unlikely to obtain an accurate picture of the place of unconscious and automatic biases in judgement because they won't be able to use these methods. They are likely to be unaware of their own biases and won't be sensitive to the biases of others that are only revealed by close analysis of response speeds. As this phenomenon demonstrates, not all epistemically relevant facts are accessible from the armchair. It seems that when it comes to automatic and unconscious processes of belief formation and revision, epistemologists should leave their armchairs and consult results from psychology.

It is not only the literature on unconscious and automatic processes that reveals facts about human psychology relevant to epistemology but unavailable in the armchair. The literature on human reasoning also reveals surprising facts about human psychology by using methods that are unavailable in the armchair. Gerd Gigerenzer and his colleagues at the ABC research group, for example, have revealed that people use heuristic reasoning strategies and that the strategies are highly efficient and truth tracking[17].

[17] For general expositions of the project of the ABC research group see Gigerenzer, G. & Selten, R. (eds.) (2001) *Bounded rationality: The adaptive toolbox* (Cambridge, MA, US: The MIT Press, 2001) as well as

Heuristics are mechanisms of reasoning which use only a highly limited amount of information. The 'take the best' heuristic is one example. People can use this heuristic to decide between two solutions to a reasoning problem. Those using this heuristic check one cue at a time to see whether it supports either of two options. Gigerenzer describes a case in which a person, Marty, uses take the best to decide which of two schools has the lowest drop-out rate (Gigerenzer, 2007). Marty cannot get hold of information about drop-out rates from the schools, but he does have access to other information about the schools, e.g. attendance rates, parental participation rates. If Marty were to use take the best, then he would focus on one piece of information about the schools that he has available, such as the attendance rates, and see whether one school rates better according to that information. Marty only has to consider other information if the first piece of information does not discriminate between the schools. Taking one cue at a time, Marty would carry on until he found a factor that favoured one school. Take the best is therefore representative of heuristics in general because it requires a person to retain and process little information at any one time. In contrast with more complex methods, it does not involve weighting various factors and making a complex calculation about how certain evidence supports a particular conclusion.

Gigerenzer and his colleagues suggest that heuristics like take the best are used widely. This might be surprising. But perhaps more surprising, and more important to epistemology, is the suggestion that the heuristics are highly effective. Gigerenzer and his colleagues have presented empirical results suggesting that the use of heuristic methods can be not only as effective as, but actually more effective than, the use of other, more computationally demanding methods. The results present a challenge to

---

Gigerenzer, G. & Todd, P. (1999). *Simple Heuristics That Make Us Smart* (Oxford: Oxford University Press, 1999).

epistemologists who are committed to the standard picture of human reasoning, according to which rationality requires applying the principles of logic and/or probability theory to all relevant information, to ensure the consistency of one's beliefs (Stein, 1996). They suggest that people frequently are successful at reasoning, though they do not check whether a conclusion is consistent with all or even much of the relevant evidence, and in fact, they suggest that people can be *more* successful if they do not make such checks. As a result, Gigerenzer and his colleagues are lead to conclude that these reasoning processes, which they call 'fast and frugal methods of reasoning', are rational.

This challenge is once again based upon evidence gathered using sophisticated techniques from psychology that are unavailable to armchair epistemologists. Gigerenzer and his colleagues support their claims about the success rates of people using heuristic strategies by developing and utilising computational models (Gigerenzer & Todd, 1999: pp. 16-18). They use computers to model the steps that would be taken by a person using take the best and similar heuristics. Then they apply these models to real life problems, using the information available in ordinary situations. They compare the rates of success of models of heuristic reasoning to the rates of success of models of other methods of reasoning. By using this technique they have discovered that heuristics perform as well as, or better than, other more complex strategies. When this discovery is presented alongside arguments suggesting that humans are unlikely to be able to use more complex models of reasoning successfully (Cherniak, 1986), Gigerenzer and his colleagues present quite a compelling case for rethinking our standards for what it is to be rational. And this is a case that armchair epistemologists are unlikely to be able to make, or even challenge and reject, because they won't have the techniques required to

simulate the computations of the various strategies and to assess their relative success rates. So if epistemologists are interested in facts about human reasoning and its successes and failures, they should be willing to consider results from psychology because the relevant facts will not always be available in the armchair.

One more case in which techniques unavailable in the armchair have been used to make a strong case against an epistemological theory is Michael Bishop's challenge to the idea that epistemic responsibility and epistemic virtue should be associated with each other. Bishop (2000) characterises epistemic responsibility as having the internalist virtues associated with warrant and justification of 'coherence, having good reasons, fitting the evidence' (p. 180). He argues that the value of epistemic responsibility comes largely, if not primarily, from its relation to truth. In his view, it is a mistaken assumption of epistemologists who defend a notion of epistemic responsibility that epistemic responsibility is associated with reliability. These epistemologists assume, in Bishop's view, that a responsibly held belief will be more likely to be true than an irresponsibly held one. Bishop challenges this assumption by presenting empirical results suggesting that the most reliable reasoning strategies are not the ones that would be deemed responsible according to the traditional notion of epistemic responsibility advocated in epistemology. He suggests that the notion of epistemic responsibility should be sidelined in response to empirically gathered results because responsibility and reliability do not come together.

Bishop finds support for his arguments in findings presented by Paul Meehl in his book *Clinical Versus Statistical Prediction: A Theoretical Analysis and Review of the Literature*. Meehl compared the predictions made by experts to those made by actuarial statistic prediction rules. Actuarial predictions are made by the mechanical application

of formulas or rules, which are developed without the weighing of factors used by experts. The formulas or rules are such that they can be given to unskilled administrators to apply, removing the need for experts. So a person trying to work out how well students will perform in university would use a model where they are presented with figures of school performance and told to use certain weights to predict the university performance where the weights are based on empirically gathered data about, for example, the relationship between students' exam scores and their academic performance in university.

According to the empirical findings cited by Bishop, actuarial statistical prediction rules such as those used in university admissions outperformed experts in all tests of predictive ability. For example, actuarial statistical prediction rules predicted criminal recidivism based on past criminal and prison record better than expert criminologists (Dawes 1994, p. 83, cited in Bishop, 2000: p. 184). According to Bishop, these findings show that expert prediction, which follows the model of epistemic responsibility, is not the best strategy to use and that we can improve our epistemic practices by adopting something like actuarial statistical prediction rules rather than using those methods that would traditionally be deemed epistemically responsible, such as those used by experts[18]. They also suggest that our epistemic appraisals should be adjusted because we would normally associate epistemic responsibility with less effective, expert methods of reasoning. He suggests that we should instead associate epistemic responsibility with more effective and reliable methods more closely replicating actuarial statistical prediction.

[18] I say 'something like' the actuarial prediction rules just described because the empirical findings suggest that actuarial prediction rules based on less demanding processes, e.g. that do not involve weighting to provide the optimal relationship between a cue and the target outcome, will still be more effective than expert prediction.

We find in Bishop's arguments the suggestion that epistemologists should significantly rethink the traditional notion of epistemic responsibility—a suggestion about the modification of an epistemological analysis—supported using statistical modelling. This methodology, used by Meehl and his colleagues to reveal the increased success that comes from using statistical prediction rules, involves complicated statistical modelling and analysis; methods that aren't available in the armchair. Once again, then, we have found an example in which distinctive methods of psychology can reveal facts about human psychology of which epistemologists should take note because they present a challenge to epistemological theories.

What can we conclude, then? Feldman is wrong to suggest that all facts about human psychology relevant to epistemology will be available in the armchair. Some of the relevant facts might be accessible in this way. But there are important facts that will only be accessible to those who are willing to be informed about what is going on in the sciences. His assumption 2 is false and only those who are informed by the scientific study of psychology will have access to all of the facts about human psychology which are relevant to epistemology.

4. Assumptions 3 & 4

In the sections 2 and 3 I have presented some arguments suggesting that methodological naturalists will be in a better position to access facts about human psychology than armchair epistemologists. I think these arguments provide good motivation for epistemologists to reach for some science journals. But there is a further assumption underlying Feldman's position which, if correct, could be used to undermine methodological naturalism. Recall that Feldman claims that facts about human psychology are relevant to epistemology because they highlight possible ways human

psychology could be and that possibilities should be accommodated by successful analyses. He based this claim upon two assumptions:

> Assumption 3. Epistemologists are interested in empirically supported facts only because they highlight possible ways human psychology might be.
>
> Assumption 4. Epistemologists are able to imagine all relevant possibilities about human psychology, that is, the space of possibilities of human psychology, from the armchair.

Feldman could attempt to challenge methodological naturalism by arguing that possible ways human psychology might be are relevant to epistemology, and that an understanding of these possibilities is available in the armchair.

My first response to this argument is that it doesn't seem at all plausible that epistemologists should be interested in all possible ways human psychology could be. Epistemologists aim to understand human knowledge (as well as rationality and justified belief, etc.). They aim to provide an account of what is going on when people have knowledge; what beliefs are like when people know something. Epistemologists may use imagined thought experiments, but when they do so they are really interested in the features of beliefs that make them knowledge in the actual world. A successful epistemological theory describes the features of beliefs that merit the positive epistemological appraisal associated with the attribution of knowledge in the actual world filled with real people and epistemological issues.

Given this understanding of the aims of epistemology, facts about human psychology contribute by enabling epistemologists to develop theories of knowledge that reflect the

actual cognitive states people are in when they have knowledge. An account developed in response to psychological facts can avoid placing knowledge beyond the reach of people with ordinary cognitive capacities. It is therefore valuable for epistemologists to consider facts about human psychology because doing so will enable epistemologists to develop theories to reflect how human knowledge really is.

On the other hand, epistemologists won't necessarily benefit from considering all possible ways human psychology might be. Quite to the contrary, an epistemologist who tries to accommodate all possibilities is likely to provide a less accurate and useful theory as a result. For example, an epistemologist might consider the possibility that all humans are able to reason perfectly in accordance with logic and probability, checking the consistency of all the information available to them. At the same time she considers psychological studies suggesting that people are best able to achieve successful results by using heuristics because they are poor at checking the consistency of their beliefs in accordance with logic and probability theory. She develops an epistemological theory that is able to accommodate both of these possibilities. Perhaps she develops a theory according to which rationality involves using any reasoning strategy that leads to successful results. But she does not provide specific details about how people can achieve successful results because her theory needs to be able to accommodate different ways that people could achieve good results given the various ways human psychology might be.

Compare this situation to one in which an epistemologist considers empirically supported facts about psychology but ignores other possible ways that human psychology might be. Such an epistemologist could provide a detailed account of how people can go about reasoning to achieve successful results. She could specify the

specific reasoning strategies that people use when they are successful. In my view, the epistemologist would provide a better account of what it is to be rational when she considers results from psychology. She would weaken her position by being accommodating to possible psychological facts. She would present a less detailed account and one that less accurately reflects what actually goes on when people obtain and revise beliefs well. This suggests, in my view, that epistemologists should only aim to accommodate empirically established facts about human psychology because in doing so they can accurately reflect the reality of human cognition, its successes and failures.

I would therefore argue that epistemologists should only aim to accommodate empirically established *facts* about human psychology. But let us imagine that epistemologists are unconvinced by my argument and accept Feldman's suggestion that their theories should be able to accommodate all possible facts about human psychology. Unless it is possible to somehow restrict the relevant possibilities, epistemologists will become involved in an endless task when considering all sorts of possibilities that seem to be relevant to epistemology. One possibility is, for example, that Michael Jackson controls all human cognition from beyond the grave and makes people excellent at reasoning when they are listening to Thriller. Should epistemologists accommodate this possible fact about human psychology? What about the possibility that all humans have the same cognitive powers as crocodiles? If crocodiles have some unexpected cognitive quirk then should we accommodate this quirk in our epistemological theory because it is possible that human cognition could be the same as crocodile cognition? Surely these cannot be relevant possibilities. And yet Feldman does not describe a mechanism for epistemologists to use to dismiss these possibilities

as irrelevant. In fact he suggests that all possibilities should be accommodated, so if we give him an unsympathetic reading we could maintain that he is committed to the idea that these possibilities should be accommodated.

On a more sympathetic reading, we could assume that Feldman means to suggest that epistemologists should only focus on possible ways that human psychology might be that are sufficiently similar to how human psychology really is. However, this raises two new problems for Feldman. Firstly, it requires him to provide some account of how much similarity is required and it is not clear where this will come from. Secondly, it re-introduces psychological facts as relevant to epistemology through the back door. In order for epistemologists to work out how similar a psychological possibility is to reality they have to have a clear idea about the real nature of human psychology; that is, they need a comparison class and it will be provided by facts about human psychology. But if epistemologists use facts about human psychology as a similarity measure they imply that what is important is that epistemology reflects how human psychology really is. Otherwise, why would it be important to focus on possibilities that are similar to the reality? It seems that if we choose to focus on possibilities that are sufficiently similar to reality then this will be because we are really interested in how human psychology is. And if epistemologists are required to have a clear notion about human psychology, and are really interested in how human psychology is, then they could focus their energies on ensuring that their theory is compatible with the reality rather than speculating about various close possibilities. Under such circumstances, speculation about possibilities once again appears to have little benefit.

It seems, then, that epistemologists who followed Feldman's advice would either end up involved in the endless task of considering various possible ways human psychology

might be, or they would implicitly accept that facts about actual human psychology should really be the focus of some investigation. I think, therefore, that epistemologists should focus on empirically established facts about human psychology rather than aiming to accommodate all possible ways human psychology might be. But even if this argument is unconvincing, and one maintains that epistemologists should accommodate all possible facts about human psychology in their theories, Feldman's argument faces a further problem; it is implausible that epistemologists will be able to get a proper understanding of the space of possibilities of human psychology in the armchair.

We do not have to look far to find reasons to question the imaginative abilities of epistemologists. Just consider the empirically supported facts about human psychology discussed in section 3.2 (unconscious and automatic biases, successful use of heuristics, etc.). In two thousand years of armchair epistemology philosophers did not imagine these psychological facts and their relevance to epistemology. It was only after psychological investigation using methods specific to psychology that epistemologists considered the possibility that traditional views of rationality are challenged by people successfully using heuristic methods of reasoning. Yet this is clearly a relevant psychological possibility that ought to be considered—what can be a more relevant possibility about human psychology than an empirically supported fact?

The shortcomings of epistemologists' imaginations are thus not difficult to discover. But it should not be surprising that epistemologists have not been able to imagine the facts revealed by scientific psychology given that advances found in other sciences such as modern physics would not be expected to be imagined in the armchair. Fortunately there is a way that epistemologists can overcome the shortcomings of their imaginations. Unfortunately for Feldman it involves methodological naturalism.

Epistemologists can simply consider results from psychology whenever they are available, attempting to find out all information that is available about those possibilities that are definitely relevant to epistemology, i.e. the facts about human psychology. That way they will get the best available idea of the geography of human cognition rather than relying on their imaginations and ability to access the space of possibilities of human cognition.

The arguments I have presented in this section suggest that Feldman's assumptions 3 and 4 are both false. I have challenged assumption 3 by arguing that epistemologists are not interested in empirically supported facts only because they highlight possible ways human psychology might be. They are interested in psychological facts because these facts can contribute to epistemology by enabling epistemologists to get a better grasp of how people really can have knowledge than is available in the armchair. Then I have challenged assumption 4 by suggesting that even if epistemologists were interested in all possible facts about human psychology they would not find satisfaction in the armchair because they would not be able to imagine all relevant possibilities.

Conclusion

Methodological naturalism is in a good position at the end of this discussion. None of the arguments that Feldman presents to challenge the methodological naturalist claim that improvements have been made to epistemological theories because epistemologists have consulted results from psychology have succeeded. Feldman suggests that results from psychology are only relevant to epistemology if they reveal facts unavailable in the armchair. I have responded that results from psychology provide the easiest access to reliable facts about human psychology and that this provides good reason for epistemologists to consider them. He suggests that armchair epistemologists would be

able to access the relevant facts about human psychology in the armchair, but I have presented examples that suggest that some facts about human psychology relevant to epistemology are only available to those who are willing to embrace the scientific methods of psychology. Finally he suggests that epistemologists should be concerned with the possible ways human psychology might be, rather than only the ways that it actually is, and that possibilities can be imagined in the armchair. I have rejected the idea that epistemologists should be interested in possibilities and maintained that they should be interested in facts about human psychology and argued that they would only be able to guarantee that they consider all relevant possibilities if they consult results from psychology to overcome the limits of their imaginations.

All of these arguments point towards the conclusion that epistemologists should engage with results from psychology and modify or abandon their theory in response to empirically supported facts that provide a challenge to it. By taking results from psychology into consideration, epistemologists are able to develop theories that reflect how people really are able to obtain beliefs successfully. They will be able to ensure that their theories accommodate all and only relevant possibilities. And they will do so in an efficient way.

By considering in some depth these responses to Feldman we have found that there is good reason for thinking that epistemologists should engage with results from psychology. The examples of unconscious and automatic processes really do seem to provide a challenge to internalism, and the evidence that people can get more successful results by using heuristic methods of reasoning than by using the methods traditionally associated with rationality provides a real challenge to those traditional ideas. Generally, it seems that epistemologists should aim to present theories of knowledge (as

well as rationality and justified belief, etc.) that reflect how people are really able to obtain and revise their beliefs. If one is willing to accept this much then one should be willing to accept that facts about human psychology are relevant to epistemology. And hopefully the arguments presented in response to Feldman in this chapter show that if facts about human psychology are relevant to epistemology, epistemologists should use results from the scientific study of psychology when developing their theories.

Chapter Two: Methodological Naturalism and Its Critics, Objections and Replies

In chapter one I provided a defence of methodological naturalism. This is a controversial position. In this chapter I am going to address some of the concerns that are often expressed by philosophers critical of methodological naturalism and suggest that the position should not be as controversial as it has been considered to be. I will argue that if methodological naturalism is given a fair hearing it can be an attractive approach that can address the concerns of its critics.

1. Is does not imply ought

Perhaps the most widely discussed worry is that it is not possible to derive normative claims from descriptive ones. Advocates of naturalized epistemology suggest that scientific findings are relevant to epistemological claims, but it requires some work to establish how descriptive claims relate to normative ones. Jaegwon Kim's (1988) critique of Quine's 'Epistemology Naturalized' provides the most famous version of this criticism. Kim argues that the descriptive nature of Quine's position prevents it from being normative and hence from being epistemology at all. By scrutinising Quine's position, Kim's response, and how they relate to other versions of methodological naturalism, we shall see that it is perfectly consistent to advocate methodological naturalism *and* to accept that epistemology is essentially a normative discipline.

Quine (1969) defends his naturalized approach to epistemology as an alternative to foundationalist approaches, which he deemed to fail. According to foundationalism, beliefs are justified just as long as they are either basic beliefs or are based on a warranted inference from basic beliefs. Quine claims that the search for foundationalist justification has failed because basic beliefs cannot justify the beliefs founded upon them. He argues that basic beliefs always underdetermine the conclusions that are based upon them, just as observation underdetermines theory, so justification cannot come simply from the relationship between basic and non-basic beliefs. Quine's recommendation, in light of the failure of foundationalism, is that epistemologists should not seek an account of how justification is conferred on a belief and they should aim to get an understanding of the relationships that hold between evidence and the theories based on them.

> The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world. Why not just see how this construction really proceeds? Why not settle for psychology? Such a surrender of the epistemological burden to psychology is a move that was disallowed in earlier times as circular reasoning. If the epistemologist's goal is validation of the grounds of empirical science, he defeats his purpose by using psychology or other empirical science in the validation. However, such scruples against circularity have little point once we have stopped dreaming of deducing science from observation. If we are out simply to understand the link between observation and science, we are well advised to use any available information, including that provided by the very science whose link with observation we are seeking to understand (Quine, 1969: pp. 75-6).

Quine suggests that to understand how observation relates to theory in sciences we should use information from psychology or other empirical sciences about the psychological mechanisms that lead us to move from the observations and evidence that we have to accept the conclusions that we endorse. In Quine's view, then, epistemology

should be descriptive, it should be a part of psychology; the part of psychology which considers how theories come out of observational evidence.

It is Quine's concentration on describing psychological mechanisms that Kim criticises because he thinks that this is incompatible with the normative aspects of epistemology.

> Just as it is the business of normative ethics to delineate the conditions under which acts and decisions are justified from the moral point of view, so it is the business of epistemology to identify and analyze the conditions under which beliefs, and perhaps other propositional attitudes, are justified from the epistemological point of view. It probably is only an historical accident that we standardly speak of "normative ethics" but not of "normative epistemology". Epistemology is a normative discipline as much as, and in the same sense as, normative ethics (Kim, 1988: p. 383).

Epistemology, according to Kim, is a normative discipline. It is about how a belief can become justified. Its main aim is not to describe the processes that lead people to accept certain propositions because it aims to show how beliefs can be justified rather than simply stimulated by the evidence available. Quine's picture, as it is merely descriptive, doesn't have this normative element. It does not provide an account of how beliefs become justified, and how people ought to go about obtaining beliefs if they want to be justified. For Kim, a theory that fulfils the aims that Quine attempts to achieve will not really be an epistemological theory because it will lack the normative element:

> ...it is difficult to see how an "epistemology" that has been purged of normativity, one that lacks an appropriate normative concept of justification or evidence, can have anything to do with the concerns of traditional epistemology. And unless naturalized epistemology and classical epistemology share some of their central concerns, it's difficult to see how one could replace the other, or be a way (a better way) of doing the other. To be sure, they both investigate "how evidence relates to theory". But putting the matter this way can be misleading, and has perhaps misled Quine: the two disciplines do not investigate the same relation. As lately noted, normative epistemology is concerned with the evidential relation properly so-called-that is, the relation of justification—and Quine's naturalized epistemology is meant to study the causal-nomological relation.

> For epistemology to go out of the business of justification is for it to go out of business (Kim, 1988: p. 391).

Kim's criticism outlined here is often viewed as decisive against naturalized epistemology[19] and in my view his criticisms are decisive against Quine's version of naturalized epistemology. An epistemology that merely describes the psychological processes that lead people to obtain beliefs will not fulfil the ordinary aims of epistemology. It will not provide reasons for using particular methods of belief formation. It can provide an understanding of the consequences that follow from using a particular belief forming process, but it cannot tell us which consequences we should value and therefore which belief forming processes we ought to adopt. As long as one thinks that it is important to understand how a theory can be justified and to provide some guidance about how people ought to go about obtaining beliefs if they want them to be justified, Quine's approach simply will not do.

It is vitally important to note, however, that Quine's version of methodological naturalism, although highly influential on more recent versions, is not representative of methodological naturalism as a whole. To say that Kim's criticisms are decisive against Quine is not to say that they are decisive against methodological naturalism. To understand the full impact of Kim's criticisms on naturalized epistemology in general it is necessary to consider the relationship between Quine's position and those of other methodological naturalists. If all methodological naturalists were committed to the idea that epistemology should be descriptive and a part of psychology then Kim's argument would have a deep and wide impact. But all prominent methodological naturalists since Quine have presented accounts of the methodology to use in epistemology able to avoid the criticisms expressed by Kim. These approaches will be the subject matter of my

[19] This criticism of the Quinean project is often rehearsed. See, for example, Bonjour, 1994: p. 287.

broader thesis, and although I will ultimately provide some criticisms against each, in this section I will show how they can be defended against Kim's accusations. I will show that methodological naturalism can be normative, even if it has a descriptive element.

Let us begin by considering the version of methodological naturalism provided by Alvin Goldman. Goldman defends a naturalistic approach to epistemology which focuses on the normative project of developing an account of epistemic justification (e.g. 1976)[20]. He outlines a position that has a substantial descriptive component but holds onto the normative component that comes with using traditional philosophical methods. He argues that a belief is justified if and only if it is produced by a reliable cognitive system—a position that has become known as 'reliabilism'. His arguments in support of reliabilism focus on the concept of justification, so they are essentially normative; they attempt to provide an account of how we *should* go about obtaining beliefs. This approach is representative of Goldman's general view of the methodology that should be used in epistemology. Goldman (1993) argues that it is important to examine ordinary folk concepts of epistemic standards such as knowledge, rationality and justified belief. Only after we have evaluated our folk concepts can we hope to build upon and improve them. In his view, even methodological naturalists should respect the folk conceptions of epistemic norms. The naturalistic aspect of the epistemology should allow us to understand and expand upon our ordinary conceptions of epistemic norms,

[20] Kim recognises that Goldman's version of naturalism differs from Quine in a significant way, and acknowledges that his criticisms are levelled at Quine's naturalism rather views like Goldman's which he thinks differ significantly from Quine's:

> ...a salient characteristic of the naturalistic approach has already emerged, which can be put as follows: justification can be characterized in terms of causal or nomological connections involving beliefs as psychological states or processes, and not in terms of logical properties or relations pertaining to the contents of those beliefs.
>
> If we understand current epistemological theorising in this way, how closely related to Quine's conception of naturalized epistemology? The answer, I think, is obvious: not very closely at all. (Kim, 1994: p. 48)

but it should not replace our normative epistemic concepts. It is clear, then, that Goldman aims to keep a normative element in his epistemology, and that he is able to do so as well as proponents of more traditional epistemological methods. As long as epistemologists have distinctively normative concepts as their objects of study, the theories concerning those concepts should have a normative element[21]. Goldman is thus our first example of a methodological naturalist who respects the normative aspect of epistemology. Unlike Quine, he does not want to subsume epistemology under psychology, he merely suggests that philosophers should be willing to take scientific findings and theories into consideration when developing a view about what is required for knowledge, rationality or justified belief.

Edward Stein is another methodological naturalist with the aim of balancing the normative component of epistemology with a form of naturalism. In *Without Good Reason* he presents a picture of how people can and should go about evaluating the normativity of other people's reasoning. He claims that epistemologists who are interested in ascribing rationality should use a method of reflective equilibrium, taking as input intuitions about cases in which people form beliefs, the principles of theories of reasoning, and scientific findings about the cognitive limitations on human reasoning, to decide how to evaluate belief forming processes (Stein, 1996: p. 252-4). He claims that we should make adjustments to the intuitions we have about whether an act of reasoning is rational to fit with the principles of reasoning we find acceptable after some consideration and evaluation, and vice versa. We should make adjustments to each of these to fit with the scientific findings about human cognitive limitations. The outcome of the process of reflective equilibrium, with all these adjustments, will be a set of

[21] For more on the strengths and limitations of Goldman's approach to developing epistemological theories see chapter three.

standards to use to evaluate reasoning which are within the reach of ordinary human beings (because the limitations will have been factored into the process of reflective equilibrium). The picture is, as a result, a naturalized picture of rationality. Yet it is a naturalized picture that shares some normative features of traditional approaches: it is based upon intuitions about what is normative and takes intuitively acceptable principles about how people should obtain beliefs into consideration. This means that the naturalized picture of reasoning that Stein develops is a normative picture because it is influenced by our notions about the normativity of certain cognitive processes and their products. The introduction of scientific findings about human cognitive limitations will lead to adjustments to the normative principles that are accepted, but this does not mean that normativity will be removed from the picture. (For more on Stein's approach once again see chapter three.)

It is more difficult to show how the version of methodological naturalism defended by Hilary Kornblith can be reconciled with the normativity of epistemology. Kornblith's position is closest to Quine's because he argues for the replacement of traditional philosophical methods with the methods of the sciences. Kornblith (e.g. 2002) claims that knowledge is a natural kind. He claims that knowledge is the category of beliefs implicated in scientific theories about the selection and retention of cognitive systems (ibid., pp. 62-63). As such, he argues, the nature of knowledge can be understood by understanding the causal properties of the world which enable us to make successful inductive inferences about cognition. That is, it can be understood by considering how scientists categorise beliefs when they make successful inductions and provide successful explanations of the selection and retention of cognitive systems. Kornblith's meta-epistemology thus encourages epistemologists to search for a description of the

actual causal structure of the world. It might therefore seem that he, like Quine, is guilty of providing a descriptive and not a normative account of knowledge.

Kornblith denies, however, that his epistemology is stripped of its normative element. He claims that an epistemology that studies the causal structure of the world, including the successful behaviour produced by certain belief forming processes, can identify a normative category of beliefs. In his view, the category of beliefs identified by epistemologists using his method will include beliefs instrumental to achieving an animal's biological goals.

> Consider the cognitive ethologist's notion of knowledge. It gets purchase on the causal structure of the world because knowledge so conceived is conducive to fitness; such knowledge is instrumental in producing behavior that satisfies a creature's biologically given needs. Is it any wonder that we should value such a condition, that we should see it as worthy of pursuit? Knowledge is of extraordinary instrumental value, for it allows us to achieve our biologically given goals, as well as our more idiosyncratic individual goals, whatever those goals may be.
>
> ... The category of knowledge is able to play its normative role precisely because it plays the causal role it does; it is valuable because it provides the means by which animals may satisfy their needs, as well as their desires. One and the same category may do the work of both prescription and description (Kornblith, 2002: p. 160-161).

Kornblith has a fundamentally normative as well as a descriptive project; his project describes how epistemologists can discover a category of beliefs that have instrumental value for attaining certain goals. While this will be a descriptive project, it will have a normative element because the beliefs identified will be those that should be valued and that people should want. They will be beliefs that enable those who possess them to achieve their goals.

We shall see in chapter four that Kornblith's picture ultimately fails, but it does not fail because it has no room for a normative element. This means that each of the examples

discussed in this section demonstrates how methodological naturalism can be normative and that it can provide prescriptions about how people should go about forming beliefs. Methodological naturalists are committed to the idea that there should be some descriptive element to epistemology—epistemology should take into consideration the descriptions provided by scientists about how beliefs are formed—but this does not preclude them from presenting epistemological theories that prescribe ways of obtaining beliefs. Since Quine, those who defend naturalistic approaches have sought, as an essential part of their naturalistic project, to provide an account of how people should go about obtaining beliefs rather than merely describing how people obtain beliefs. The positions cannot therefore be dismissed on the basis of Kim's objection.

2. Ideals and Aspirations

In the previous chapter I argued that methodological naturalism is valuable and psychological results are relevant to epistemology because by taking these results into consideration it is possible to reflect the psychological reality of people who have knowledge. Epistemologists who are informed by results from psychology can develop analyses of knowledge, rationality and justified belief that reflect how people can revise their beliefs and can consequently provide advice about how people should improve their belief revision. Those who are not informed by psychological results will be liable to present inadequate theories, providing epistemic standards that are beyond the reach of ordinary thinkers and are never really met by any human being.

One potential criticism of this argument is that our epistemic notions and concepts can be worthwhile at the same time as being beyond the reach of ordinary thinkers and that it can be useful to have an idealised concept of knowledge even if people are never in the requisite state. According to this argument, theories could present standards for

epistemic notions like knowledge, rationality and justification and other epistemic notions that ordinary humans can't hope to achieve. But standards that humans cannot hope to achieve could be useful because they provide an ideal for people to aspire towards. People could aim to get beliefs that are as close as possible to fitting into each of the ideal epistemic categories. The imperative that comes from the normative theories could be an imperative to get as close as possible to the ideal. The normative aspect of any epistemological theory could come from prescribing standards we can aspire towards, rather than aims that we can actually achieve.

The inadequacy of this response can be understood by considering in some more detail the rationality case which was introduced in chapter one, which has been a particularly fruitful area of research for advocates of naturalistic approaches to epistemology. It has been established by cognitive scientists that our ordinary concept of rationality is far from the reality of how it is that we can achieve the best beliefs (e.g. Stich, 1990; Bishop, 2000; Samuels, Stich and Bishop, 2002). As a result of these findings, it is inappropriate to assume that the standard picture of rationality might even provide an appropriate ideal for people to pursue. As discussed in the previous chapter, people are not able to follow the principles of the standard picture of reasoning, according to which they check the consistency of each of their beliefs by applying logic and probability theory, because they are limited in their cognitive capacities and the information available to them. But this is not all that psychological studies tell us about how people can go about reasoning. They provide positive evidence about how people are able to achieve desirable results that suggest that people are able to make correct inferences by using what has been described as 'fast and frugal methods of reasoning' or 'heuristics'. Fast and frugal methods of reasoning do not meet the demands of the standard picture of

reasoning, because they require only limited information search. People using the methods consider only a restricted number of options when checking which of a number of new beliefs they should accept. Rather than considering all information relevant to a belief, they focus on certain readily available pieces of information to make a quick judgement. Because they do not consider much information, they cannot be thought to be applying logic and probability theory to check the consistency of any new beliefs with their prior beliefs. Yet the outcome of their reasoning is often the efficient, accurate acquisition of beliefs. Consequently, it seems right to conclude that a person can be rational without achieving what is described in the standard picture of reasoning. They can be rational using fast and frugal methods of reasoning because these are the best methods available to achieve accurate beliefs (see, for example, Stich, 1990; Samuels, Stich and Bishop, 2002).

An epistemologist might maintain that the standard picture of rationality provides aspirations that people could strive to achieve in their reasoning by arguing that the outcomes of fast and frugal methods of reasoning can only be called rational if they have the same outcomes as reasoning in accordance with the standard picture of rationality. *Prima facie* this argument has some force. The psychologists and philosophers who argue that humans can be rational using fast and frugal methods often establish this point by making a comparison between the results achieved by using fast and frugal methods of reasoning and the results that could be achieved by following the standard picture of rationality, if it were realistic. Fast and frugal methods are deemed as contenders for the status of 'rational' because they achieve results similar to those that would be achieved by people who successfully checked the consistency of their beliefs using formal models such as logic and probability theory. But this would be to miss the

point about why the naturalistic approach to epistemology is valuable. Methodological naturalism has normative force because it provides guidance about the *particular strategies and cognitive processes* that can be used to achieve beliefs with positive epistemological status. The methods of the cognitive sciences are particularly good methods to use to provide this advice because they can give specific advice about the processes that people should use rather than merely describing desirable outcomes. What the psychological findings from the cognitive sciences suggest, then, is that people who want to achieve certain desirable outcomes are best able to achieve these ends by departing quite radically from the standard picture of reasoning; by basing beliefs on limited information that is readily available rather than attempting to use a large amount of information to check the consistency of various beliefs.

What the rationality case thus suggests is that it is not possible to get around the criticisms of idealised epistemological theories by saying that they provide ideals to aim towards. Theories about knowledge, rationality and justified belief that are able to provide guidance for humans will take scientific findings from the study of human psychology into consideration. The idea that it might be useful to use idealised epistemological standards as targets to aspire to doesn't withstand scrutiny because epistemology will really benefit from highlighting principles of good reasoning that people can follow if they want to acquire the best beliefs possible. These principles will be significantly different from those presented within the standard, armchair picture of rationality.

I can foresee an objection to this approach that is worth addressing here. One might think that it begs the question about the relevance of psychological findings to epistemology to try to establish their relevance by considering results from

psychological studies about how people can go about reasoning. But I would argue that the circularity is virtuous rather than vicious. What I have attempted to show is how scientific findings and theories can positively impact upon the epistemological theories we develop. I have attempted to show how much poorer our epistemological notions would be if we did not exploit scientific theories and findings. Rather than begging the question about the merits of methodological naturalism when discussing the rationality debates, I have aimed to establish the merits of the approach by showing how progress can be made by using it and highlighted how much poorer alternative approaches would be.

3. The scope of the epistemological project

My defence of naturalistic approaches to epistemology so far has often focused upon the idea that only methodological naturalism can provide an adequate account of knowledge, rationality and justified belief that is within the reach of ordinary thinkers. One consequence of using scientific findings in this way is that it is possible to develop positive advice for people about which reasoning strategies to use. This aspect of the project of methodological naturalism has been described as the 'meliorative' (Kitcher, 1992) or 'ameliorative' project (Bishop & Trout, 2005) in epistemology (henceforth, ameliorative project). In my view, the ameliorative aspect of naturalized epistemology is one aspect that makes it stand out as exceptional and particularly worthy of pursuit. But some philosophers object to methodological naturalism on the basis that the ameliorative project isn't, in their view, a central part of epistemology. They argue that epistemology is about developing ways of assessing justification and that the ameliorative project is about applying the assessments developed through epistemology to actual cases of reasoning.

> ...it is presumably an important part of this general meliorative effort to provide critical assessments of inferential patterns for which purpose some knowledge of the relevant facts about such practice, psychological or otherwise, is clearly needed. All these points, however, have to do again with *applying* epistemological assessments to actual practice, not with how those assessments are themselves arrived at and justified (Bonjour, 1994: p. 291).

In light of these arguments it is worth considering what the scope of epistemology ought to be. It might be that methodological naturalists are simply too ambitious about what they think people from within their sub-discipline of philosophy ought to be doing. Should they concentrate on developing standards others can use to assess beliefs and belief forming processes rather than assessing these by considering results from psychology themselves? Should epistemologists accept that they have little to contribute to critically assessing what is discovered by psychologists and that they should simply supply standards of evaluation to use in critical assessments, rather than getting involved with assessing the implications of the empirical data gathered by scientists? Richard Feldman (1999) develops an argument along these lines:

> …information of two sorts is needed to assess the epistemic status of beliefs resulting from perception, memory, or testimony. We need to know what the general standards for such evaluation are. And we need empirical information about either the reliability of processes or the evidence people actually have about those sources. The latter is a merely empirical matter. Theorizing about the first topic is the province of armchair epistemology, and we haven't seen good argument for the role of psychology in identifying and clarifying these general standards. Empirical information is needed to determine if the standards are met in particular kinds of cases (p. 183).

Feldman and BonJour both suggest that there should be a division of labour amongst epistemologists and scientists; that epistemologists should concentrate on 'armchair methods' for identifying general standards, leaving scientists to gather empirical evidence about whether people are able to meet these general standards. These critics of methodological naturalism seem to be committed to the idea that the work of

epistemology is done when general standards of assessment are developed which can subsequently be used by psychologists.

In the chapter one, however, we found that much progress has been made in epistemology as a result of epistemologists considering results from psychology. To some extent, it seems that it doesn't matter whether epistemologists or psychologists consider how psychological results should influence our epistemological theories, but it seems valuable for someone to take our epistemological theories and psychological theories and combine their insights. It is only by combining such insights that the best possible epistemological theories can be developed.

What is important, then, is that someone engages with both the epistemological theories and psychological theories because the most complete account of what is required of knowledge, rationality and justified belief can only come about if people consider both inputs. At the present moment there are many epistemologists interested in these issues, and they seem well equipped to deal with both the conceptual and empirical side of the debate, so I would suggest that we should accept that this important work should be done by epistemologists. Moreover, there are many people currently working in epistemology who are developing epistemological theories that are inadequate because they fail to reflect the reality of human cognition. It seems only right that this should be prevented. They should be encouraged to expend their efforts on developing theories that do reflect psychological facts; it only seems sensible that their resources should be reallocated so that they take psychological facts on board given that, as I showed in chapter one, by considering results from psychology it is possible to develop epistemological theories most efficiently and effectively.

4. Self-refuting naturalism?

Laurence Bonjour (1994) argues that the critical position that methodological naturalists take with regard to the use of conceptual analysis and a priori justification, and the claim that any theory based on these methods alone will be incomplete, means that the position is self-refuting[22]:

> ... if, as the naturalist claims, there are no a priori reasons for thinking anything to be true... the inevitable result is that we have no reason for thinking that any of the beliefs whose content transcends direct observation are true.
>
> This is epistemological disaster in itself, but a further consequence is that the vast majority of the claims about the nature of the world, the nature and reliability of human psychological processes, etc., upon which naturalized epistemology so lovingly focuses, are things that we have no reason at all for thinking to be true- as, indeed, are the very theses that epistemology must be naturalized or that traditional epistemology is untenable (together with all claims of any sort). In this way, naturalized epistemology is *self-referentially inconsistent*: its own epistemological claims exclude the possibility of there being any cogent reason for thinking that those claims are true (ibid. p. 296).

Bonjour suggests that the claims in favour of methodological naturalism are 'self-referentially inconsistent' because they provide a priori justification for the claim that a priori justification is unreliable. As the importance of methodological naturalism can only be established by making claims that go beyond the evidence provided by the sciences, the normativity of methodological naturalism appears to derive from a priori claims about how epistemologists ought to proceed that are not based on empirical evidence but rather on a priori theorising.

How can the methodological naturalist respond to this claim? She could attempt to show that many of the moves made within methodological naturalism are direct responses to

[22] Similar arguments have been presented by Bealer (e.g. 1998), Jackson (1998), Kaplan (1994), and Siegel (1984). A more detailed naturalist response to these arguments can be found in Kornblith, 2002: chapter one.

psychological findings and therefore have empirical rather than a priori support. For example, the suggestion that there are problems with internalism discussed in the previous chapter is supported by empirical findings about how people revise their beliefs. One could argue that these empirical facts demand that epistemologists respond and revise their theories to reflect them. But this argument does not work because ultimately the claim that we should attempt to develop standards for knowledge, rationality and justified belief that are within the reach of ordinary thinkers, which is central to the naturalist argument, is underdetermined by observational experience. It is a normative claim that goes beyond the descriptive claim.

There is, however, another response available to the naturalist. She might accept that her position has some a priori component while maintaining that epistemology also requires some a posteriori element. This is perfectly consistent with the methodological naturalist views expressed in this thesis so far. The methodological naturalist is committed to the rejection of epistemological theories which are wholly a priori because they think that scientific theories and findings can contribute to our understanding of what is going on when people have knowledge, rationality and justified belief. We are committed to the idea that any account of knowledge, rationality and justified belief should take results from the scientific study of psychology into consideration, or else they will be incomplete. It doesn't follow that we must be committed to the idea that there should be no a priori element in epistemology[23]. As long as we methodological naturalists accept that there can be some a priori element to epistemology our position is not self-refuting.

[23] As Phillip Kitcher (1992) is interpreted as doing.

5. Is naturalism dependent upon reliabilism?

So far I hope to have persuaded the reader in chapter one that there are some good reasons for thinking that a naturalized epistemology is preferable to a non-naturalized one; we can improve our epistemological theories by taking scientific theories and findings into consideration. In this chapter I hope to have assuaged some of the worries that people often raise with regard to methodological naturalism. However, one might maintain that methodological naturalism is intrinsically associated with a particular position within epistemology, namely reliabilism, and that those who are committed to some other epistemological position could be satisfied with a non-naturalistic meta-epistemology[24]. This argument has been suggested by Feldman (1999) in response to naturalistic arguments such as those presented by Kitcher (1992). Kitcher argues that naturalism re-entered into the philosophical tradition with the development of causal theories of knowledge, 'Analyses of the concept of knowledge (and, later, that of justification) were no longer confined to specifying the logical relations among propositions believed by the subject but could take into account the processes, including inevitably the psychological sub-processes, that causally generates states of beliefs' (p. 60). According to causal theories of knowledge, what justifies belief is its causal history; it must be caused by the fact known. Causal theories of knowledge have lead to the development of theories referring to the reliability of psychological processes (see Goldman, 1976). According to the view expressed by Kitcher, then, the re-introduction of theories such as reliabilism made naturalism seem attractive again, or, as Brandom put it, 'part and parcel of reliabilism's appeal in epistemology is […] the idea that reliabilism provides at least the raw materials for a naturalized epistemology- one that

[24] Feldman (1999) suggests that naturalism might have a better status when it comes to reliabilism than other positions in epistemology such as evidentialism, although he ultimately concludes that neither reliabilism nor evidentialism would benefit from a naturalistic approach.

will let is exhibit states of knowledge as products of natural processes fully intelligible in broadly naturalistic terms' (Brandom, 2000: p. 111). It is worth considering the possibility, then, that the appeal of methodological naturalism in epistemology is intrinsically linked with the appeal of reliabilism, so that anybody who rejects reliabilism thereby rejects methodological naturalism.

Methodological naturalism and reliabilism do fit neatly together—they are mutually reinforcing[25]. Reliabilists suggest that something is knowledge or justified belief if and only if it is produced by a reliable belief forming process. Science comes into the picture because it can tell us which belief forming processes are reliable. It can tell us, for example, whether our reflective belief forming processes are reliable. This, in turn, tells us whether the beliefs we acquire through reflection constitute knowledge or justified belief. I have argued in chapter one and the present chapter that methodological naturalism is valuable because epistemologists should search for epistemological theories which are psychologically plausible. This implies that reliabilists should be methodological naturalists: they should take into consideration how it is psychologically possible for people to obtain reliably produced beliefs. But it does not imply that people who reject reliabilism should also reject naturalism.

We can see that it is possible to reject reliabilism without also rejecting naturalism by considering how the evidentialist might embrace naturalism. According to evidentialism, 'the epistemic justification of a belief is determined by the quality of the believer's evidence for the belief' (Feldman & Conee, 1985). Whereas according to reliabilism the justification of a belief comes from the reliability of the cognitive

[25] This point will become especially important in chapter four in the discussions of Hilary Kornblith's methodology.

processes used to produce it, according to evidentialism the justification of a belief comes from the evidence upon which the belief is based. My suggestion is that a person who accepts the main tenets of naturalism, as described in this chapter, can consistently also accept evidentialism; she can combine the two views. A person who thinks that the evidence a person has determines whether she is justified in holding a belief can maintain that scientific findings and theories should be used to develop epistemological theories about what evidence is required for justification.

To illustrate this point let us consider once again the findings presented in the rationality debates by Gerd Gigerenzer and his colleagues, about how people go about reasoning (e.g. Gigerenzer and Todd, 1999; Gigerenzer and Selten, 2001). The findings suggest that people are able to make successful inferences by using reasoning strategies that involve the consideration of only a small amount of relevant information. My suggestion is that someone who defends evidentialism, but also accepted the general arguments in favour of naturalism, could develop a version of evidentialism using methodological naturalism and taking into consideration findings such as Gigerenzer's about the sorts of evidence people can use to make successful inferences. Rather than just advising a person that their beliefs will be justified as long as the beliefs are supported by all the evidence that they have available, the evidentialist could provide the advice that people should aim to obtain beliefs based on certain types of evidence, i.e. the types of evidence that psychological findings suggest that people can make good use of, such as the evidence that a particular company is recognisable. It is not only reliabilists, then, who can benefit from being committed to methodological naturalism. Evidentialists can also benefit from taking a naturalistic approach to epistemological theories.

Conclusion

In my view, epistemologists should be methodological naturalists. They should use results from psychology when developing their epistemological theories to ensure that their theories reflect the psychological states that people can really be in when they have knowledge, are rational and have justified beliefs. However, methodological naturalism is controversial. Most epistemologists continue to solely use armchair methods rather than engaging with scientific findings. I think that the reason that methodological naturalism is so controversial is because of various misconceptions that there are about it. These misconceptions are reflected in the criticisms I have outlined and responded to in this chapter. In addressing some of these criticisms I have pointed out a number of these misconceptions with the aim of showing that methodological naturalism ought not to be so controversial if properly understood. With a proper understanding of the variety of positions all included under the umbrella of methodological naturalism, of what exactly methodological naturalists are and are not committed to and of the role that methodological naturalism has in describing the strategies that people should use to revise their beliefs, methodological naturalism should be an attractive methodological position.

Chapter Three: The Trouble with Traditional Methods of Epistemology

Introduction

So far my arguments have focused upon the merits of methodological naturalism as a general methodology. I have suggested that once it is properly understood, methodological naturalism should seem an attractive prospect and that other, non-naturalistic approaches to epistemology are inadequate in comparison. In the current chapter the focus will be switched onto some particular versions of methodological naturalism. In it I shall begin the search for the best possible account of the methodology to be used in epistemology consistent with methodological naturalism.

The accounts of the methodology to be used in epistemology that will be the focus of this chapter attempt to integrate traditional methods of analytic epistemology with an approach using results from the study of human psychology. They integrate conceptual analysis or reflective equilibrium into a version of methodological naturalism, suggesting that successful epistemological theories can be developed by making adjustments to the results produced by traditional epistemological methods. The fact that these approaches have been presented by methodological naturalists might seem surprising given that, as has been suggested in previous chapters, methodological naturalism is often presented as an alternative to the traditional methods of analytic epistemology, such as conceptual analysis and a priori theorising. The methods of traditional analytic epistemology are often categorised under the umbrella term 'armchair methods', as distinguished from the methods of methodological naturalism

which involve the study of the sciences. Methodological naturalism is often defended on the basis of criticisms of armchair methods. However, some methodological naturalists do maintain that the traditional methods of analytic epistemology should be integrated into a naturalistic picture. Alvin Goldman argues that epistemologists should engage in conceptual analysis as a part of the methodological naturalist project. Edward Stein suggests that theories of rationality should be based upon the method of wide reflective equilibrium, with an added naturalistic component. The purpose of this chapter is to evaluate these approaches.

I shall consider some of the arguments against traditional methods. I shall show that a compelling case can be developed against armchair methods based on these arguments, one that should be particularly compelling for methodological naturalists. Then I shall examine the prospects for Stein and Goldman's attempts to integrate traditional methods with methodological naturalism. We shall see that there are some obvious attractions to epistemological methodologies that retain some element of traditional epistemology, not least that the methods of traditional epistemology are able to provide some degree of continuity with folk concepts, but that ultimately the problems with traditional methods are not avoided by the attempts available in the literature to integrate the traditional with the naturalistic.

1.1 The proper target of epistemology

The first criticism of armchair methods I will consider is that armchair epistemologists take concepts as their targets when they should not. The idea that traditional analytic epistemologists are engaged in conceptual analysis or at least in the study of their concepts is supported by the prevalence of thought experiment methodology. Traditional analytic epistemologists often engage in the practice of presenting examples

and counterexamples to epistemological analyses with the aim of establishing whether an analysis is correct or incorrect. They aim to induce intuitions or a priori responses to these thought experiments and use the responses as evidence for or against an analysis. So, for example, some epistemologists used to think it is necessary and sufficient for knowledge that a person has justified true belief. Then Gettier famously presented some counterexamples to this claim (Gettier, 1963). For instance, he described a scenario in which two people had applied for a job. Smith has been told by the company president that Jones will get the job. He has also counted the coins in Jones's pocket and found that there were ten. He therefore forms the justified belief that the man who will get the job has ten coins in his pocket. As it happens, Smith actually gets the job. But he also has ten coins in his pocket. So Smith seems to have had a justified and a true belief, but doesn't seem to have knowledge. Readers of this Gettier example are asked to engage with the thought experiment and consider whether they think that it is appropriate to describe the conditions found in the scenario as knowledge. They are asked 'is it appropriate to apply the concept of knowledge in this case?'. Gettier's suggestion is that philosophers should not be inclined to apply the concept of knowledge under the conditions described in the scenario, and that this implies that people can have justified true belief without having knowledge. Gettier's methodology is characteristic of much of traditional analytic epistemology: epistemologists engaged in it use thought experiment methodology to seek an understanding of the necessary and sufficient conditions for the application of a concept by considering whether it seems appropriate to apply a concept under various stipulated conditions.

Two prominent methodological naturalists have defended their methodology by arguing that the main problem with armchair methods is simply that they have the wrong target.

While armchair epistemologists aim to characterise our epistemological concepts, in Hilary Kornblith's and Stephen Stich's view, they should aim to characterise the real phenomenon of knowledge, rationality and justified belief that is beyond our concepts. This claim is articulated by Hilary Kornblith and Stephen Stich in the following passages:

> 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 ethics is the right and the good, not our concepts of them. The subject matter of philosophy of mind is the mind itself, not our concept of it. And the subject matter of epistemology is knowledge itself, not our concept of knowledge. (Kornblith, 2002: p. 1)

> There is a substantial anthropological literature aimed at describing the concept of disease as it is used in various cultures. And if you are interested in how people conceive of disease, this is the place to look. But if you are interested in what disease is then it is biology or medicine you should be studying, not cognitive anthropology. An entirely analogous point could be made about gold, or space, or mass, or heredity. If you want to know how people conceive of them, then the description of common sense concepts or knowledge structures is the project to pursue. But if you want to know what gold, or space, or mass, or heredity is really like, then you should be studying chemistry or physics or genetics. (Stich, 1992: p. 251)

These quotes suggest that Kornblith and Stich are committed to the idea that the subject matter of philosophy cannot be reached through conceptual analysis because all we will come to understand through this method is what humans conceive of. While this might be interesting to anthropologists or ethnographers, philosophers ought to be interested in the items in the world that our concepts relate to. And to understand the nature of these items philosophers should take into consideration what is said by people working in other disciplines. For an understanding of what knowledge is, epistemologists should, in the view of Kornblith and Stich, turn to cognitive science rather than to conceptual analysis. Imagine what the success criteria for conceptual analysis in epistemology would be: it would be an analysis of our concept of knowledge. For Kornblith and Stich

the success criteria of epistemology should be something else: it should be an account of what knowledge actually is.

To argue his case, Kornblith (2002) focuses on the positive nature of an epistemological category, knowledge. He argues that knowledge is a natural kind; a natural category that can be discovered by the sciences (for more on this see chapter four). He claims that epistemologists should pursue an understanding of its nature in the same way a rock collector searches for items sharing particular characteristics. They should begin with certain clear cases of knowledge and then work out which common characteristics are shared by all the cases.

> What we are doing, as I see it, is much like a rock collector who gathers samples of some interesting kind of stone for the purpose of figuring out what the samples have in common. We begin, often enough, with obvious cases, even if we do not yet understand what it is that provides the theoretical unity to the kind we wish to examine. Understanding what the theoretical unity is is the object of our study, and it is to be found by careful examination of the phenomenon, that is, something outside of us, not our concept of the phenomenon, something inside of us. In short, I see the investigation of knowledge, and philosophical investigation generally, on the model of investigations of natural kinds (Kornblith, 2002: p. 11).

For Kornblith, the process of identifying knowledge involves the study of the natural phenomenon of belief formation. Once we have been able to identify some clear cases of knowledge, we can extrapolate out from the clear cases to gain an understanding of which features are constitutive of the kind. Then, once we know the kind-constitutive properties, we can identify other phenomena that are a part of that kind. It is not concepts that are the target of this process but instead the real phenomenon of belief acquisition and revision that underlies knowledge, rationality and justified belief. For Kornblith, an epistemologist who focuses on the analysis of epistemological concepts will, as a result, risk presenting a picture of knowledge that misrepresents the real

natural phenomenon. So, in Kornblith's view, it is the natural kind knowledge rather than our concepts thereof that should be the target of epistemology.

The psychological studies presented in chapter one of this thesis provide good reason to think that Kornblith is correct that our epistemological concepts fail to reflect the reality of what goes on when people successfully develop and revise their beliefs. For example, the concept of knowledge as requiring only partial access to the basis for our beliefs has come under some pressure in light of experimental findings suggesting that people are often able to access some of the bases of their beliefs through reflection even when their actions are ultimately decided by implicit biases. Meanwhile the concept of rationality that associates it with the preservation of consistency among beliefs has come under some pressure in light of findings about how people can use shorthand rules of thumb to great success in their reasoning. The idea that epistemic justification should be associated with epistemic responsibility has also been criticised on the basis that people can achieve desirable results by using less responsible methods. We therefore have good reason for thinking that something beyond our epistemological concepts should be the target of epistemology; that is, we have reason to think that our pre-theoretical epistemological concepts fail to reflect the reality of successful human cognition and that epistemologists should aim to understand that reality. Whether we should adopt Kornblith's view of knowledge as a natural kind will be the subject of chapter four, so I will not dwell upon it now. But something certainly seems right about how Kornblith switches the focus of epistemology onto the phenomena of belief revision and away from our concepts. As our concepts fail to accurately represent how people can successfully revise their beliefs, and scientific theories are better able to represent this,

epistemologists who are interested in developing a picture of how positive appraisal should be associated with beliefs should be interested in the latter and not the former.

Stephen Stich (1990) bases his criticisms of conceptual analysis not so much upon the merits of looking beyond our concepts but on the problems associated with conceptual analysis. He claims that epistemologists should not aim to study their concepts because concepts are merely 'idiosyncratic and arbitrary' (p. 94) products of our culture. He points out that the epistemic concepts that we have are only some among a number of possible concepts that we might have had, and that 'there is no obvious virtue that distinguishes our concepts from the alternatives, apart from the fact that we happen to have inherited them' (ibid.). He suggests that epistemologists who aim to analyse their concepts are guilty of a form of epistemic chauvinism.

> If the principle reason that our evaluative epistemic concepts, concepts like rationality and justification, stand out from the crowd is that they happen to be the evaluative notions passed on by our language and culture, it's hard to say why anyone but an epistemic chauvinist would much care whether his cognitive processes were sanctioned by those notions (Stich, 1990: p. 130).

For Stich, then, concepts are inadequate sources of evidential support for a theory because of their origin. They should not be given an elevated status in epistemology because we could have ended up with quite different concepts. In his view, epistemologists should not accept an epistemological theory simply because it matches their concepts. If they did accept an epistemological theory for this reason they would be guilty of epistemic chauvinism; that is, they would be dismissive of other concepts that disagree with their own, accepting that their concepts are better simply because they are their own. For these reasons, Stich suggests that the project of epistemology should be to focus upon the real phenomena of belief formation revealed by results from psychology rather than to focus on epistemological concepts.

The arguments presented by Kornblith and Stich present a compelling case to support the idea that methodological naturalists ought to worry about the target of traditional analytic epistemology and the armchair methods it uses. As I have argued in previous chapters, methodological naturalists should be interested in providing advice for thinkers about how to revise their beliefs. To do so, they should reflect the realities of human cognition; including how people can successfully obtain and revise their beliefs. Kornblith and Stich highlight the possibility that our epistemological concepts might not reflect the cognitive reality. They might simply reflect misguided, arbitrary and/or idiosyncratic epistemological principles that have become embedded in our culture. As a result of these observations, methodological naturalists should evaluate conceptual analysis with a critical eye. They should be suspicious that those who take concepts as a suitable target for their investigations will present theories which fail to reflect the best way for people to revise their beliefs. Results from the sciences can challenge some of our most deeply embedded concepts of what it is to have knowledge, be rational, or be justified in our beliefs. It is clear that conceptual analysis is inadequate in these cases because pure analysis of our concepts would not have revealed the facts that challenge them.

1.2 Intuitions as the evidential base of traditional methods

The arguments of section 1.1 have focused on whether concepts can be an appropriate target for epistemologists. In this section I will focus on the evidence that is provided for conceptual analysis by intuitions. I will suggest that even if methodological naturalists were willing to accept that concepts can and should be the target of analysis, we should worry that intuitions are the source of epistemological theories that adopt the methods of traditional analytic epistemology.

In recent years the traditional methods of analytic epistemology have come under fire because of the extent to which they depend upon intuitions. Recall the earlier description of the method of conceptual analysis and its use of thought experiments. Thought experiments are often used in epistemology, and when they are used the intuitive responses that people have to the experiments are taken to be evidence for or against an epistemological theory. There has been a large body of research into the nature of the intuitions that are the psychological responses involved when philosophers respond to thought experiments and use similar methods (see Nagel, 2007 for a summary). However, unfortunately, this body of research has yielded no consensus. So, like other philosophers working in this area, I am going to work with a very rough characterisation of what intuitions are. I take intuitions to be immediate responses to philosophical questions. They are not easily overridden. And, most importantly, they are taken to be evidence for philosophical positions. They might be similar to other modal judgements (Williamson, 2007), or they might be distinctive 'intellectual seemings' (Sosa, 1998), but what is important is that they are easily identifiable immediate responses to philosophical thought experiments or problems that are taken to be evidence for a philosophical position.

The use of intuitions in epistemology has been criticised on a number of bases. Some of the most vocal opposition to the use of intuitions in epistemology has come from those influenced by Stich's idea that the study of our concepts could be epistemically chauvinistic (see section 1.1). In recent years, additional support for the claim that we should worry about cultural and social variation of intuitions has come from the ranks of experimental philosophy. Stich's prediction that people from different socio-cultural backgrounds are likely to have different intuitions has found support in results from

experiments probing the intuitions people have to philosophical thought experiments. For example, Weinberg, Nichols and Stich (2001) presented participants with a number of thought experiments familiar to epistemologists, including one to test their intuitions about Gettier cases. Gettier cases were discussed in section 1.1, they are cases in which people seem to have justified beliefs which are true but which also seem, at least to a western audience, not to be knowledge. Weinberg, Nichols and Stich developed the following version of Gettier cases and presented it to undergraduate students at Rutgers University.

> Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American Car. He is not aware, however, that her Buick has recently been stolen, and he is not aware that Jill has replaced it with a Pontiac, which is a different American car. Does Bob really know that Jill drives an American car, or does he only believe it?
>
> REALLY KNOWS ONLY BELIEVES
>
> (Weinberg, Nichols and Stich, 2001: p.20)

This example aims to probe people's intuitions about whether Gettier-style cases are knowledge. The aim of the experimental philosophers was to see whether people from different cultural backgrounds share intuitions about thought experiments such as the Gettier example. They took students from Western and East Asian backgrounds from within Rutgers University and asked them to respond to the example.

The outcome of the studies undertaken by Weinberg, Nichols and Stich suggest that there is variation in the intuitions of people from different cultures. Here is how the authors of the study describe its results,

> The striking finding in this case is that a large majority of Ws [Westerners] give the standard answer in the philosophical literature, viz., "Only believes." But among EAs [East Asians] the pattern is actually *reversed*! A majority of EAs say that Bob really knows (ibid.).

Out of a sample of 66 Western participants, 49 or approximately 74% said that Bob only believes, compared with a sample of 23 East Asian participants, of whom only 9, or approximately 39% said that he only believed. 60% of East Asian participants rejected the standard response to the thought experiment found in Western philosophical literature, saying that the character described in a Gettier case has knowledge. Results such as these, suggesting that there is variation in intuitions, were replicated across other epistemological thought experiments, such as experiments usually used to support reliabilism, and in other studies the same experimental philosophers found similar variation in intuition across social groups within Western society (Nichols, Stich, and Weinberg, 2003). The investigations of experimental philosophy are still in their early stages but the data do give reason to take the epistemic chauvinism complaint seriously. They suggest that philosophical intuitions can be influenced by social and educational background and therefore should be treated as evidence not about philosophical truths but just about the values embedded in our culture and educational systems.

The epistemic chauvinism complaint has been supplemented by a number of other worries about the use of intuitions in philosophy. As Jennifer Nagel describes, some of the literature in epistemology using traditional intuition-driven methodology suggests that intuitions about potential cases of knowledge can shift based on the context in which beliefs are evaluated:

> Difference in the alternatives made salient can also cause epistemic intuitions to waver: for example, just after a discussion of misprints in airline schedules or last-minute changes in flight plans, one might be more reluctant to judge a subject as having knowledge on the basis of a quick glance at a schedule (Nagel, 2007: p. 798).

In this passage, when Nagel describes how intuitions about knowledge waver, she is concerned with the intuitions underlying contextualist and relevant alternatives

approaches to knowledge. According to these views, people's attributions of knowledge are context-dependent. Some philosophers, for example, maintain that it is the stakes at play in a situation which determine whether a belief is classified as knowledge: the higher the stakes, the less likely a person is to attribute knowledge. Others argue that whether a person judges a belief to be knowledge will depend upon which relevant alternatives are available: where there are many alternatives that cannot be eliminated—such as that the flight schedule a person glances at is inaccurate—people are less willing to attribute knowledge than at other times when it is easier to eliminate relevant alternatives. Nagel suggests that the 'shiftiness' of the intuitions described by contextualists is indicative of problems with intuition-driven methodology in general. Because intuitions about knowledge are influenced by context, one might think that they are inappropriate as sources of epistemological theories. However, some philosophers would disagree. For example, Goldman and Pust (1998/2002) suggest that it should not be surprising that epistemological intuitions are shifty because intuitions and the epistemological concepts that they reflect are context sensitive. Evidence of shiftiness of intuition might be viewed as providing support for contextualist views that suggest that knowledge attributions are context-sensitive.

This response to the evidence that intuitions are shifty doesn't seem applicable to all the examples used to establish the 'shiftiness' of intuition, however. Recent studies undertaken by experimental philosophers suggest that intuitions can be influenced by factors far less plausibly relevant to epistemological theorising than the context in which a belief described in a thought experiment is evaluated. In the examples described in support of contextualism, relevant alternatives accounts, and other similar views, the shiftiness of intuitions follows as the result of the details of the examples. If people are

more willing to attribute knowledge to a person in a high stakes situation than in a low stakes situation, the attribution that they make will be determined by the stakes in the situation. But Swain et al. (2008) present findings that suggest that philosophical intuitions can be influenced by factors which seem as if they should be irrelevant to the stipulated example.

Swain et al. present findings suggesting that people's intuitions about the thought experiments used to criticise reliabilism vary according to the order in which the thought experiments are presented and considered. Whether people are willing to attribute knowledge to the individuals described in a thought experiment can depend upon the thought experiments presented before and after it. This means that a single thought experiment can elicit variable responses, where the variation is caused by something other than the details of the situation described.

One might think, once again, that the ordering effect that Swain et al describe is a consequence of how participants are considering relevant alternatives to the cases described. When they are considering whether a person described in a thought experiment has knowledge they make a comparison between beliefs about the prior and present scenarios. But when these findings are considered alongside those presented by psychologists studying how people form moral judgements, the shiftiness of intuitions becomes even more worrying and the claim that it is a positive thing less plausible. Schnall et al. (2008) present findings suggesting that people's moral judgements of actions that they find disgusting are less severe if they have had their 'concept of cleanliness activated' or if they have physically cleansed themselves (p. 1219). Perhaps most striking is the finding that people who have been made to watch a physically disgusting scene from a film tended to make less severe moral judgements if they had

washed their hands between watching the film and being asked to provide a moral evaluation. For current purposes the most important point to note is that the study seems to provide evidence that moral intuitions are subject to the kind of 'shiftiness' that it has been suggested influences epistemological intuitions. It suggests that moral intuitions are influenced by factors that ought to be irrelevant to the evaluations; it is highly implausible that whether an evaluator has cleaned his or her hands ought to be relevant to her evaluation of the morality of an action. As we have seen that epistemological intuitions can be affected by factors such as the ordering effect, these findings suggest that similar sorts of irrelevant factors can influence our epistemological intuitions.

Those who criticise the use of intuitions in philosophy maintain that all epistemologists ought to be concerned about the epistemic chauvinism complaint and the apparent shiftiness of intuitions, but methodological naturalists ought to be particularly concerned about whether the intuitions they use are trustworthy. Why is this? Because they are committed to the idea that we ought to take empirical evidence and the theories based upon it into consideration when developing philosophical theories relating to the mind. This should include meta-philosophical theories about how people should go about developing philosophical theories. Philosophical theories are produced by human beings so any account of how they should be developed should reflect the nature of the minds of humans and how they can go about successfully doing philosophy. If there is psychological evidence of the sort presented by Schnall et al., or empirical studies of the sort presented by Swain et al., or even simply strong anecdotal evidence suggesting that these intuitions are untrustworthy, it seems that the naturalistic approach demands that concerns about the trustworthiness of intuitions be addressed.

Results from both experimental philosophy and psychology thus provide reason to think that the intuitions used by epistemologists who examine thought experiments could be affected by factors that are seemingly irrelevant to the philosophical questions. They present a challenge for epistemologists to show that the intuitions that they use as evidence are unlikely to be subject to undesirable influences. Arguments presented by Robert Cummins provide even more reason for worrying about the use of intuitions in epistemology.

Cummins (1998) describes and criticises the use of intuitions in philosophy from a theoretical rather than an empirical angle. He claims that the use of intuitions as evidence for a position is problematic because it is impossible to calibrate intuition. He compares the use of intuitions to the use of scientific data in support of a theory. He claims that scientific data provides a stable basis for a theory because it can be calibrated; that is, evidence produced by one instrument can be checked against evidence produced by another. With intuitions, on the other hand, it is generally taken to be that there is no way to check for correctness. Intuitions are supposed to provide a basic evidential source for a theory, a source that cannot, and also need not, be given independent verification. If, for example, an intuition suggests that we can have justified true belief without having knowledge, this provides evidence against the idea that justified true belief is sufficient for knowledge. Philosophers do not calibrate the intuition through some other non-intuitive source. Philosophers who use intuitions seem to assume that calibration is not necessary, and for Cummins this is a serious problem with intuition-driven methods.

One might worry that Cummins' argument runs the risk of leading to radical scepticism. Many of the beliefs that people ordinarily form are based on judgements that are not

calibrated. For example, it is not unusual to obtain beliefs about the weather from the newspaper without checking them against other sources of information. We would not want to deny that all beliefs of this sort, which are derived from a single source without calibration, lack justification. So why should we think that intuition-driven theories are troublesome because intuitions are not calibrated?

To understand why it is important that intuitions cannot be calibrated we can note how this analogy breaks down. In the newspaper case it is possible to calibrate the mechanisms used for belief acquisition, even if it is not possible to calibrate the judgements or beliefs themselves, whereas in the intuition case it is not possible to do so. We have good reasons for thinking that the newspaper is a good source of evidence about the weather because we have been able, in the past, to check the prediction made by the paper with the actual weather. We will have been able to calibrate the results provided by the newspaper and the actual weather we have experienced. We might not calibrate our beliefs about the weather in every instance but we are able to make checks on the mechanism of reading the weather in the newspaper which is underlying our beliefs about the weather. In contrast, there seems to be no way to calibrate the mechanisms underlying the development of intuitions that are used in philosophy. Because they are brute sources of evidence we cannot ever check whether they accurately reflect some independent reality, so we cannot check whether the mechanisms underlying them are generally reliable.

Overall, then, the prospects for traditional intuition-driven epistemology are bleak. There seems to be little reason for thinking that our intuitions can provide access to an accurate picture of what it is to have knowledge, rationality and justified belief. To assume that our intuitions are correct would be to ignore the empirically supported fact

that people's philosophical intuitions are influenced by factors that seem to be irrelevant to philosophical skills, such as cultural or ethnic background or even the cleanliness or comfort of the setting in which a judgement is made. There is a genuine demand for methodological naturalists who maintain that intuition-driven methodology provides the best way forward to provide a strong defence of the use of intuitions against these criticisms.

2. Integrating traditional methods with science

If I am right, then, methodological naturalists have a particular demand upon them to provide reasons why conceptual analysis and intuition-driven philosophy in general are valuable, that is, if they wish to use these methods. This means that those methodological naturalists, namely Alvin Goldman and Edward Stein, who hold onto the idea that traditional methods should be used in epistemology, maintaining that methods like conceptual analysis should be integrated into an approach that is informed by results from the cognitive sciences, have a challenge ahead. The goal of subsequent sections is to show how they fail to meet this challenge.

2.1. Stein's naturalized picture of rationality

Edward Stein presents his version of methodological naturalism while discussing the rationality debates in the cognitive sciences, which have been discussed a number of times in previous chapters. Let us have a quick recap of the nature of the debates. In the rationality debates psychologists have presented findings suggesting that humans cannot be rational according to the standard picture of rationality adopted in philosophy and psychology, which associates rationality with checking the consistency of one's beliefs while following the principles of logic and probability theory when revising one's beliefs. These findings have stimulated discussion about which alternative standards

should be used to assess rationality because it seems inappropriate to use standards that are beyond the reach of human beings due to the very nature of their cognitive capacities. Stein's answer to this question is that standards for rationality should be developed by using a version of reflective equilibrium. He suggests that those who want to develop a theory of what is required for rationality ought to take intuitive principles of rationality, intuitions about cases of reasoning, and philosophical theories and scientific theories from the cognitive sciences showing how people are able to go about revising their beliefs, putting all of them into a process of reflective equilibrium. The output of this process will be principles of good reasoning or rationality.

To understand Stein's position it might be helpful to get an understanding of what exactly reflective equilibrium is. Nelson Goodman (1955) originally described the method as a way for particular inferences and rules of inference to become justified. He suggested that any inference can become justified by bringing it into agreement with a rule of inference that one finds acceptable.

> The point is that rules and particular inferences alike are justified by being brought into agreement with one another. A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend. The process of justification is a delicate one of making mutual adjustments between rules and accepted inferences; and in the agreement achieved lies the only justification needed for either (Goodman, 1955: p. 64).

For Goodman, an inference becomes justified by being brought into agreement with considered principles of good reasoning which have been adjusted and modified so that they fit with other principles and inferences. If one is inclined to make an inference which clashes with a rule of inference that one generally finds acceptable, or to accept a rule of inference that clashes with particular inferences that one is willing to make, then one has to make some adjustment to either the inference or the rule. Once mutual

adjustments have been made to the rules of inference and the inferences, until the two are in agreement, both are justified. No independent check is needed on whether any particular inferences or rules of inference are justified.

A number of other thinkers (Rawls, 1972; Daniels, 1979), particularly those working within moral and political theory, have expanded upon Goodman's theory. Norman Daniels, for example, proposes a way that reflective equilibrium can be used in moral theory. He advances Goodman's notion of reflective equilibrium by arguing that moral judgements should be based upon wide reflective equilibrium rather than Goodman's narrow reflective equilibrium. Narrow reflective equilibrium is achieved when there is agreement between the rules of inference that a person finds acceptable and the judgements that the person is willing to make, but wide reflective equilibrium requires that there is agreement between the rules of inference a person finds acceptable, the judgements that a person is willing to make *and any relevant theories*:

> We do not simply settle for the best fit of principles with judgments, however, which would give us only a narrow equilibrium. Instead, we advance philosophical arguments intended to bring out the relative strengths and weaknesses of the alternative sets of principles (or competing moral conceptions) (Daniels, 1979: p. 258).

The advantage of wide reflective equilibrium is that the judgements and inferential rules that we approve are not simply adjusted to fit with one another; they are scrutinised under the glare of established philosophical theories. Arguments are provided and theories developed to increase the understanding of the strengths and weaknesses of the judgements and inferential rules. Once the strengths and weaknesses of the judgements and rules have been illuminated, adjustments will be made. In wide reflective equilibrium, inferences and rules will only be justified if they can withstand the independent scrutiny of philosophical theories.

It is a version of wide reflective equilibrium that Stein defends as the appropriate method for developing a theory of rationality. In Stein's view, philosophers developing a theory about what it is to be rational should take as their starting point (i) the first order judgements that they are willing to make about cases in which people go about reasoning and (ii) the principles that they think should be followed by people who are reasoning well. They should make adjustments to both their judgements about cases and the principles that they find attractive until the two are brought into agreement, i.e. reflective equilibrium. However, Stein is not satisfied with this narrow reflective equilibrium; he thinks that a theory of rationality is only justified if the judgements and principles that philosophers find acceptable are brought into wide reflective equilibrium. To make the reflective equilibrium wide rather than narrow, Stein argues that philosophers should input both philosophical theories about rationality and—most importantly for current purposes—scientific theories about the limits on human cognition into the process. They should make adjustments to the theory of what it is to be rational produced by judgements about cases and the principles of good reasoning so that they are also in agreement with philosophical theories about what it is to be rational and scientific theories about the nature of human cognition.

> Here, I will argue that including scientific evidence as input to the wide reflective equilibrium process of justifying the normative principles of reasoning is useful… I call the picture of reasoning that results the naturalized picture of rationality. The naturalized picture of rationality says that there are normative principles of reasoning, and that these principles come from a process of wide reflective equilibrium that balances our first-order judgements about what counts as good reasoning, our more general intuitions about what the normative principles of reasoning are, and various philosophical and scientific theories (Stein, 1996: p. 255).

Stein labels the output of the process of wide reflective equilibrium that he suggests should be used by theorists who are interested in rationality as 'the naturalized picture of rationality'.

There are various features of Stein's account that should be attractive to the methodological naturalist. By taking scientific findings into consideration, those theorists who apply the 'naturalized picture of rationality' will ensure that their theories reflect the psychological capacities of those individuals to whom their standards will be applied. Theorists applying the naturalized picture of rationality will therefore avoid the criticism that is levelled against some theories of rationality that they are beyond the reach of humans with ordinary psychological capacities.

The naturalistic appeal of Stein's approach can be understood by seeing his model as a reaction to and an attempt to improve upon what he describes as the 'standard picture of rationality',

> According to this picture, to be rational is to reason in accordance with principles of reasoning that are based on rules of logic, probability theory, and so forth. If the standard picture of rationality is right, principles of reasoning that are based on such rules are normative principles of reasoning, namely they are principles we ought to reason in accordance with (Stein, 1996: p. 4).

The standard picture of rationality presents a lofty ideal about how people should go about revising their beliefs. According to the standard picture of reasoning, it is a requirement for rationality that the consistency preservation principle should be followed. According to this principle, a rational person will check whether any new belief they add to their set of beliefs is consistent with all their other beliefs (Stein, 1996:p. 235). If a person is rational she will follow the rules of logic and probability theory to ensure that the beliefs that she allows into her set of beliefs are consistent. As

has been established by Christopher Cherniak, though, people will never be able to meet these lofty standards because of the human finitary condition (Cherniak, 1986); that is, humans are never able to check the consistency of any target belief (a belief that is a candidate for entry into their belief system) with all the other relevant beliefs that they are committed to because of time limits and limits on their computational capacities. Stein recognises that theorists who do not take the human finitary condition into consideration, and maintain that the standard picture of rationality is correct, will fail to present standards for rationality that ever apply to anyone. They will not be able to distinguish between rational and irrational people, because all people will be irrational simply as a result of being in the human finitary condition. Stein suggests that methodological naturalists who want to avoid these problems can do so by adopting the naturalized picture of rationality that adapts the lofty standards of rationality so that they reflect how people really can go about revising their beliefs.

It seems, then, that given the methodological naturalists' aim to provide standards to use to evaluate people with ordinary reasoning capacities, the naturalized picture of rationality, with its concern for producing a theory that reflects how people really can go about reasoning, ought prima facie to be attractive to them. Some might even see it as a blueprint for methodological naturalism in general. However, it is important to note that the naturalized picture of rationality is far from free from the influence of those problematic aspects of traditional methods of epistemology. Recall that the method of reflective equilibrium used in the naturalized picture requires that philosophical theories and theories from the sciences are brought into agreement with intuitive principles of good reasoning and intuitions about cases of reasoning. The output of the naturalized picture will therefore be highly influenced by intuitions. It therefore fits into the

category that I am labelling in this chapter as 'integrative approaches'. Integrative approaches draw heavily upon the traditional methods of epistemology but they also take scientific findings on board. Advocates of integrative approaches suggest that it is possible to naturalize epistemology by simply making adjustments to the theories that would be produced using traditional methods until the theories reflect the limits of human psychology. They suggest that the problems with the use of traditional methods like intuitions can be overcome by integrating them into a naturalistic framework. As we shall see in a moment, this suggestion is mistaken and that integrative approaches are vulnerable to the criticisms levelled at other intuition-driven approaches, but before we evaluate Stein's position let us consider another similar position that has been presented and defended by Alvin Goldman.

2.2. Goldman's conceptual analysis

A second integrative approach is Goldman's attempt to integrate conceptual analysis into methodological naturalism. Goldman defends the use of conceptual analysis in epistemology. He claims that epistemology should aim to be continuous with folk concepts and that conceptual analysis allows epistemologists to ensure that they present theories which are continuous with these concepts. But rather than advocating straightforward conceptual analysis, Goldman suggests that epistemologists should 'transcend' folk concepts by developing a 'scientific epistemology' (Goldman, 1993). They should first engage in conceptual analysis to get an understanding of what folk concepts of knowledge, rationality and justified belief are like and then they should attempt to improve upon these concepts. Rather than merely describing the concepts that people have and the assessments that they are likely to make, something that can be done at the stage of conceptual analysis, Goldman suggests that epistemologists should

also engage in a normative project of describing what people's concepts ought to be like and how people should make assessments.

> Scientific epistemology, we have seen, has two branches: descriptive and normative. While descriptive scientific epistemology aims to describe our ordinary epistemic assessments, normative scientific epistemology continues the practice of making epistemic judgments, or formulating systematic principles for such judgments. It is prepared to depart from our ordinary epistemic judgments, however, if and when that proves advisable (Goldman, 1993: p. 273).

Goldman suggests that epistemologists should turn to cognitive science to engage in this normative stage of epistemology in a way that improves upon folk concepts, thereby getting an understanding of the psychological mechanisms underlying belief revision.

> On my view, epistemic concepts like knowledge and justification crucially invoke psychological faculties or processes. Our folk understanding, however, has a limited and tenuous grasp of the processes available to the cognitive agent. Thus, one important respect in which epistemic folkways should be transcended is by incorporating a more detailed and empirically based depiction of psychological mechanisms (Goldman, 1993: p. 273).

For Goldman, then, our folk epistemological concepts and notions are useful but they are not adequate to provide an account of how people should get knowledge, rationality or justified belief because they do not accurately reflect the psychological faculties or processes that an agent uses when forming and revising her beliefs.

Goldman therefore suggests that epistemologists should aim to understand folk concepts and notions and then also aim to transcend them to reflect what cognitive science says about the psychological processes underlying belief revision. The parallel between Goldman's position and Stein's is remarkable. Both are methodological naturalists who maintain that epistemologists should use some of the methods of traditional epistemology. Moreover both think that adjustments should be made to folk concepts or intuitively appealing notions to reflect the psychological processes that are available to

people when they form and revise their beliefs. This means that Goldman's and Stein's account share the same attractions for the methodological naturalist. However, it also means that they have the same drawbacks. For the rest of the chapter we shall see how integrative approaches are problematic because they fail to overcome the troubles with traditional methods of epistemology outlined in section 1.

3. Can an integrative approach avoid the troubles of traditional methods of analytic epistemology?

So far we have learnt that there are various troubles with the traditional methods used in analytic epistemology, but some advocates of methodological naturalism nevertheless attempt to integrate these methods into the methodologies they defend for use in epistemology. Stein and Goldman both maintain that traditional methods should be used in epistemology, as long as they are integrated with some naturalistic component; that is, as long as results from the scientific study of psychology influence the epistemological theories produced. The suggestion of these methodological naturalists is that epistemology can be insulated from criticism as long as some naturalistic component is added to its methodology. In this section we shall consider whether the attempt to naturalize traditional methods succeeds. We shall consider three ways that Goldman and Stein attempt to defend their use of traditional methods: (a) Goldman's 'Continuity Claim' that we need to engage in conceptual analysis to ensure that our epistemological theories are really epistemology; (b) Goldman's suggestion that intuitions are useful because they indicate the nature of our concepts; and, finally, (c) Stein's claim that the method of reflective equilibrium will 'weed out' untrustworthy intuitions. I shall show that none of these arguments are successful. If we want to defend a successful version of methodological naturalism we must recognise and

attempt to overcome the problems that come with the use of traditional methods in epistemology rather than simply integrating the methods with some scientific results.

3.1. Goldman's Continuity Claim

Goldman suggests that methodological naturalists should use conceptual analysis, even though there might appear to be problems with it, because epistemologists can only ensure that their theories are truly epistemological theories if they study folk concepts.

Continuity Claim:

> Whatever else epistemology might proceed to do, it should at least have its roots in the concepts and practices of the folk. If these roots are utterly rejected and abandoned, by what rights would the new discipline call itself 'epistemology' at all? (Goldman, 1993: p. 272)

Goldman's suggestion is that epistemology has to be continuous with folk concepts of knowledge, rationality and justified belief because otherwise it cannot really claim to be doing the work of epistemology at all.

Goldman's continuity claim can be seen as a response to positions such as those defended by Stephen Stich (1990). Stich defends a view according to which truth should not be a concern of epistemology. In Stich's view epistemologists should not aim for an account of truth or even knowledge, they should aim to provide an account of how people can achieve their cognitive goals, which can include health, wealth and happiness. He argues that there is no clear account of what truth is that makes it something that people should desire. When Goldman argues that epistemology should aim to be continuous with folk concepts he rejects the idea that epistemology could be what Stich suggests it is. In order for epistemology to be epistemology in Goldman's

view it has to be continuous with what the folk think about knowledge, rationality and justified belief.

The continuity claim is attractive. It seems right that epistemologists should aim to provide an account of how we can achieve distinctively epistemic goals, such as knowledge, rationality and justified belief. They should aim to provide an account of how we can achieve these goals because they are of both intrinsic and instrumental value to us.

Epistemic classifications such as knowledge, rationality and justified belief are of value to humans. We are interested in the truth, we want to know things and fairly and accurately ascribe knowledge; we want to ensure that we are rational and to fairly and accurately evaluate other people's rationality; and we want justified beliefs and to fairly and accurately evaluate the justification of other people's beliefs. We tend to want to know how we can get to the truth, what it is to know and be rational and how we can ensure that our beliefs are justified. We want to understand how to make epistemic classifications because we want concepts such as knowledge, rationality and justified belief to guide our acquisition of belief and our assessments of others. We want an account of knowledge to be able to guide scientists as they seek an understanding of diseases so that they can develop medicines to cure them. We want a notion of how evidence can provide justification for our beliefs as we aim to get facts about war crimes committed when our governments intervened and when they failed to act. There are various good reasons why we value having distinctively epistemic concepts. It is as a direct result of this intrinsic value that we place on epistemic concepts that epistemologists should attempt to provide an account that is continuous with folk concepts. They should attempt to develop theories that are continuous with folk

concepts to the extent that they are able to perform the functions performed by the concepts of knowledge, rationality and justified belief which have just been described, and many other similar functions. Epistemologists who aren't at all constrained by folk concepts might fail to present accounts of knowledge, rationality and justified belief that could enable us to achieve important epistemic goals because they might not even see the value of various epistemic notions.

Even if, however, it weren't the case that humans did intrinsically value epistemological notions such as knowledge, rationality and justified belief, the continuity claim would still have appeal because it seems that folk epistemic concepts are important because truth is instrumentally valuable. Stich rejects the claim that truth is instrumentally valuable. To support this point he describes a situation in which a person is taking a flight. Harry checks the time of his flight and forms the true belief that his flight leaves at 7.45. He catches the flight but then it crashes. Stich suggests that truth is not of instrumental value because Harry would have been better off if he had formed the untrue belief that the flight was at 8.45. However, as Susan Haack suggests, although Harry might have been better off if he had formed the false belief that his flight left at 8.45, he would have been even better off if he had the more complete belief that the flight was due to leave at 7.45 but would crash (Haack, 1993: p. 200). If he had that complete information he could have saved all the other people on the flight. Haack's response to Stich's example therefore suggests that although sometimes having a true belief can be detrimental, it is generally more beneficial to have a more complete set of true beliefs. Sometimes when it appears that we benefit from having false beliefs, we would benefit even more if we had additional true beliefs. This suggests that truth really is instrumentally valuable. Even if it is that what we truly want is health, wealth and

happiness, we are more likely to get what we want if we have true beliefs about the world.

Prima facie, then, Goldman's continuity claim has some appeal. It seems that we should value truth and other related epistemic concepts such as knowledge, rationality and justified belief, and epistemologists can ensure that their theories reflect the importance of these concepts by studying folk concepts. But Goldman's continuity claim, while appealing in itself, is not successful as a defence of conceptual analysis. This is because it is consistent to maintain that there should be some degree of continuity between folk concepts and epistemological theories *and* that conceptual analysis will lead epistemologists astray.

We can understand this point by considering a distinction made by Kornblith between strong and weak continuity between folk concepts and epistemological theories.

> The idea that a scientific epistemology must, in some sense, be continuous with our folk epistemology if it is to be worthy of being called 'epistemology' at all is certainly a plausible one, but there are weak and strong versions of this idea, depending on just how much continuity is required. Consider, for example, the relationship between alchemy and chemistry. Chemistry did not, of course, come about in an instant, and the transition from alchemy to chemistry, on any reasonable account, shows a continuous development. Some alchemical concepts were retained or modified as chemistry emerged; others were entirely abandoned. If the relationship between our folk epistemology and scientific epistemology is thought of as comparable to the relationship between alchemy and chemistry, then a similar sort of continuity should be expected. And I certainly have no objection to this at all.
>
> But Goldman seems to have something far stronger in mind. On Goldman's view, contemporary epistemologists need to begin their work by engaging in a conceptual analysis of our folk epistemological concepts. And if the relationship between folk epistemology and scientific epistemology is modelled on the relationship between alchemy and chemistry, it is hard to see why we should demand any such thing. Chemists do not need to begin their investigations by providing detailed analyses of folk chemical concepts. A scientific chemistry certainly aims, among other things, to

> transcend the confines of our folk chemical concepts, but this does not mean that effort needs to be spent on elucidating those folk categories before the real work of chemical theorizing can be done (Kornblith, 2007: p. 38-9).

In this extended passage, Kornblith describes two notions of continuity that Goldman could be adopting when he claims that epistemology should be continuous with folk concepts. The first notion of continuity is a weak notion, according to which the epistemological theory is to folk concepts like chemistry was to alchemy; it took roughly the same subject matter as its target but did not involve the analysis of the concepts of alchemy and had different outcomes to alchemy as a result. The second notion of continuity is a strong notion, according to which the epistemologists should study folk concepts. It involves the analysis of concepts and the search for the necessary and sufficient conditions for the application of the concepts.

What I would suggest, along with Kornblith, is that what the continuity claim supports is the weak notion of continuity. The arguments given above in defence of the intrinsic and instrumental value of epistemological concepts do not show that epistemologists should engage in conceptual analysis, they only show that epistemologists should aim to provide an account of how people can attain the truth, get knowledge and justified belief and be rational. Sometimes Goldman seems to only defend the weak version of the continuity claim:

> Examining folk epistemic concepts should reveal how truth (true belief) is a primary basis of epistemic evaluation and epistemic achievement. This is indicated, for example, by the truth-condition on knowledge and the reliability desideratum associated with justifiedness. When moving from folk epistemology to scientific epistemology, we should retain the concern with truth-related properties of methods and practices. If we never studied folk epistemic concepts, or studied them without proper understanding, this desideratum might elude us (Goldman, 2007: p. 22).

In this passage, Goldman suggests that we should study concepts so that we can ensure that our epistemological theory reflects the truth-related properties of epistemological concepts. It seems to me that this aim could be achieved without in depth consideration of the application conditions of a concept, just by considering what general functions our epistemological concepts tend to play, i.e. our epistemological concepts usually provide us with some guidance about how to get to the truth. This means that Goldman, in this passage at least, only provides reason to aim for weak continuity. But at other times he suggests that conceptual analysis is required to ensure the requisite continuity with folk concepts. The continuity claim doesn't support this stronger claim because it seems that the important continuity with our folk concepts can be achieved without conceptual analysis.

If we recall the arguments presented in section 1.1 about the problems associated with conceptual analysis we find another reason why only weak continuity should be the aim of epistemologists. Humans are prone to misunderstandings of how people can best revise their beliefs and these misunderstandings are reflected in people's epistemological concepts. If the strong continuity claim were upheld then theories of knowledge, rationality and justified belief would have to be continuous with these concepts which are laden with misunderstandings. But this is incompatible with the suggestion that psychological results about the processes underlying cognition should be taken into consideration when developing epistemological theories. This suggestion implies that people should discard concepts based on mistaken beliefs about how human cognitive systems work. So it seems that if we take Goldman as saying that the strong notion of continuity is correct then we must construe him as defending a contradictory or least confused position: that we should develop theories that both reflect the

misunderstandings underlying our epistemological concepts and correct these misunderstandings when science tells us that we are wrong about belief revision.

To be sympathetic to Goldman, then, it seems best to construe him as defending the weak continuity claim according to which epistemologists should aim to develop theories that reflect the general purposes of folk epistemic concepts and aim to avoid the misunderstandings underlying the concepts. But if Goldman's continuity claim is given this sympathetic reading then it provides no good reason for engaging in conceptual analysis. Goldman's continuity claim therefore fails as a defence of conceptual analysis.

3.2. Intuitions as an evidence source

Goldman's attempt to defend conceptual analysis seems even less promising when one considers how it relates to claims about the use of intuitions in epistemology. Recall that in section 1 we found that that conceptual analysis is a poor method to use in epistemology not only because concepts should not be the target of epistemology but also because untrustworthy intuitions are often used in conceptual analysis. In his work with Joel Pust (Goldman and Pust, 1998/2002), Goldman suggests that he can avoid the criticisms of intuition-driven epistemology because he only uses intuitions as evidence for folk concepts which are psychological entities. In the previous section I argued that folk concepts should not be the target for epistemology. In the present section I will argue that given the unstable nature of philosophical intuitions, either they do not provide access to our concepts, or we have additional reason for thinking that our concepts should not be the target of epistemology.

In Goldman and Pust's view, intuitions are useful because they track folk concepts. They accept that there are good reasons for thinking that intuitions would fail to track

mind-independent items, but maintain that they are useful to the extent that epistemologists are interested in what is inside the head, i.e. are engaged in what they call 'mentalism':

> Mentalism interprets philosophical analysis as trying to shed light on the *concepts* behind philosophically interesting predicates, whether the term "concept" refers to a psychological structure or state that underpins a cognizer's deployment of a natural-language predicate. [...] our form of mentalism is very congenial to the notion that intuition might be a basic evidential source. Under mentalism, moreover, its status as a basic evidential source can be sustained in a fully naturalistic framework (Goldman & Pust, 2002: p. 83).

Goldman and Pust suggest that a mentalist interpretation of epistemology, according to which mental entities are the object of study, is congenial to the idea that intuitions can be a basic source of evidence for an epistemological theory. They claim that it is likely that there will be a direct causal relationship between possessing a particular concept and having a related intuition because it will be a part of having a concept that people will have dispositions to have certain intuitions:

> The concept associated with a predicate 'F' will have many dispositions, but among them are dispositions to give rise to intuitive classification judgments such as "example e is (is not) an instance of F". Thus it is not only possible, but almost a matter of definition, that if the concept possessor were fully informed about the features of e, then if e satisfied the concept he expresses through 'F', his intuitive response to the question of whether e satisfies this concept would be affirmative; and if e did not satisfy the concept he expresses through 'F', then his intuitive response to the question of whether e satisfies this concept would be negative (ibid. p. 84).

Goldman and Pust's suggestion is that if one possesses a concept, one will necessarily have a disposition to have certain intuitions, so these intuitions can be used as a basic source of evidence for the concept. The intuitions provide a reliable indication that a person has a particular concept. Goldman and Pust therefore present the following response to the challenge to the use of intuitions in epistemology:

1. Intuitions are unreliable sources of evidence about extra-mental items, but reliable sources of evidence about mental items,
2. If epistemologists are concerned with extra-mental items then intuitions are a poor source of evidence, but if they are concerned with mental items then intuitions are a good source of evidence,
3. Epistemologists are (or at least should be) concerned with mental items,

Conc. Intuitions are a good source of evidence for epistemologists.

In section 3.1 I disputed premise 3. In this section I will argue that the empirical evidence about intuitions suggests that either 1 and 2 are incorrect or we have additional reason for rejecting 3.

The claim that intuitions are reliable indicators of our concepts is problematic because of the criticisms levelled against intuitions in section 1.2. The empirical results presented by experimental philosophers and psychologists suggest that epistemological intuitions can be influenced by situational factors unconnected with the philosophical issues being considered. Recall that, for example, evidence from experimental philosophy suggests that the intuitions that people have about whether a belief is knowledge can depend upon what a person thinks about before considering the thought experiment (i.e. whether it is a clear case of knowledge or non-knowledge). Moreover, psychological results suggest that a person's intuition about how wrongful an act is can be influenced by whether that person is in a clean room and has been able to wash his or her hands. There are two potential responses that can be made about the relevance of these empirical findings to claims about the relationship between philosophical intuitions and concepts. The first response is to maintain that if epistemological

intuitions can be influenced by situational factors, concepts are also influenced by those factors. The second response is to deny that concepts are influenced by situational factors but then one also has to deny that intuitions reliably indicate our concepts. As we shall see, each of these responses causes problems for those who defend the use of intuitions as a source of information about concepts.

As Goldman and Pust insist that intuitions provide a reliable source of evidence for folk concepts, it seems that they are committed to the idea that folk concepts are influenced by situational factors irrelevant to epistemology. They are committed to the following:

1. Intuitions are influenced by situational factors irrelevant to epistemology (from empirical evidence).
2. Intuitions track folk epistemological concepts (from G&P's defence of the use of intuitions in epistemology).

Conc. Folk epistemological concepts must be influenced by situational factors irrelevant to epistemology.

We can see in the above argument that it is a direct result of Goldman and Pust's defence of the use of intuitions as evidence in epistemology coupled with the empirical evidence about intuitions that they are committed to the idea that concepts are influenced by situational factors. Their argument has the important consequence that people's epistemological (and moral) concepts can change as quickly as they can think about some other thought experiments, change rooms, or wash their hands.

Goldman and Pust's attempt to provide some account of how intuitions can be reliable evidence in epistemology therefore has some rather bizarre consequences. It implies that

our epistemic concepts can change to reflect the conditions under which we think about a thought experiment. It implies that our everyday applications of epistemological concepts are influenced by situational factors to the extent that it might be correct to say 'if I had spoken to more knowledgeable people this morning my concept of knowledge would be different to what it is now' because our current concept of knowledge would be dependent on past situational factors such as our experiences of speaking to others. Our concept of rationality could be one thing when we are in one set of surroundings and another thing when we are in another. If I am in a clean room I might have a different concept of moral wrongdoing than if I am in a dirty one. These things sound implausible, but stranger things have been found to be true by empirical investigation, so we should not dismiss the possibility from the armchair. However, even from the armchair it is still possible to show that this response won't do as a defence of conceptual analysis.

It is not possible to defend the use of intuitions in epistemology on the basis that they reveal epistemological concepts that are influenced by situational factors because if concepts are influenced in this way they seem to provide the wrong kind of target for epistemology. Epistemologists should not base their theories upon epistemological concepts if those concepts are influenced by factors irrelevant to the questions they are interested in. Otherwise an epistemological theory might have as a part of its basis the fact that its advocate was in a particular room when developing the theory. To avoid irrelevant factors from deciding the outcome of the philosophical theorising, epistemologists should avoid targeting concepts if concepts are influenced by such factors.

If intuitions were only influenced by factors relevant to epistemology then things might be different. As mentioned in section 1.2, Goldman and Pust actually defend the idea that epistemological theories should reflect how the categorisation judgements underlying concept application are context-sensitive. They appeal to psychological results suggesting that people's judgements can be affected by the context in which they make those judgements. So, for example, people are more likely to say that they are happy if they have just spent time with a disabled person. They interpret these findings as suggesting that people are influenced by the relevant alternatives when they make categorisation judgements. Then they suggest that epistemological theories that reflect how attributors of knowledge are influenced by context should be viewed positively in light of these findings.

In contrast to this argument, the suggestion that all of the variation in epistemological intuitions demonstrated by experimental philosophy could reflect acceptable variation in epistemological concepts is highly implausible. Whereas a reasonably good case can be made that some aspects of the context of the attribution of knowledge can be relevant to epistemology, a good case cannot be made that the selection of an epistemological theory should reflect the room in which an epistemologist develops a theory (or a moral theory should reflect whether a moral philosopher has washed her hands). It seems, then, that if changeable epistemological intuitions reflect changeable epistemological concepts then we should not target those concepts because they are so easily influenced by situational factors that should not influence epistemological theorising.

Another plausible interpretation of the experimental philosophy data is that intuitions do not provide a reliable indicator of concepts. They provide a faulty and unreliable picture of what our concepts are. Our immediate responses to thought experiments can provide

useful evidence about whether we are inclined to apply a concept to a particular example at a particular point in time but they can't necessarily tell us what that concept is. According to this view, our concepts remain stable and it is just our intuitions that change, for some reason other than the nature of our concepts. This interpretation of the experimental data avoids the criticism that concepts are bad targets for epistemology because they are unstable, but it also undermines the claim that intuitions are good evidence because they provide direct evidence for our concepts. On this view, intuition-driven methodology is problematic not only because it takes our concepts as a target but also because the method used to understand the nature of our concepts is unreliable—intuitions simply aren't good indicators of our concepts.

Whichever of these responses to the experimental findings is taken, they still suggest that intuitions provide a poor source of evidence for philosophical theories. Either philosophical intuitions reflect concepts but concepts are a poor target for epistemology because they are unstable and influenced by situational factors irrelevant to epistemology, or intuitions do not reflect our concepts and then they provide poor evidence for even these mental items. Either way we have additional reason for rejecting the idea that it is our concept of knowledge that should be the target of epistemology because either our concepts are unstable or we are unable to get reliable access to them.

3.3. 'Weeding out' troublesome intuitions

Those who accept my suggestion that Goldman fails to successfully defend the use of intuitions as evidence for our epistemological concepts might find a claim found in the work of Stein appealing; that is, the claim that his methodology 'weeds out' the influence of troublesome intuitions. Unlike Goldman, Stein acknowledges that the

dependency of his picture on intuitions makes him vulnerable to criticisms like those outlined in section 1.2 about the untrustworthy nature of intuitions as a source of philosophical theories. But he is optimistic that the 'naturalized picture of rationality' will weed out untrustworthy intuitions.

> The naturalized picture of rationality does give weight to our intuitions about rationality, but our intuitions are brought into balance with philosophical theories and scientific evidence. Our intuitions are continually being evaluated, re-evaluated, and modified in light of these considerations. Given this, whatever epistemic chauvinism is initially present in our intuitions about rationality is likely to be cancelled out by considerations that are not chauvinistic, namely philosophical and scientific considerations (Stein, 1996: p. 256-7).

> It is possible that these considerations are chauvinistic too- science and philosophy might be biased because they give undue weight to our 'local' intuitions. This is a possibility that we should take seriously, but full consideration of it is beyond the scope of this project. For now, suffice it to say that since critical reflection, evidence, and argumentation are central parts of both philosophy and science, epistemic chauvinism is likely to be discovered and weeded out through the wide reflective equilibrium process involved in the naturalized picture of rationality (Stein, 1996: p. 257).

In these passages, Stein recognises the weight given to intuitions in his picture, but remains optimistic that the method of reflective equilibrium used to create the naturalized picture of rationality will ensure that our intuitions aren't decisive to the final outcome of the procedure. In his view, the intuitions that are influential in the method of reflective equilibrium will be evaluated and re-evaluated so many times that any epistemic chauvinism will be removed. Stein maintains that the input of philosophical and scientific theories into the process of reflective equilibrium will provide some sort of guarantee that troublesome intuitions won't be decisive in the naturalized picture of rationality.

If Stein's argument here worked it could be used as a defence of approaches that integrate traditional methods with some naturalistic component. The intuitions which

are influential in epistemology could be deemed troublesome when they are used in isolation as evidence for a philosophical position but not when philosophical theories and scientific theories are also used to develop an epistemological theory. But for the argument to work it would have to be the case that either philosophical or scientific theories are free from the influence of intuitions and able to perform the functions of intuitions. It would have to be possible for intuitions to be overcome by these theories. However, I shall suggest that the position fails because philosophical theories are far from free of the influence of intuitions, and scientific theories, which are relatively free from the influence of intuitions, do not perform the same functions as intuitions so cannot override them.

We can understand the extent to which philosophy is influenced by intuitions by considering once again the extent to which conclusions in philosophy are often based upon responses to thought experiments. Traditional analytic philosophy has historically been based upon thought experiment methodology. Epistemologists have proposed analyses of knowledge, justification and other epistemological concepts by using thought experiments. Other thought experiments have been used to challenge those same analyses. Philosophers have viewed it as a sign of their success if they have elicited intuitive responses that support their positions. Recent papers that challenge the use of intuition in philosophy, such as those presented in the compilation *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry,* have been met with a race by traditional analytic epistemologists to defend intuition-driven methodology because it is accepted as so central to philosophical thinking (e.g. Williamson, 2007; Sosa, 2009). There is therefore good reason for thinking that much of philosophy, including many of the epistemological theories likely to be influential in

Stein's naturalized picture of rationality, are accepted because they are in agreement with intuitions. Intuitions are, at least at the present time, highly influential in philosophy.

It is because intuitions are so influential in philosophy that it seems inevitable that they will be decisive in the naturalized picture of rationality. Even if philosophical theories challenge the intuitions that people initially input into the process of reflective equilibrium, intuitions will ultimately be decisive. Either the intuitions underlying the support for the philosophical theories will be influential because they will overcome the initial intuition or the person applying the method of reflective equilibrium will stubbornly retain the original intuition. And in some respects a person applying the method could not be blamed for retaining their original intuition when it is only other intuitions that support the position being presented in opposition. Whether philosophical theory overcomes our initial intuitions will depend upon which intuitions are more compelling; those that support the initially intuitive picture or those that support the philosophical theory. This means that the process of evaluating and re-evaluating theories of knowledge, rationality and justified belief through philosophical evaluation and dialogue cannot be expected to produce a theory free from the influence of intuitions. Either philosophical theories will overcome other intuitions, in which case the intuitions supporting the theories dominate, or the intuitions will withstand scrutiny under the glare of philosophical theories. Stein is mistaken when he suggests that philosophy will have the power to overcome untrustworthy intuitions.

One might think that science can succeed where philosophy fails because science is not dependent on intuition in the same way that philosophy is. Scientific theories get support from empirical data so they do not wholly rely on support from intuition.

Perhaps, therefore, the introduction of scientific theories into the integrative approaches can overcome the influence of untrustworthy intuition. Although prima facie appealing, this move ultimately fails. This is because scientific theories underdetermine the evaluative or normative aspect of epistemology and untrustworthy intuitions can influence those parts of epistemology which science underdetermines.

We can understand this point by considering the criticisms levelled against Quine's position in his earlier work on naturalized epistemology first discussed in chapter two. Quine (1969) called epistemology a descriptive project and described it as a branch of psychology. He claimed that epistemologists should be like scientists. Epistemologists have responded to this argument by saying that epistemology cannot be a purely descriptive enterprise because what is required of epistemology is a set of normative principles. Normative principles should provide guidance about how people should behave rather than how they do behave. For these principles of right and wrong, epistemologists need to make evaluations about how people behave. These evaluations may be informed by science, as methodological naturalists suggest, but they will not be fully determined by science. On top of science we need some standards against which to assess the ways that we find that people behave. It is in this normative aspect of epistemology that intuitions can be really influential. Our intuitions guide our normative evaluations, which is where science falls silent. Science will not therefore be able to override intuitions because it plays a different role to them.

It might be helpful to consider a case study. Let us return to the rationality debates and use the discussions between psychologists about whether humans are rational as our case study. In recent times two groups have emerged in the rationality debates, the heuristics and biases group and a cluster of evolutionary psychologists. In the first

group are psychologists who discovered that humans tend to use heuristics and biases in their reasoning, that is, they use shorthand rules of thumb rather than checking the consistency of their various beliefs with each other. Among the evolutionary psychologists is Gerd Gigerenzer. He and his colleagues have taken the evidence presented by the heuristics and biases group suggesting that people use heuristic shorthand methods and suggested that they could be rational to do so[26]. They suggested that it could be rational to use such methods because they do not require much information or cognitive resources to succeed (for more on the debate see chapters two and five).

What matters about this dispute for current purposes is that the proponents of the two opposing positions base their arguments on very similar scientific data about how people go about reasoning. They each accept that people use shorthand methods of reasoning and can sometimes make errors doing so. But Gigerenzer and his colleagues maintain that people can nonetheless be rational when using these methods because of the evolutionary psychologists' disposition to associate rationality with the successful exploitation of environmental resources. Those who are predisposed to maintain that it is valuable to attempt to preserve the consistency of various beliefs would have produced a very different notion of rationality to the one presented by evolutionary psychologists. If they had a strong enough intuition that their notion of rationality is the correct one, no amount of scientific evidence might change their minds. The scientific findings alone do not determine whether someone accepts that the use of shorthand methods is rational. Some combination of an understanding of the scientific findings

[26] For examples of the work of Gigerenzer and his colleagues see Gigerenzer and Todd (1999) and Gigerenzer and Selten (2000).

and a predisposition to certain intuitions about rationality will decide which theory of rationality one adopts.

The rationality example shows that Stein is mistaken if he thinks that the introduction of scientific findings will eradicate the influence of intuitions because the scientific results do not determine which normative theory is accepted. A person with a strongly held intuition that consistency is required for rationality would remain committed to the heuristics and biases position while a person who has the intuition that rationality should be associated with successful exploitation of environmental resources might be convinced by the evolutionary psychology position.

In seems, then, that advocates of integrative approaches cannot assume that they overcome the influence of intuitions just by defending the idea that results from the cognitive sciences are relevant to epistemology. The introduction of results from psychology can present a challenge to certain concepts and intuitions, because it can suggest that certain concepts and notions of knowledge, rationality and justified belief are beyond the reach of ordinary thinkers, but it cannot ensure that the influence of all troublesome intuitions is overcome. There is also no guarantee that philosophical thinking, even when under the influence of scientific theorising, will overcome the influence of untrustworthy intuitions. In fact, it seems that if intuitions are weeded out in the process of making adjustments to our concepts, philosophical theories and scientific theories as a part of the integrative approach, they are likely to be supplanted by philosophical theories influenced by other intuitions and it will be these that decide which epistemological theory is accepted. Thus integrative approaches are not able to ensure that untrustworthy intuitions are not influential.

We can conclude, then, that intuitions are inevitably going to influence the output of Stein's naturalized picture of rationality. One might think, however, that Stein's position is still defensible because a theory based on the right kinds of intuition can still be valuable. But things are not so simple because there is no mechanism available to distinguish reliable from unreliable intuitions. Recall that intuitions cannot be calibrated; it is not possible to check intuitions against some independent measure. This means that there are not resources available to distinguish good from bad intuitions. We cannot ensure that our theories are only influenced by desirable intuitions.

There is a further problem with Stein's account which becomes clear once one realises that unwieldy intuitions cannot be weeded out by the process of reflective equilibrium. Recall that in section 3.1 we discussed the distinction between weak and strong continuity as an aim for epistemologists; epistemologists can either aim to ensure that there is some general continuity with their folk concepts, for example, that they ensure that their concepts are able to perform the same functions as ordinary folk epistemic concepts, or they can aim at the sort of strong continuity with concepts achieved through conceptual analysis. Stein seems to want to defend the idea that both weak and strong continuity with folk concepts should be the aim of epistemology. He suggests that strong continuity is required by arguing that intuitions about cases should be input into the process of reflective equilibrium, influencing our epistemological theories. But he also suggests that intuitions should be weeded out, suggesting that he only aims for weak continuity with folk concepts. Now it is not possible to achieve both weak and strong continuity at the same time and this is the source of Stein's failure to present a theory that truly 'weeds out' intuitions—it is not possible to weed out intuitions at the same time as using them as a source of an epistemological theory. Neither weak nor

strong continuity is obviously compatible with Stein's aims and methodology. If he aims for only weak continuity with folk concepts then he has to explain why intuitions should influence his naturalized picture. By arguing that intuitions should be input into the process of reflective equilibrium Stein seems to defend the idea that epistemologists should aim for strong continuity with our concepts. However, if he aims for strong continuity with folk concepts, then he has to explain why intuitions and the concepts they reflect, which can be epistemically chauvinistic and deeply mistaken, should be influential. By suggesting that epistemologists should aim for both weak and strong continuity with ordinary folk epistemic concepts, Stein provides an inconsistent position which ultimately fails in its aim to integrate traditional methods into a satisfactory version of methodological naturalism.

In sum, then, it is not possible to "weed out" the influence of undesirable and untrustworthy intuitions while maintaining that traditional methods of epistemology are valuable. It is an inevitable consequence of the use of intuitions and conceptual analysis that epistemological theories are open to the various criticisms outlined in section 1.

Conclusion

There are various problems with conceptual analysis and intuition-driven epistemology. These problems are particularly evident from a naturalistic perspective. Concepts do not seem to be an appropriate target for a theory which aims to reflect how people can best revise their beliefs if those concepts are mistaken about successful belief revision. But even if we were to accept that concepts should be the target of epistemology it seems that the intuitions we use to reveal those concepts are unstable and influenced by factors irrelevant to philosophy, so provide a poor basis for an understanding of our epistemological concepts.

Some methodological naturalists who recognise the force of the arguments against the use of traditional armchair methods have nevertheless advocated their use as a part of a wider naturalistic project. We have seen that they provided three general solutions to the problems with the methods of traditional analytic epistemology, but we found that each of these solutions failed.

Conceptual analysis cannot be defended on the basis that epistemological theories need to be continuous with folk concepts because some degree of continuity can be attained without engaging in conceptual analysis. Given that there are strong reasons for thinking that folk concepts are mistaken it seems that epistemologists should seek an alternative to engaging in conceptual analysis as a way to ensure their theories are continuous with those concepts.

Intuitions cannot be defended on the basis that they provide evidence for our concepts because recent results from empirical studies suggest that intuitions are unstable. This suggests that either concepts are also unstable, and therefore should not be the target of epistemology, or intuitions do not give access to concepts, so they are not useful tools in conceptual analysis.

Finally, an integrative approach cannot get around the problems associated with traditional analytic epistemology simply by taking existing philosophical and scientific theories into consideration. Neither philosophical theorising nor scientific studies will guarantee that problematic intuitions are overcome.

Overall, then, those who have defended the adoption of an integrative approach in epistemology have failed to supply good reasons for thinking that in doing so we can

acquire a theory that avoids the criticisms of traditional analytic epistemology and its use of intuitions and conceptual analysis.

Chapter Four: Why Knowledge is not a Natural Kind

Introduction

Hilary Kornblith claims that knowledge is a natural kind. He uses this claim to motivate his account of the methods to use in epistemology consistent with methodological naturalism, which is posed as an alternative to the integrative approaches discussed in chapter three. In this chapter I will outline a number of problems for the claim. I argue that epistemologists should consider the goals that people aim to achieve through their epistemic classifications, so the process that leads to something being classified as knowledge or non-knowledge will not be wholly natural, it will be dependent upon human interests[27]. The strategy I will use to establish this claim will be to show how things can go wrong if we treat knowledge as if it is a natural kind in the way that Kornblith suggests we should, using intuitions about obvious cases of knowledge and/or results and explanations from the sciences as a guide while we aim to uncover the natural kind. I shall argue that ultimately we should only accept something into the category of knowledge if this categorisation will answer to human needs, so knowledge is not a natural kind. Epistemologists should not, in my view, expect to discover a natural kind of knowledge.

1. Knowledge as a natural kind

[27] My arguments in this chapter will therefore suggest that epistemologists should take something similar to the approach adopted by Edward Craig (1990) when he suggested that at least a part of epistemology should involve examining the purposes of our epistemic classifications. Kornblith (2011) criticises Craig's position on the basis that epistemologists should not be interested in their concepts and should be interested in the naturally existing category of knowledge. My arguments in this chapter suggest that Kornblith is wrong to criticise Craig in this way because there are various problems with assuming that knowledge is a natural kind which can be avoided by making a move like Craig's and studying the goals we aim to achieve through our epistemic classifications.

The conception of natural kinds that Kornblith adopts when he argues that knowledge is a natural kind is Richard Boyd's[28]. Boyd (1988, 1991) argues that specific items belong to a natural kind as long as the success of scientific predictions and explanations depends upon assuming that they have properties which consistently cluster together in nature. He defines natural kinds as 'homeostatic property clusters'. He assumes that it is because properties consistently cluster together in a law-like manner that it is possible to make successful predictions about the presence of one property based on evidence of the presence of another. For example, it will be possible to identify whether an individual is a member of the species 'human being' by identifying in that individual certain properties shared by many members of the species. And it will be possible to explain why an individual has certain traits by explaining how the species as a whole acquired those traits. As long as it is possible to make successful scientific predictions and explanations like these by assuming that the properties cluster together in this law-like way, we have reason to think that those clustered properties form a natural kind. The success of the predictions and explanations can be explained by how the properties cluster together in a causally governed way independent of human conventions, reflecting the causal laws of nature.

> …the use we make of reference to the kind in induction and explanation requires that it be defined by a set or cluster of properties whose membership is determined by the causal structure of the world and is thus, in a relevant sense (which I propose to specify), independent of our conventions or our theorizing (Boyd, 1991, p. 129).

[28] Given that there are difficulties locating essential properties that are required for membership of some natural kinds, for example the natural kind of species, Boyd's conception of natural kinds is attractive. It does not require that there are some properties essential to all members of a species, unlike other prominent account of natural kinds, such as those presented by Putnam (1975) and Kripke (1972). Boyd's conception of natural kinds merely requires that there are law like clustering of properties in nature—something that is implied by the success of scientific predictions and explanations.

A natural kind is thus, according to Boyd, a cluster of properties which features in causal laws because of the clustering. If causal explanations depend for their success on the idea that certain properties cluster together then it is right to treat those clustering properties as a natural kind.

Kornblith's adoption of this conception of natural kinds leads him to claim that knowledge is a natural kind on the basis that there is a cluster of reliably produced true beliefs which are assumed to exist by scientists while those scientists successfully develop predictions and explanations. He claims that when scientists who study animal cognition (i.e. cognitive ethologists) try to explain how members of a certain species have survived it becomes important to attribute all members of the species with knowledge.

> If we wish to explain how it is that members of a species have survived, we need to appeal to the causal role of the animals' knowledge of their environment in producing behaviour which allows them to succeed in fulfilling their biological needs. Such explanations provide the basis for accurate inductive inference. The knowledge that members of a species embody is the locus of a homeostatic cluster of properties: true beliefs that are reliably produced, that are instrumental in the production of behaviour successful in meeting biological needs and thereby implicated in the Darwinian explanation of the selective retention of traits (Kornblith, 2002: p. 62).

When describing how the members of a species have been able to survive and reproduce, scientists speak about an animal's ability to acquire the information that they require from their environment. So, for example, if asked to explain how the plover species survived[29], they could explain that the plovers' cognitive capacities are able to pick up on information about intruders, the danger they pose to the nest, and how it is possible to distract them. They could explain the behaviour of members of the plover

[29] I am borrowing Kornblith's example of the plover here. He introduces the example of how scientists can predict and explain the behaviour of this species of bird in *Knowledge and the Place of Nature.*

species in terms of how the cognitive capacities of members of the species are attuned to their environment. According to this view, the species survived because its members have cognitive mechanisms which consistently enable them to pick up on the information they require for their survival, such as information about intruders. For Kornblith, these beliefs referred to by scientists as they develop their successful scientific explanations are knowledge because they are reliably produced true beliefs, so successful explanations of the survival of animal species depend upon the attribution of knowledge.

On Boyd's conception of natural kinds, these arguments support the claim that animal knowledge is a natural kind. The success of the explanations of animal behaviour depends upon assuming that the behaviour is underwritten by knowledge, so animal knowledge is a natural kind: it is the natural kind of reliably produced true belief assumed to exist by scientists developing explanations of the survival of animal species.

At this juncture one might ask how epistemologists are supposed to respond to the claim that knowledge is a natural kind, and that animals have knowledge, given that they generally aim to understand *human* knowledge. What are they supposed to conclude about human beliefs from the claim that there is a natural kind of animal knowledge? Kornblith (1998, 2002: chapter 1) accompanies the claim that knowledge is a natural kind with a description of how philosophical methodology in general should reflect this. He claims that epistemologists should be like rock collectors who aim to discover what is common to all the instances of the natural kind of knowledge:

> When we appeal to our intuitions about knowledge, we make salient certain instances of the phenomenon that need to be accounted for, and that these are genuine instances of knowledge is simply obvious, at least if our examples are well chosen. What we are doing, as I see it, is much like the rock collector who gathers samples of some interesting kind of stone for the

> purpose of figuring out what it is that they have in common. We begin, often enough, with obvious cases, even if we do not yet understand what it is that provides the theoretical unity to the kind we wish to examine. Understanding what that theoretical unity is is the object of our study, and it is to be found by careful examination of the phenomenon, that is, something outside of us, not our concept of the phenomenon, something inside of us. In short, I see the investigation of knowledge, and philosophical investigation generally, on the model of investigations of natural kinds (Kornblith, 1998: p. 134).

In Kornblith's view, then, epistemologists should begin with their intuitions about obvious cases of knowledge. In doing so, they will identify obvious instances of the natural kind. Then they should seek an understanding of those features that unify the kind. Then they will be able to identify other, less obvious members of the kind.

Kornblith contrasts this sort of investigation with conceptual analysis, where philosophers seek an understanding of the necessary and sufficient conditions under which it is appropriate to apply their concept of knowledge.

> 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. […] the subject matter of epistemology is knowledge itself, not our concept of knowledge (Kornblith, 2002: p. 1).

According to Kornblith, philosophers should study the real nature of the phenomenon of knowledge rather than anybody's concept of knowledge. By studying the theoretical unity of the natural kind, a kind that can be discovered at the early stages of investigation by considering intuitions about obvious cases, it is possible to understand what knowledge is, without seeking out an understanding of human concepts, or how they reflect human ideas or purposes.

2. Criticising Kornblith: An inconsistent position

I have a number of criticisms of Kornblith's position. Firstly I would like to highlight a tension between his claims about the place for intuitions and his claims about the place for science in epistemology. He claims both that knowledge is a natural kind which can be discovered by considering the classifications made by scientists and that intuitions should be used as a starting point in epistemology. Epistemologists should take the claims made by scientists about knowledge seriously because 'a priori arguments against well established scientific research programs have a history of failure' (Kornblith, 2002: p. 32) and yet, at the same time, intuitions are supposed to provide a reliable guide about obvious cases of knowledge. On the one hand, then, science seems to be doing a significant amount of work for the epistemologist because belief classified as knowledge as part of scientific explanation really is knowledge. On the other hand, epistemologists are supposed to be guided by brute intuitions, uninformed by science; these can provide evidence for or against an account of knowledge. He seems to assume, then, that epistemologists' brute intuitions and results from the sciences will complement each other, so that they can be used together to develop a theory.

The first problem with this approach is that it seems as if it will only be available to some people, who share certain assumptions. In many cases it won't be possible to use both intuitions and scientific explanations as a starting point, not least because very many epistemologists assume that scientists are ill-equipped to identify beliefs which are epistemically valuable. For example, many epistemologists have the view that epistemology should be an armchair enterprise and that we can understand the nature of knowledge by considering thought experiments. They take the view that results from the sciences are irrelevant to the project of developing an epistemological theory. This

means that they will find it unintuitive that science can identify knowledge[30]. And this is an intuition about knowledge, after all. It is the intuition that belief based on armchair reflection and not on reflection on results from science constitutes knowledge in epistemology. But if this intuition is taken as a credible starting point for epistemology there seems to be little hope that many epistemologists will be both guided by their intuitions *and* accept Kornblith's claim that the beliefs to which scientists refer are knowledge. My first criticism of Kornblith's claim that knowledge is a natural kind discovered by intuitions *and* science is, then, that it fails to supply a consistent account of how epistemologists should work if they do not share the assumption that results from science can supply an account of knowledge.

Kornblith might respond to this criticism by saying that people's intuitive responses are likely to be reliabilist and that if they are reliabilists they should accept that the reliably produced true beliefs discussed by scientists are knowledge[31]. Those with reliabilist intuitions might accept that results from the sciences are relevant to epistemology because they reveal which beliefs are reliably produced by telling us about how reliable particular cognitive mechanisms are. The intuitions of these reliabilists would fit neatly with the view that knowledge is a natural kind that can be discovered by seeking an understanding of the reliably produced true beliefs referred to in scientific explanations. It is possible, then, to imagine a scenario in which intuitions and scientific theories could complement each other while being used to develop an epistemological theory. However, the trouble for Kornblith is that many epistemologists don't share his

[30] This intuition is likely to withstand Kornblith's suggestion that scientists are able to make successful predictions and explanations by attributing knowledge.

[31] That Kornblith himself is a reliabilist and has reliabilist intuitions cannot be denied, 'The account of knowledge I am offering here is a reliability account, and it is, at the same time, a naturalistic account' (2002: p. 63).

reliabilist intuitions, and he fails to supply an account of how those who do not share the intuitions should do epistemology.

For example, some epistemologists, namely internalists, think that a belief cannot be knowledge without the believing person being introspectively aware of the reasons why she has the belief. *Their* intuitions will provide a quite different starting point for epistemology and there is no guarantee that they will accept that results from psychology can inform them about the requirements for knowledge[32]. If intuitions have weight as a starting point in epistemology then it seems that these intuitions should have as much weight as the contrasting reliabilist intuitions. But then Kornblith must explain why the natural kind that cognitive ethologists discuss, or any other natural kind referred to by explanations in the sciences, should be of interest to epistemologists who do not have the intuition that the specific beliefs which he highlights that scientists discuss are knowledge[33].

It seems, then, that Kornblith's claim that knowledge is a natural kind, identifiable through our intuitions and the natural kind that scientists refer to in their explanations, will only have weight for those who share a number of his assumptions. Where

[32] Kornblith actually supplies some arguments against internalism in chapter four of *Knowledge and its Place in Nature*. He claims that beliefs produced by reflective cognitive processes are less likely to be accurate than beliefs produced by unreflective processes. He therefore argues, on the basis of empirical evidence, that internalism is not a good theory. However, to make internalists take note of these arguments, based on empirical evidence, Kornblith has to first supply reasons for thinking that results from psychology are relevant to epistemology. His method of establishing the relevance of results from psychology is to argue that knowledge is a natural kind. This is problematic because he does not provide reasons for those who are not intuitively reliabilist to think that knowledge is a natural kind. He therefore fails to supply good reasons for (a) thinking that knowledge is a natural kind, (b) thinking that results from psychology are relevant to epistemology, and therefore (c) that internalist intuitions are wrong because of what results from psychology say.

[33] Kornblith does provide some arguments for reliabilism at other parts in *Knowledge and its Place in Nature*, so he could argue that epistemologists should adopt reliabilism on the basis of these arguments and then accept that knowledge is the natural kind of reliably produced true beliefs discussed by scientists. However, if epistemologists accept that knowledge is a natural kind on the basis of arguments of this sort they won't be adopting the methodology Kornblith explicitly outlines because they won't be guided by their intuitions about obvious cases. Instead they will be guided by the arguments presented by another philosopher.

epistemologists do not share the intuition that all and/or only cases of reliably produced belief are knowledge, it is unclear why they should be interested in any natural kind, so it is unclear how they can both depend upon intuitions to discover a natural kind of knowledge and look to science to understand it. Without some better understanding of how and why epistemologists can combine their intuitions and results from the sciences while they develop their theories, it still seems as if epistemologists must either ultimately be guided by their intuitions about what counts as knowledge or by what science says about human cognitive systems because they will not be guided by both. Even those who accept that knowledge is the reliably produced true belief discussed by scientists, and who therefore buy into the idea that some animal belief is a part of a natural kind of knowledge, must do so either on the basis that they find it intuitive that knowledge has the characteristics of being true belief caused by a reliable mechanism or because they have a strong conviction, independent of their intuitions about cases, that science can identify a natural kind of knowledge. So it seems that if Kornblith is right that knowledge is a natural kind, it must either be the case that (a) intuitions are a good source of evidence about the natural kind, knowledge, or (b) there are independent reasons for accepting that the natural kinds discussed by science are knowledge. I aim to show that neither of these claims is convincing.

3. The Use of Intuitions as Evidence about Natural Kinds

Let us begin by considering the prospects for using intuitions to discover the proposed natural kind of knowledge. The first thing to say about this interpretation of Kornblith's methodology is that it is doesn't seem to be compatible with his rejection of conceptual analysis. There is a tension between the claim that epistemologists should use their

intuitions about obvious cases of knowledge as evidence and the claim that epistemologists should aim to understand the real phenomenon of knowledge rather than our concept of knowledge, especially as the criticism of the study of our concepts is based on the assumption that our concepts do not track the real phenomenon of knowledge. Kornblith does not explain why, when people study intuitions about obvious cases, they are not just learning about their disposition to apply their concept. When we have an intuition that something is knowledge, we seem to be acknowledging that it is appropriate to apply our concept. So if our concepts are unlikely to track reality, there is little reason for thinking that our intuitions will do any better, as intuitions seem merely to track our concepts. Kornblith might argue that our intuitions about obvious cases of knowledge are just judgements that an item is a member of the natural kind of knowledge[34], but this requires some account of why we should think that these judgements accurately reflect any real phenomenon rather than merely our concepts.

For those who deny that knowledge is a natural kind, it need not be a problem that our intuitions about knowledge track our concepts and not the real phenomenon of knowledge. This is because they can maintain that there is no natural kind independent of human interests to discover. They can allow that what we want to discover is the

[34] Some authors discussing this issue make this sort of move. Williamson (2007), for example, argues that our intuitive responses to thought experiments are really only ordinary modal judgements, based on a faculty reliably used for other modal judgements. A similar sort of defence of the use of intuitions as evidence in epistemology is given by Nagel (forthcoming). She argues that our intuitions about whether people described in thought experiments have knowledge are the product of a reliable mechanism of mind reading that we often use. If either of these theories is correct then there is some reason to accept that the cognitive mechanisms we use to assess whether someone has knowledge are reliable. However, even if Kornblith adopted a position like these he would not have an easy response to my criticism because it is still possible that our ordinary modal judgements and our mindreading capacities are guided by our concepts (the latter is perhaps more likely on a theory theory view of mind reading), in which case Kornblith would have to face the difficulty that we might be finding out about our concept of knowledge when we use our intuitions even if Williamson or Nagel are correct.

nature of our concept of knowledge[35]. But Kornblith explicitly denies that concepts should be the target of philosophy so he is required to supply an account of how and why our intuitions do track the real phenomenon of knowledge. He needs to explain why the intuition that some belief is obviously knowledge provides good evidence that it is.

Even if, however, we concede to Kornblith that intuitions somehow track obvious cases of knowledge, we do not thereby concede that he has provided a viable description of how epistemologists should construct their theories. This is because, even if our intuitions about obvious cases do track reality, there seems to be good reason to doubt that the generalisations based upon those intuitions about obvious cases will provide us with a trustworthy guide to the nature of knowledge (or other epistemic categories) because, for one thing, these intuitions do not seem to provide guidance about how to deal with less obvious cases.

For example, take the belief I have that I perceive a computer screen in front of me. Cases of perceptual belief where there is no impediment to the person perceiving seem to be obvious cases of knowledge. If we can hope for agreement in intuition about any belief being knowledge, it is likely to be the intuition that standard cases of perceptual belief are. But how do we generalise from obvious cases like these to other claims about knowledge in other, less obvious, cases? How do we ensure that we focus on relevant features of the example of perceptual knowledge, and other obvious cases of knowledge, making generalisations about the requirements of knowledge which truly

[35] This is the response that Alvin Goldman and Joel Pust would make (e.g. Goldman and Pust, 1998/2002; Goldman, 2007). As we saw in chapter three, they defend a position very similar to Kornblith's but maintain that concepts should be the target of intuition-driven epistemology. They can therefore avoid the accusation of inconsistency because they do not maintain, as Kornblith does, that intuitions can be evidence for epistemological theories *and* that concepts should not be the target of epistemology.

reflect the putative natural kind? A successful theory of knowledge has to be able to distinguish less clear cases of knowledge from non-knowledge as well as obvious cases of knowledge from obvious cases of non-knowledge. But even if, for some reason unspecified by Kornblith, we should be confident that we can recognise obvious cases of knowledge, there is reason to doubt that we can rely upon our intuitions about obvious cases to identify whether something is knowledge in other cases.

One reason that the thought experiment methodology in epistemology is often so stimulating is because those engaged in the methodology present cases which challenge our intuitive picture of knowledge by presenting controversial cases. So, for example, Keith Lehrer (1990) describes a situation in which a person has a small device implanted into his brain which is an accurate thermometer and produces correct thoughts about the current temperature. The implant causes Mr Truetemp to have accurate beliefs but he does not reflect upon how he ends up with the beliefs. Lehrer's thought experiment is designed to make the reader think that reliabilism does not work because they are supposed to have the intuition that Mr Truetemp does not have knowledge even though his belief is caused by a reliable process. What matters about this example for current purposes is that it is supposed to challenge reliabilist intuitions that people might have in response to seemingly obvious cases of knowledge, such as perceptual beliefs. And it is not clear how Kornblith would suggest that the epistemologist should respond to such complex scenarios which challenge intuitions about obvious cases. If knowledge is a natural kind, there should be an answer to the question of whether beliefs like Mr Truetemp's are knowledge. And if intuitions about obvious cases of knowledge provide evidence about the natural kind then they should tell us whether it is knowledge. But intuitions about obvious cases of knowledge (like

our perceptual beliefs) could be challenged on consideration of this counterexample to reliabilism; even those who prima facie think that knowledge is reliably produced true belief might ultimately have wavering or anti-reliabilist intuitions in response to the question of whether Mr Truetemp has knowledge. It is also easy to imagine how two people might agree in their intuition about an obvious case of knowledge, such as perceptual knowledge, but have different intuitions about whether Mr Truetemp has knowledge. This is not the result we would expect if knowledge were a natural kind that could be understood through our intuitive responses to obvious cases. It should be troubling to Kornblith that the intuitions that he suggests provide access to a natural kind of knowledge provide inadequate guidance about whether new cases should be classified as a part of the kind. It should also be troubling that people could have such different intuitions about non-obvious cases, even though they share the same intuitions about obvious cases, given that intuitions about obvious cases are supposed to provide evidence about the natural kind and therefore determine how we should respond to less obvious cases. All of this suggests that intuitions about obvious cases of knowledge provide a poor source of generalisations about what is required for knowledge because they can easily be challenged when we consider other cases. They therefore cannot be seen as providing a source of understanding about whether any belief is a member of the natural kind of knowledge.

I can imagine a number of ways Kornblith could respond to this criticism. Firstly, he might dispute the value of the example I have chosen to illustrate my point. He might argue that epistemologists should not be interested in examples such as Lehrer's Mr Truetemp because they should consider real life cases of knowledge and our intuitions will provide reliable guidance about these cases. He could claim that our intuitions

about whether real life cases are reliable, even though intuitions are unreliable guides when it comes to imagined thought experiments. There is, however, an easy way to respond to this argument: by providing a real life example which presents the same problem.

Here is such an example. Vinnie is an ordinary individual using an unconscious method of reasoning. His method of reasoning has been examined by psychologists who have discovered that it is highly reliable. It is so reliable that some psychologists claim that it must have been endowed by evolution through the process of natural selection. But people are not aware of how they come to a conclusion when they use this method of reasoning. This example is not far-fetched. People do use unconscious methods of reasoning and some evolutionary psychologists do argue that some of these are so reliable that they are likely to have been selected for (e.g. Gigerenzer, 2007), so Vinnie's case is just like a real life example of a person with true beliefs produced by a reliable mechanism without reflective awareness of the reasons for the beliefs. Kornblith's methodology should provide an account of how epistemologists should decide whether there is knowledge in cases like Vinnie's even if it is not required to provide an account of whether a person has knowledge in some imagined thought experiment. However, it seems as if our intuitions about obvious cases, such as our intuitions about perceptual beliefs, fail to provide clear guidance about Vinnie's case. One might have the intuition that a perceptual belief is knowledge because it is reliably produced but, when they are asked to respond to Vinnie's case, also have the intuition that a person cannot have knowledge even if they have a belief produced by a reliable mechanism without knowing the source of their beliefs. It also seems as if different people could have the same response to the obvious case of perceptual knowledge but

different responses to Vinnie's case; they could agree that perceptual beliefs are knowledge while not all of them accept that Vinnie has knowledge. So the problems that arise for Kornblith's claim that knowledge is a natural kind discoverable via intuitions about obvious cases arise with respect to real life cases as well as cases described by thought experiments.

A second response Kornblith could make would be to say that intuitions about the cases just described do not provide reasons to accept or reject any epistemological position because there is no general consensus in people's intuitions about the cases. He argues, after all, that 'our intuitions are merely obvious cases of the phenomenon under study. That they are obvious, and thus uncontroversial, is shown by the wide agreement that these examples command' (Kornblith, 2002. p. 12). He might therefore claim that where there is a clash of intuitions about the nature of the phenomenon of knowledge, none of the intuitions serve as evidence for or against a theory, because the case is controversial. According to this view, only intuitions about which people are in general agreement provide evidence for epistemological theories. It seems that Kornblith should resist this move, however, because it would be counterproductive for a philosopher who defends the use of intuitions in epistemology to claim that only intuitions about obvious cases are valuable. This is because much of the valuable work achieved by epistemologists involves considering the details of non-obvious cases. Intuitions about obvious cases, such as standard perceptual knowledge, are often compatible with various theories of knowledge. The intuition that we have knowledge in standard perceptual cases is compatible with the idea that there is knowledge because the perceiver is using a reliable process to form the belief *and* with the idea that there is knowledge because the knower is aware that the reasons for the beliefs are good ones

(she is aware that her visual system is usually reliable). This means that the intuition about this obvious case of knowledge is compatible with reliabilism but also with some versions of an opposing view, internalism. When our intuitions about obvious cases of knowledge are compatible with various opposing accounts of knowledge, examples such as Lehrer's Mr Truetemp which are less obviously knowledge or non-knowledge enable us to distinguish between various theories. The fact that people aren't all in agreement about whether Mr Truetemp has knowledge is not a sign that intuitive responses to the thought experiment are less useful than intuitive responses to obvious cases of knowledge; it is a sign that the thought experiment highlights important issues about what is and is not required for knowledge.

A final response Kornblith could make would be to say that intuitions are only used as corrigible evidence by those who adopt his methodology,

> What should we say about the rock collector's judgments at early stages of investigation, i.e. prior to any deep theoretical understanding of the features that make his samples samples of a given kind? Such judgments are, of course, corrigible, and they will change with the progress of the theory. What seemed like a clear case of a given kind in the absence of theoretical understanding may come to be a paradigm case of some different kind once the phenomena are better understood (ibid.: p. 13).

He claims that intuitions about the natural kind of knowledge are much like the rock collector's judgements at the early stage of his investigation; they are used as a starting point for investigation but can be discarded through the course of development of a theory. Kornblith could therefore claim that intuitions are not supposed to provide guidance about how to make generalisations which apply to non-obvious cases of knowledge because they only serve as a starting point. However, it is difficult to see how intuitions can really be of use even as a starting point in epistemology unless they provide an indication of how to make generalisations and draw conclusions about what

is required for something to be knowledge. If intuitions about obvious cases of knowledge provide poor guidance about how to make generalisations about the requirements for knowledge it is difficult to see how they can really contribute to progress in epistemology.

In sum, if we interpret Kornblith as claiming that intuitions will be decisive in epistemology because they will uncover the natural kind of knowledge then we see that his account runs into trouble. By relying on intuitions as a source of support for epistemological theories, Kornblith makes his position vulnerable to the criticism that it will provide theories that reflect our concepts rather than any natural kind. But even if we grant that intuitions can get some purchase on obvious cases of knowledge we have little reason to think that these intuitions will provide firm foundations for a general theory about knowledge, providing an indication of which beliefs are knowledge where there is some degree of ambiguity. We might concede that intuitions about obvious cases are reliable indicators about those cases but we have seen that intuitions about obvious cases do not provide guidance about how to classify less obvious cases, or to deal with disagreement between people with different intuitions about cases. So it seems that Kornblith is wrong to think that epistemologists should trust intuitions because they can track the real nature of the natural kind of knowledge. Even if they can reliably identify obvious members of a kind, they cannot be trusted to provide generalisations which can be used to identify whether other, less obvious cases are knowledge.

4. Why science doesn't have the answers

The arguments presented in section 3 challenged one specific aspect of Kornblith's account of knowledge as a natural kind, i.e. the claim that the natural kind can be

discovered via intuitions about obvious cases. Even if these arguments are successful, and intuitions cannot be guaranteed to identify the natural kind of knowledge, that natural kind might still exist. But if methodological naturalists are going to accept that knowledge is a natural kind they do not just need to see that it is possible that such a natural kind exists, they need good reasons to think that it really does. To establish that the natural kind does exist, Kornblith relies heavily upon the analogy between animal knowledge and human knowledge.

> I have argued that the category of reliably produced true belief plays an important explanatory role in cognitive ethology (Kornblith 2002). Reliable belief-producing processes are important to the science of ethology because there is an evolutionary explanation for their existence. … It is not surprising that ethologists interested in the evolutionary explanation of behavior have a term for reliably produced true belief: they call it knowledge. I have argued that this is not just a technical term within cognitive ethology. Rather, by looking at the ethology literature, we come to understand what knowledge is (Kornblith, 2011: p. 45-6).

In this passage Kornblith suggests that by looking at the literature on animal knowledge it is possible to understand what knowledge in general is. In the animal literature knowledge is associated with having true belief produced by a mechanism selected by the process of evolution that produces behaviour that satisfies our goals. His suggestion therefore seems to be that we can identify knowledge by understanding which beliefs are produced by cognitive equipment that is the product of evolution.

This interpretation is also supported by the following passage in which Kornblith associates knowledge with the proper functioning of a cognitive system:

> Our cognitive equipment should thus be viewed as having a function, and with any such defensible claim about function comes a commitment to a notion of proper functioning, the idea that there are certain standards that must be met if the equipment is to fulfil its proper function. To repeat: these standards are not a projection of human parochial interests and concerns. They are, instead, objective features of the world. But this is just to say that the mere fact that talk of knowledge embodies certain standards provides us

> with no argument at all that knowledge is not a natural kind (Kornblith, 2011: p. 48).

His claim here seems to be that human knowledge is true belief produced by cognitive equipment fulfilling its proper function as designed by evolution, just as it is in the animal case. On this view of knowledge, epistemologists should aim to understand what fits into the category of true belief produced by a mechanism fulfilling its proper function[36]. It seems, then, as if there is potential for Kornblith to avoid the criticism that his view is based on the idea that intuitions are likely to reveal the natural kind of knowledge because he is actually committed to the idea that science has the answers.

If it is right that the beliefs produced by cognitive equipment fulfilling its proper function are knowledge then epistemologists could focus on identifying the beliefs produced in this way rather than trying to make generalisations about knowledge based on their intuitions. However, I aim to show that epistemologists should reject this move as well. More generally, they should not expect to be able to identify knowledge without consideration of how the concept is required to meet human needs. To persuade the reader of this point I will present two examples aimed to show that we should not classify all beliefs meeting Kornblith's criteria as knowledge, and certainly not simply

[36] Ram Neta (2008) makes a claim very similar to Kornblith's one here:

> Epistemology, so naturalized, is both empirically informed and substantively prescriptive. Just as it is an empirical question how our other functional systems (e.g. respiratory, circulatory) are supposed to work, so too is it an empirical question how cognition is supposed to work (p. 338).

In this passage Neta suggests that we can find out about epistemic excellence by discovering the proper function or 'goal-state' of a cognitive system, that is, the proper function of those systems of the human body designed to discover facts about the world. His main point of departure from Kornblith is in his characterisation of the goal state of human cognitive systems. Whereas Kornblith thinks human cognitive systems have the same proper function as animal cognitive systems, and human knowledge can therefore be discovered by understanding the proper function of animal cognitive systems, Neta thinks that the cognitive systems of inquisitive creatures like humans is different to those of non-human animals. He thinks that we can understand the proper function or goal-state of human cognition by understanding success in the sciences. But this departure from Kornblith's ideas does not change the fact that Neta, like Kornblith, is committed to the idea that we can understand knowledge by understanding the proper function of a cognitive system. This means that if my arguments against Kornblith's claim that knowledge is belief produced by a cognitive system functioning properly are successful, then they will challenge Neta as well as Kornblith because they will challenge the relationship between knowledge or epistemic excellence and proper functioning.

because they meet the criteria, and without considering what work can be done by our epistemic classifications. Hopefully the reader will agree that the best possible response to these examples is to reject the idea that knowledge is a natural kind altogether because human needs decide whether something should count as knowledge.

5. Example One

I take Kornblith's claim to be, then, that a cognitive system used to perform a task that it was designed by natural selection to perform would produce knowledge whenever it produced true belief. According to this view, being produced by cognitive equipment fulfilling its proper function is sufficient for being knowledge. My goal in this section is to show that this is not the case because whether or not any belief deserves the status of knowledge will depend upon whether human needs are answered by categorising it as such, even when that belief is produced by cognitive equipment fulfilling its proper function. Consider the following example which illustrates this point.

A couple use the recognition heuristic to decide upon which of various companies would be a good one in which to invest all their savings. The recognition heuristic is a belief forming process or reasoning strategy described by psychologists, including those who are concerned with developing evolutionary explanations of human cognition. When using this heuristic, people select the first option they recognise from a number of different options. Gigerenzer and his colleagues used the heuristic on the stock market, building a portfolio of shares based not on detailed information about the companies they were investing in or the nature of the market in general, but based solely on whether they recognised the name of the company (Borg et al., 1999). Their portfolio did well, producing a good profit, suggesting that people can achieve good results by using the recognition heuristic. On the basis of findings such as these about how

successful people can be using heuristics such as the recognition heuristic, Gigerenzer and his colleagues have made a compelling case that people can obtain true beliefs while considering only a limited amount of information. They argue that these belief forming processes are likely to be the products of evolution, selected for during the process of natural selection. Beliefs produced by the recognition heuristic are therefore candidates for the status of beliefs produced by reliable mechanisms which are performing their proper function. The proper function of the cognitive mechanism is enabling people to get the information they need from the environment without requiring them to consider a large amount of information. This all suggests that on Kornblith's view the couple using the recognition heuristic would acquire knowledge if their beliefs are true.

Now further imagine that the couple are using the recognition heuristic immediately after the BP oil spill in 2010 and they are considering investing in an oil firm[37]. They happen to browse over a list of oil companies and see ESSO on the first page. They decide to invest in this company because they recognise the brand name. They invest in the company and make a tidy profit. According to Kornblith's view, it seems as if they *knew* that ESSO were a good company to invest in. Prima facie this conclusion might have some appeal because successful oil companies are likely to be high profile so people will recognise them and can use their recognition as evidence that the company is successful, and thus a good prospect to invest in.

But consider a further detail. The couple happen to have browsed only the second page of the list of companies. The pages of the list happened to be stapled in the wrong order. Had the list been stapled in the right order they would have seen the BP company name

[37] Thanks to George Botterill for pointing out the relevance of examples such as this.

first. They would have recognised that. They would then have concluded that BP were a good company to invest in and acted on this conclusion. Would they have known that BP provided a good investment opportunity? Well, given that there was an oil spill which increased the profile of the company but also decreased its profitability, it seems that they would not have known that the company was a good one to invest in on the basis of the recognition heuristic. They could have recognised the name of the oil company for a reason other than that it was a profitable company. And in the end they would have had a false belief because the company would not have been a good one to invest in because its profits were likely to drop as the full consequences of the oil spill became apparent.

I think this example can be used to illustrate that something does not acquire knowledge-status simply by being produced by cognitive equipment fulfilling its proper function. Whether one classifies the belief that ESSO is a good company to invest in as knowledge will depend upon what goals one wants to achieve through this classification. One thing that might make it seem implausible that the belief should count as knowledge is if one values being able to distinguish beliefs that are only accidentally true and which would have been false in many close possible worlds from those that are not only accidentally true. If this influences one's classification of something as knowledge then the ESSO belief might not count as knowledge because there are close possible worlds in which the couple ended up considering and recognising BP first, concluding that it was a good company to invest in. There are also close possible worlds in which it is ESSO rather than BP who had the oil spill, so the couple could come to the same conclusion in the same way but end up with a false belief. Even though the ESSO belief is produced by cognitive equipment which is

fulfilling its proper function, it will not necessarily count as knowledge. What will decide how it is classified is what use we can make of classifying it either way. It depends upon what work we can do by distinguishing between beliefs which are only accidentally true and those which would be true in all close possible worlds, or some other work that could usefully be done by either including or excluding the belief from the category of knowledge. This all suggests that being a part of the natural kind of beliefs produced by cognitive equipment fulfilling its proper function is not sufficient for knowledge because a belief can meet these criteria but only deserve its status as knowledge/non-knowledge because humans are able to answer to their own needs by classifying it in this way.

One way that Kornblith might respond to the criticism just mounted against him would be to claim that not all beliefs produced by the recognition heuristic are knowledge. He could argue that the heuristics described by Gigerenzer and his colleagues would have been selected to deal with certain types of information and that they would only produce knowledge when used to process this sort of information. They could deny that the heuristics are performing their proper function when processing information about oil companies so, even according to Kornblith's view, they do not produce knowledge about what oil company is best to invest in. Evolutionary theory actually provides a way to articulate this response. It is said that our cognitive systems evolved to aid the survival and reproduction of Stone Age Man. The environment of the Stone Age man would have been very different to our own, for example, it would not have included oil companies. This means that if the proper function of the heuristic reasoning processes was to process information in the Stone Age it would not have been performing its proper function in the ESSO case. One way Kornblith might respond to the criticism

levelled against his theory, then, is to say that in all cases where cognitive equipment performs its proper function and produces true belief that belief is knowledge because of its kind membership but in the ESSO case the cognitive system is not fulfilling its proper function. Then the ESSO case would not work as a counterexample to the claim that belief produced by a mechanism fulfilling its proper function is knowledge because it is a member of a natural kind. The problem with this response is that rather than showing the strength of Kornblith's position it highlights another undesirable consequence it has.

6. Example Two

If only those beliefs which are a part of the natural kind of beliefs produced by cognitive equipment fulfilling its proper function are knowledge, and the proper function is decided by the environment of Stone Age our ancestors, then it seems as if beliefs about many ordinary modern day objects and issues could be completely excluded from the category of knowledge. On the strict construal of the proper function of a cognitive mechanism described at the end of section 2, the mechanism is only performing its proper function when it is processing information that would have been available to Stone Age man. But given how different conditions for Stone Age man would have been, the proper function of the cognitive system would not be to process information about modern objects or issues such as laws, courts, schooling, etc (Pinker, 1997: p. 42). This means that if we accepted as knowledge only beliefs produced by cognitive systems performing their proper function we would have to deny that we could have beliefs about laws, courts, schooling, etc.

This is a highly undesirable consequence of any account of knowledge. If it is not possible to have knowledge about laws, courts or schooling, then you cannot claim to know that the law should protect the rights of ethnic minority groups. You cannot claim to know that the principle of 'innocent until proven guilty' should be applied in all fair courts. You cannot claim to know that your child would be better off going to a __school (fill in the blank with 'private/independent/public/state/faith/non-faith'). This is just the tip of the iceberg. On this view of knowledge, it is impossible to know about all those things that would not have been around for Stone Age man to deal with. All beliefs about these matters should be excluded from the category of knowledge because a cognitive mechanism cannot be fulfilling its proper function by processing information about these matters.

Once again, I would suggest that the problem that the natural kinds approach encounters here is a result of the failure to recognise how the concept of knowledge has to answer to human needs. We require the concept to provide distinctions between beliefs about all sorts of matters that Stone Age man would not have had exposure to. An account that fails to take this into consideration risks producing seriously unacceptable consequences. It could imply that we simply couldn't make claims about whether our own or other people's beliefs about many important matters are knowledge. Epistemologists should, in my view, attempt to avoid this sceptical consequence and conclude that it is not necessary for a belief to be produced by cognitive equipment fulfilling its proper function for it to be knowledge because if this were the case then the concept of knowledge would not apply to beliefs about many important matters. Knowledge has certain roles in human life and epistemologists should take these into consideration when developing an account of knowledge. They should recognise that an

account of knowledge is only satisfactory if it allows the concept of knowledge to meet certain human needs. In the process of recognising this they should reject the idea that knowledge is a natural kind independent of human interests.

7. Kornblith on the Social Role of Knowledge

One might wonder why Kornblith does not worry about the human needs that are answered by the concept of knowledge. The explanation for this is that he thinks that knowledge is able to answer to human needs because it is the natural kind that it is:

> I don't think that it is at all implausible to think that knowledge, like gold, plays an important social role, and yet its ability to play this social role is explained by deeper features that it has, features it has independently of the social role it plays (Kornblith, 2011: p. 45).
>
> The social role that knowledge plays is not what makes it the kind of thing it is. Rather, the kind of thing that knowledge is explains why it is well-suited to play the social role it does for human beings, a role it does not and cannot play for non-human knowers (Kornblith, 2011: p. 46)

What Kornblith claims in these passages is that the social role of knowledge comes about as a result of knowledge being a natural kind of reliably produced true belief. He refers to the usefulness of the category of reliably produced true beliefs which is applied to animals. Humans have knowledge like animals have knowledge, on his view, although it happens that it can be useful for humans to use the concept of knowledge to identify the phenomenon, so that the concept can incidentally fulfil the additional social role that it doesn't have in the animal case. But the social role of the concept of knowledge is parasitic on the natural kind. We can understand the natural kind, thereby grasping what knowledge is, without considering the social role that the concept which identifies the natural kind has.

What my arguments in the previous sections show is that Kornblith is too optimistic when he assumes that a natural kind of reliably produced true belief fulfils the social roles it is required to perform and meets human needs. If we want our needs to be met by the concept of knowledge then we should resist the analogy with animal knowledge. We should deny that human knowledge is simply belief produced by cognitive equipment fulfilling its proper function. We should instead see that knowledge is belief that we can usefully classify as knowledge to meet certain human needs. We should not prioritise identifying the natural category of reliably produced true beliefs that Kornblith discusses, while ignoring how the category must answer to human needs, because we will not be able to meet our needs by doing so and without fulfilling our needs we will be left with an unsatisfactory account of knowledge.

Conclusion

This chapter has aimed to challenge Kornblith's claim that knowledge is a natural kind by suggesting that we cannot discover the nature of knowledge simply by using intuitions and/or borrowing classifications from the sciences and assuming that knowledge is belief produced by cognitive equipment fulfilling its proper function. In my view, rather than using intuitions or borrowing distinctions from the sciences and accepting that knowledge is a natural kind referred to by scientists in their explanations, epistemologists ought to recognise that there are certain important uses we can make of our epistemological distinctions and that we should only borrow a distinction from the sciences if we can fulfil our epistemic goals by doing so. They should engage in the project of understanding what valuable work can be done by our epistemological classifications. They should use the conclusions they draw to decide between competing intuitions and to decide which scientific distinctions to borrow from because any theory

that fails to enable us to make certain important distinctions (such as about beliefs about courts) will be an unsatisfactory epistemological theory. But if this is right then knowledge is not a natural kind, the classification of beliefs as knowledge is determined by human interests. Intuitions and distinctions from the sciences are relevant to epistemology only to the extent that they enable epistemologists to develop theories that reflect the distinctive goals we have when making our epistemic classifications. They are not relevant to epistemology because knowledge is a natural kind because it is, in fact, a kind defined so that it can answer certain human needs and *not* a natural kind.

Chapter Five: How to be a methodological naturalist

Introduction

In this thesis so far I have argued that epistemologists should be methodological naturalists. They should take the outcomes of scientific theorising into consideration when developing accounts of knowledge, rationality and justified belief. But I have found that there are problems with the various accounts of the methodology of epistemology consistent with methodological naturalism currently on the table. The task of this chapter is to outline an account of the methodology epistemologists should adopt that is consistent with methodological naturalism but avoids the various problems with other accounts.

Amongst those who are concerned with the correct methods to use in epistemology, some philosophers have recently defended the idea that epistemologists should study the goals that can be achieved through people's epistemic classifications. For example, Edward Craig (1990) and Jonathan Weinberg (2006) argue that epistemologists should consider which goals people can usefully achieve by classifying some things as knowledge, rationality and justified belief. They argue that one should only accept an epistemological theory that describes these categories in such a way that people can achieve valuable goals by applying the categories. I call these approaches 'teleological approaches to epistemology'[38]. In this chapter I aim to show that epistemologists should

[38] My account is to be distinguished from Ram Neta's similarly titled 'teleological account of knowledge'. Neta's (2007) view is very similar to Hilary Kornblith's (2002). He suggests that knowing is having a belief produced by a cognitive system achieving its goal-state or proper function. He therefore thinks that epistemologists should aim to understand the function of a cognitive system. My claim is rather different. I think that epistemologists should aim to understand the function or goal of our epistemic classifications. The *telos* that I discuss is of the concept of knowledge, whereas the *telos* that Neta discusses is of the cognitive system.

adopt a teleological approach to epistemology. However, I will argue that they should not go fully into either Craig or Weinberg's camps because neither of their accounts recognises the full role that results from the scientific study of human psychology can perform in discovering the goals that can usefully be achieved through our epistemic classifications. In other words, I shall argue that epistemologists should adopt a teleological approach to epistemology, but I shall defend an approach which is more radically naturalistic than either of the approaches currently on the table. This approach will combine all of the benefits of teleological approaches to epistemology with all the very best features of methodological naturalism, providing a strong methodology to use in epistemology.

1. An introduction to teleological approaches to epistemology

In chapter four, we found that epistemic concepts, such as knowledge, are required to answer to human needs. In my discussion of the claim that knowledge is a natural kind, I claimed that epistemic categories are not wholly natural because they are used by humans to achieve certain important goals. An account of knowledge according to which the concept of knowledge could not be used to identify good beliefs about courts, schools, policing, etc. would be an inadequate one because we require the concept to distinguish good from bad beliefs about these things. Those who defend teleological approaches to epistemology address this issue headlong, which is a distinct advantage of their position.

According to defenders of a teleological approach to epistemology, epistemologists should engage in the project of understanding the goals that can be achieved and the human needs met by our epistemic classifications. Edward Craig (1990) is a key proponent of this view of the methods to be used in epistemology. He argues that

epistemologists should not engage in the project of conceptual analysis, aiming to discover the necessary and sufficient conditions for the application of a concept, and instead they should aim to understand what work can be done by classifying something as knowledge.

> Instead of beginning with ordinary usage, we begin with an ordinary situation. We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application (Craig, 1990: p. 2).

He imagines a state of nature, adopting a framework used in political theory, and asks what purpose the concept of knowledge would have served in this state.

> [...] we may think of our starting point as state-of-nature theory: we are to apply similar methods to a different object: not political, but conceptual and linguistic institutions (Craig, 1990: p. 10).

He hypothesises that the concept would have emerged to tag reliable informants. People would not have been able to access all the information they needed themselves, so they would have needed to identify other people who could reliably inform them.

> On any issue, some informants will be better than others, more likely to supply a true belief. (Fred, who is up a tree, is more likely to tell me the truth as to the whereabouts of the tiger than Mabel, who is in the cave.) So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use. The hypothesis I wish to try out is that the concept of knowledge is one of them. To put it briefly and roughly, the concept of knowledge is used to flag approved sources of information (Craig, 1990: p. 11).

He then argues that any theory of knowledge should describe the concept in such a way that it could be used to tag reliable informants, achieving the goal that the concept emerged to perform.

Applying this method to a contemporary account of knowledge, Craig asks whether the causal theory of knowledge describes knowledge in such a way that the concept could

be used to identify and tag reliable informants. According to the causal theory, a belief is knowledge if it is caused by the fact known. Craig suggests that the causal theory of knowledge is unsatisfactory because it would suggest that the concept could not have emerged to perform the role of identifying reliable informants. When people use the concept of knowledge to identify reliable informants, they often do not know what has caused the belief, so they cannot only be applying the concept to identify beliefs caused by the fact known. On the other hand, it seems that in many cases when people are sure that a belief is caused by the fact known they do not find it necessary to apply the concept of knowledge to identify reliable informants because they already have access to all the information that they need, for example, because they have seen the phenomenon that caused the person to have the belief. Because people are unlikely to be tagged as a reliable informant because they have a belief caused by the fact known, Craig thinks that the concept of knowledge cannot be what the causal theory of knowledge says it is.

Craig's account is characteristic of what I shall call in this chapter a teleological approach to epistemology because of the concern for the *telos* or goals that can be achieved by applying the concept of knowledge. Jonathan Weinberg (2006) has also developed a teleological approach to epistemology, which he has called 'reconstructive neo-pragmatism'. He also thinks that epistemologists should consider the goals that can be achieved by classifying something as knowledge and should only accept an account of knowledge which answers to human needs. But the methodology he describes differs from Craig's in two important respects.

Firstly, Weinberg claims that epistemologists should not engage in a genealogical approach like Craig's state of nature approach. For Weinberg, epistemologists should

take a future-orientated approach; they should consider what goals they would like their epistemic classifications to achieve in the future. To see the difference between the two positions, compare the following quotes, the first from Craig and the second from Weinberg:

> Such an investigation would still have an anchorage point in the everyday concept: should it reach a result quite different from the intuitive intension, or one that yielded an extension quite different from the intuitive extension, then, barring some special and especially plausible explanation of the mismatch, the original hypothesis about the role that the concept plays in our life would of course be the first casualty. For it is not the idea to construct an imaginary concept, but to illuminate the one we actually have, though it be vague or even inconsistent; and to illuminate it by showing that a concept with the hypothesised role would have characteristics closely resembling those that it exhibits itself (Craig, 1990: p. 2).

> My operative question will be: were we to consider a radical re-constitution of our epistemic norms, what would we include, what might we strengthen, and what might we abandon as outmoded?' (Weinberg, 2006: p.66).

Craig aims to understand the everyday concept of knowledge and thinks that our understanding of the goals to be achieved through our epistemic classifications has to match the intuitive extension of the concept. In contrast, Weinberg aims to provide an account of what our epistemic concepts or norms *should* be and not how they are. He discusses how we would constitute them if we were to radically re-think them. He argues that epistemologists should imagine that they are beginning afresh and could select whichever epistemic norms or categories they wanted rather than thinking about how we ended up classifying what we do classify as knowledge, rationality or justified belief. In Weinberg's view, epistemologists need to consider all the goals that could be achieved by distinguishing between various beliefs and should accept the account of knowledge according to which humans can achieve the most valuable goals. Whether internalism, externalism, evidentialism or any other theory about the nature of knowledge (rationality or justification) is best will depend upon what goals can be

achieved by making the distinctions supplied by these theories. For example, if it is useful for us to be able to distinguish between beliefs for which reasons can and cannot be given then internalism seems *prima facie* to be valuable.

Secondly, Weinberg develops Craig's idea by suggesting that epistemology should be interdisciplinary. It should be consistent with results from the social sciences and psychology. Weinberg claims that results from the social sciences can be relevant to epistemology because they can reveal the nature of our social structures. By understanding the social structures in place in our society we can come to understand what roles our concepts need to fulfil within those social structures. He also claims that results from psychology can be relevant to epistemology because, by taking these results into consideration, epistemologists can ensure that they do not endorse an epistemological theory which is too demanding and places knowledge, rationality and justified belief beyond the reach of ordinary human beings with their limited cognitive capacities.

There are a number of important positive features of Weinberg's account. It shares the advantage of Craig's methodology of acknowledging that our epistemic classifications are required to answer to human needs (explored more in section 2.2). It also has advantages over Craig's account. By taking a future-oriented approach, those who use Weinberg's methodology are able to concentrate energies on considering what goals can be achieved by making epistemic classifications at the current time. This will be more useful than understanding what goals could be achieved in the state of nature or even the goals that we currently achieve by applying our epistemic concepts. This is because our epistemic situation has developed over time so we will require our concepts to answer to different needs to our ancestors. It is also quite plausible that our existing

epistemic concepts do not enable us to achieve valuable goals, as will be explored further in section 2.1. As the utility of epistemic classifications is paramount in a teleological approach to epistemology, it certainly seems to be an advantage of Weinberg's modifications to Craig's methodology that they are likely to produce an account of knowledge according to which our epistemic classifications will be more useful.

In addition to this, as we have seen throughout this thesis, there are great gains to be made by epistemologists considering results from the scientific study of human psychology. One of these gains is that by considering such results it is possible to ensure that one does not develop an account of knowledge, rationality or justified belief that is beyond the reach of ordinary humans so that it has the sceptical consequence that no-one can ever have knowledge, be rational, or have justified beliefs. These advantages of Weinberg's account mean that I will focus on his project rather than Craig's, although much is owed to the latter.

2. The benefits of a teleological approach to epistemology

There are numerous great benefits that come from adopting a teleological approach to epistemology, as compared to alternative approaches to the methodology to be used in epistemology. In this section we will explore these advantages.

2.1. Answering to human needs

The first advantage of a teleological approach to epistemology has already been hinted at: it ensures that the accounts of knowledge, rationality and justified belief it produces answer to human needs. There are various goals that we aim to achieve through our

epistemic classifications. We often want to identify those people upon whom we can trust. We want to be able to work out which of other people's beliefs about important matters we should adopt and which we should ignore or dismiss. The teleological approach has it built in as a criterion for success that any account of knowledge, rationality or justified belief must describe epistemic concepts that answer to these and other human needs. It involves examining the needs that can be met by applying our epistemic concepts and checking whether any theory can meet those needs. The first, and perhaps most obvious, advantage of a teleological approach to epistemology is, then, that it provides an account of knowledge, rationality or justified belief which enables us to achieve valuable goals by applying our epistemic concepts. There is always room for human fallibility, but given that people adopting the teleological approach are explicitly aiming to understand the goals we can achieve by classifying something as knowledge, the chances are greater that they will provide a theory of knowledge with utility for humans.

2.2. Not mere conceptual analysis

Aside from the central advantage of teleological approaches to epistemology described in 1 and 2.1, there are a number of advantages of adopting the approach that come from it being a viable alternative to other methods currently used in epistemology. For example, those who adopt it are able to provide an account of how people can understand epistemic categories without engaging in conceptual analysis. Recall that the aim of those who adopt the approach is not to discover the necessary and sufficient conditions for the application of the concept of knowledge. It is to understand the goals that can be achieved by classifying something as knowledge.

As we have seen, there are various reasons why conceptual analysis is a problematic method to use in epistemology. We have seen that results from the scientific study of human psychology reveal the shortcomings of some standard epistemic concepts. They show, for example, that the standard concept of rationality, as involving checking the consistency of any new information with the beliefs in our belief system (Stein, 1996), is beyond the reach of ordinary thinkers because it places a constraint on rationality that people simply cannot meet because of the human finitary condition. People cannot possibly check the consistency of all their beliefs, or even identify those beliefs that they must check the consistency of, because of limits to their computational powers (Cherniak, 1986). Because the standard concept of rationality will not apply to ordinary thinkers, the use of that concept has the sceptical consequence that people cannot be rational, so the concept cannot be used to identify some and not other people as rational. This means that, for our conception of rationality to be useful, epistemologists need to move beyond conceptual analysis, because by using this methodology they will only discover the necessary and sufficient conditions for concepts such as the standard concept of rationality, which are not necessarily very useful at all.

By considering the goals that can be achieved by applying our epistemic concepts, Weinberg and Craig are not restricted to considering the conditions under which people tend to apply their concepts. They can engage in a practice of critically evaluating epistemic concepts such as the standard concept of rationality, asking how we should use our epistemic concepts to achieve valuable goals rather than merely aiming to discover how we happen to use them. Instead of engaging in the stagnant process of revealing the nature of our potentially flawed concepts, they can actively engage in the

project of working out what goals we can achieve by classifying some things into our epistemic categories.

2.3. No search for necessary and sufficient conditions

A related benefit of the teleological approach to epistemology is that it does not require there to be necessary and sufficient conditions for the application of our epistemic concepts. Those who engage in conceptual analysis are committed to the idea that concepts have a certain structure because they aim to capture the necessary and sufficient conditions under which it is appropriate to apply a concept. This assumption has been challenged by empirical findings from psychology about people's categorisation judgements. Results from psychology suggest that people do not have a list of necessary and sufficient conditions stored in their heads. Instead they have prototypes or exemplars stored and make a judgement about the similarity of any item to the prototype or exemplar—if an item is sufficiently similar to the prototype or exemplar then it will be classified under a concept (e.g. Rosch and Mervis, 1975; Smith and Medin, 1981). These findings suggest that the search for necessary and sufficient conditions for the application of a concept is likely to be futile because there are no necessary and sufficient conditions to be found. This in turn suggests that those philosophers who engage in conceptual analysis, and are therefore committed to the idea that there are necessary and sufficient conditions for the application of any concept, are in trouble.

Brian Weatherson (2003) has challenged the relevance of these psychological findings to epistemology, arguing that epistemic concepts do have necessary and sufficient conditions. He claims that if similarity judgements determined whether a person applied

a concept, then we would find it impossible to outline necessary and sufficient conditions for the application of a concept. It should be possible to find examples that we would classify as knowledge, for example, although they do not meet the necessary and sufficient conditions described for the application of the concept of knowledge. What we actually find is that philosophers are able to propose necessary and sufficient conditions for the application of epistemic concepts relatively easily. The reason that it is difficult to provide a successful analysis of epistemic concepts is that philosophers are also able to find counterexamples where those necessary and sufficient conditions are met but where it does not seem to be right to apply a concept, for example, because the case does not seem as if it is knowledge. For Weatherson, conceptual analyses of knowledge do not fail because it is impossible to describe the necessary and sufficient conditions for the application of the concept; they fail because some things that meet those conditions do not seem as if they should count as knowledge. In his view, the psychological results do not bear upon epistemology because epistemic concepts do take the form of necessary and sufficient conditions, so epistemologists can sensibly seek an understanding of these conditions.

I have two responses. Firstly, Weatherson does not provide good reasons for thinking that the psychological theory of concepts does not apply to those specific concepts epistemologists discuss. If he could provide empirical evidence to counter the empirical results from the scientific study of human psychology, or if he could provide a strong argument to show that the psychological results are irrelevant to our understanding of epistemic concepts, then he would provide a strong case for treating epistemic concepts differently. As it stands, he presents weak armchair arguments against empirical results. He just highlights the fact that epistemologists are able to come up with a list of

necessary and sufficient conditions for the application of the concept knowledge. He does not show that the list really captures the conditions under which people do apply the concept. Given the psychological findings, there seems to be good reason for thinking that epistemic concepts do *not* have such necessary and sufficient application conditions and that epistemologists list conditions which do not really reflect people's categorisation judgements. This means that rather than assuming that epistemic concepts are different from other concepts because there are necessary and sufficient conditions for their application, it seems more plausible that epistemologists are not getting a proper understanding of people's concepts when they find necessary and sufficient conditions for the application of some concepts. They are able to come up with a list of necessary and sufficient conditions, but these are unlikely to reflect the real structure of our epistemic concepts. This would explain why no successful analysis of any of our epistemic concepts has yet been found: epistemologists are proposing conditions for concepts which actually fail to reflect people's categorisation judgements so people have intuitions which do not agree with the conditions proposed. Consequently, they reject any proposed analysis.

Secondly, results from experimental philosophy actually contradict Weatherson's claim that epistemic concepts are likely to have necessary and sufficient conditions for their application. As was first mentioned in chapter three, Swain et al (2008) presented participants in an experiment with a number of vignettes and asked them to say whether the people described in the vignettes have knowledge. Some of the vignettes were of clear cases of knowledge, some were of cases where it was unclear whether a person had knowledge or not, others were of cases that were clearly not knowledge. Participants were shown these vignettes in various orders so that the experimenters

could evaluate whether intuitions about epistemological examples are subject to ordering effects. The results of the experiment suggest that people's ascriptions of knowledge *can* be influenced by the order in which they are presented with thought experiments. They showed that people's responses to the same thought experiment could differ depending upon what thought experiment they had previously been shown. For example, people were more likely to ascribe knowledge to a character in a thought experiment presented after a case that clearly was not knowledge than they were to the same character in the same thought experiment when it is presented after an obvious case of knowledge. What is most important about these findings for current purposes is that they provide support for the claim that people are not checking a list of necessary and sufficient conditions for the application of a concept. If they were checking a list of characteristics it seems as if they would respond to a single scenario in the same way whether it is preceded by a clear case of knowledge or a clear case in which a person does not have knowledge. In each case they would check for and find the same features from the list of necessary and sufficient conditions for knowledge. They would come to the same conclusion about whether a belief belongs to the category of knowledge. If, on the other hand, participants in the experiment were assessing the similarity of the various cases to some prototype or exemplar of knowledge it would be easy to see how they could end up classifying cases preceded by clear cases of knowledge as not being knowledge—their judgements could be influenced by the fact that the case they are evaluating is less similar to the prototype or exemplar than the previous one[39].

[39] This is, of course, only one hypothesis that should be tested empirically. However, it seems to me that we are more likely to explain why people's responses can be affected by order effects if we accept that when they apply the concept of knowledge they are *not* checking the item against a list of characteristics that are required of knowledge.

I do not expect these responses to Weatherson to conclusively prove that there are not necessary and sufficient conditions for the application of the concept of knowledge and that epistemic classifications take the form of similarity judgements. But I think that they provide additional reason to prefer a methodology which isn't only viable if necessary and sufficient conditions can be found for concepts like knowledge. It is therefore an advantage of teleological approaches to epistemology that they provide an alternative to conceptual analysis. It is possible to achieve an understanding of the goals that people can achieve by classifying something as knowledge even if the application of the concept of knowledge isn't governed by a list of necessary and sufficient conditions. This means that the teleological approach looks strong no matter what the scientific study of concept attribution says about how we make categorical judgements.

2.4. Intuitions not used as direct evidence

Those who defend the teleological approach to epistemology are also able to avoid using intuitions as direct evidence for or against a theory in epistemology. There are various problems with the use of intuitions as evidence in epistemology, as we found in chapter three. The first is that intuitions only seem to provide access to our concepts and, as was suggested in section 2.2, it seems that epistemologists should be interested in knowledge (or rationality or justified belief) rather than our potentially flawed concept of knowledge (or rationality or justified belief) (Stich, 1993; Kornblith, 2002). The second problem is that intuitions seem to be influenced by factors irrelevant to epistemology, such as social and cultural background, ordering effects, and even the environment in which a person considers a thought experiment (e.g. Weinberg, Nichols and Stich, 2001; Swain et al. 2008). The third problem is that there seems to be no way to calibrate philosophical intuitions. In the sciences it is possible to calibrate data by

gathering other data using different instruments but there is no source of calibration for philosophical intuitions (Cummins, 1998). All that philosophers can use to check whether our intuitions are correct are other judgements we would make about an example, and these judgements are likely to be influenced by the same sorts of factors as our intuitions are. Given that intuitions seem to be influenced by factors irrelevant to philosophy, there seems to be a special demand to be able to calibrate intuitions to establish which (if any!) are correct, but this is a demand that cannot be met. Altogether, then, it seems to be an advantage if an epistemological methodology does not involve the use of intuitions as direct evidence for or against any account.

Happily for the defender of the teleological approach to epistemology, on their approach intuitions are not used as a source of direct evidence for or against any account. What will decide whether knowledge is justified true belief, for example, is whether people can achieve important goals by classifying beliefs which are justified and true in one category and distinguishing them from all other beliefs. The fact that people simply do or do not have the intuition that a person with justifiably produced true belief has knowledge in response to the Gettier thought experiment (Gettier, 1963) will not be decisive. So, whereas much of traditional analytic epistemology is criticised on the basis that philosophers use intuitions as a brute source of evidential support for or against epistemological theories, defenders of the teleological approach can avoid this criticism because they do not depend upon intuitions as brute evidence. What decides for them whether a theory is correct is something that can be discussed and debated, that is, an evaluation of whether it is useful to make a certain type of classification or distinction. It is not a brute response to a philosophical thought experiment.

2.5. Required continuity with folk concepts

The teleological approach to epistemology has one final, important benefit: it ensures just the right degree of continuity between folk concepts and the concepts described by epistemologists. Recall that in chapter three we saw how Alvin Goldman has highlighted the importance of ensuring that there is a degree of continuity between folk concepts and those concepts deployed in our epistemological theories[40]:

> Whatever else epistemology might proceed to do, it should at least have its roots in the concepts and practices of the folk. If these roots are utterly rejected and abandoned, by what rights would the new discipline call itself 'epistemology' at all? (Goldman, 1993: p. 272)

Goldman's claim is that there is a distinctive subject matter for epistemology—the folk concepts of knowledge, justification, etc.—and that any theory that departs too far from folk concepts should not really be called an epistemological theory. Goldman's solution is to defend the contribution of conceptual analysis. He claims that if epistemologists begin by analysing concepts, their theories will be continuous with the concepts of the folk. However, as we have found, there are various good reasons to reject the use of conceptual analysis in epistemology. And, in fact, it seems that to the extent that we want to ensure that there is continuity between folk concepts and those deployed in epistemological theories what we really want is the useful work done by the folk concepts to also be done by deploying the concepts employed in epistemological theories. We only want to ensure that our theories fit with folk concepts to the extent that those concepts have utility. And what we really want is to ensure that our classifications are useful rather than that they are identical to folk classifications. By taking a teleological approach to epistemology, and considering the goals that can be

[40] Frank Jackson (1998) makes a similar point about conceptual analysis being required to define the subject matter of epistemology.

achieved through our epistemic classifications, epistemologists attempt to ensure that epistemic classifications are useful even where they are not identical with the classifications made by the folk, which might not really be all that useful. They can strive towards the right degree and kind of continuity with folk epistemic concepts—not complete continuity but continuity of purpose.

3. Should we adopt Weinberg's reconstructive neo-pragmatism?

We can see, then, that there are various advantages to adopting a teleological approach to epistemology and in section 1 I argued that Weinberg's approach is preferable to Craig's. Does this mean that epistemologists should adopt Weinberg's reconstructive neo-pragmatism? In my view, epistemologists should not move wholly into Weinberg's camp because he gives too small a role to results from psychology.

As mentioned in section 1, Weinberg accepts that results from psychology can be relevant to epistemology and should be taken into consideration when developing accounts in epistemology. However he suggests that these results should only have a limited influence on the theories developed in epistemology. Here is an example of Weinberg summarising the role of results from psychology in the methodology that he defends:

> Any proposed strict constraint must satisfy two criteria: (i) it must not place so onerous a burden on our cognition that too little of our epistemic lives can be sustained, and (ii) whatever beliefs and inferences do pass through must be appropriately DR-promoting. If a proposed constraint fails either condition, then it is inconsistent with our DR interests and should thus be weakened (Weinberg, 2006: p. 39).

This quote requires some unpacking. Weinberg is discussing the constraints that he thinks should be placed on any account of knowledge, rationality and justified belief. In

(i) he suggests that accounts should not place the requirement for good cognition too high. In (ii) he claims that those beliefs that are judged positively given some criteria of what is required for positive epistemic appraisal should be 'DR-promoting'; by this he means that they should promote diachronic reliability and dialectical robustness. This means that an account of knowledge, rationality or justified belief must promote (a) the ability to track the truth by updating one's beliefs in line with the facts across time (diachronic reliability) and (b) the ability to communicate with others the reasons for our beliefs so that we can persuade each other to share our true beliefs (dialectical robustness). So Weinberg places a number of constraints on successful accounts of knowledge, rationality or justified belief: they should be not be too onerous, they should promote the ability to track the truth across time and they should promote the ability to communicate with others the reasons for why we have certain beliefs. But on close inspection of what Weinberg says, it becomes apparent that he actually defends the first criterion only as a way to successfully meet the second, 'If a proposed constraint fails *either condition*, then it is inconsistent *with our DR interests*' (emphasis added). For Weinberg, the problem with failing to take results from psychology into consideration is that by doing so one risks producing an account which is too demanding for human cognition and therefore cannot identify and promote states that are diachronically reliable and dialectically robust.

What this shows is that Weinberg thinks we can understand the goals that can be achieved via our epistemic classifications independently of consulting results from the scientific study of human psychology. At the very least, Weinberg does not stress the importance of using results from psychology to understand the goals that can usefully be achieved by our epistemic classifications. Notice that he does not provide empirical

support for the claim that diachronic reliability and dialectical robustness are valuable. Results from the empirical study of human psychology are used only to ensure that theories which meet the goals of DR-promotion do not place too onerous constraints on human cognition, so that humans can never meet these constraints. According to the quote from Weinberg, results from psychology should not play an operative role in deciding which epistemic norms or standards should be used to assess beliefs or belief revision, they should merely be used as a filter on previously established ideas about how these things should be assessed. My suggestion is that epistemologists should use results from psychology more readily than this: they should recognise that only by considering results from psychology is it possible to understand the goals we can achieve through our epistemic classifications. This is because by considering such results it is possible to understand new goals that people can achieve by applying their epistemic concepts. Results from psychology should not merely filter previously settled ideas about the goals of our epistemic classifications.

4. How to take a teleological approach

To see how it can be useful to take results from psychology into consideration when developing an account in epistemology, let us return to some of the examples of results from psychology which have been discussed so far in this thesis and see how these results from psychology reveal goals that can be achieved through our epistemic classifications.

Take, for instance, the scientific study of human reasoning. Recall that, as I have discussed at numerous places in this thesis, psychologists working in this area have

discovered that people often use heuristic methods of reasoning, that is, they develop beliefs based on only a small amount of information. For example, people asked to choose between two options sometimes use the recognition heuristic; they choose the option that they recognise. Students asked which of two foreign cities is larger selected the city that they recognised (Goldstein and Gigerenzer, 2002). For some time, psychologists from the 'heuristics and biases' movement thought that results showing that people use heuristics established that humans are highly irrational (e.g. Kahneman and Tversky, 1982). More recently, however, some psychologists have argued that it can be very useful to utilise heuristic reasoning strategies because they can be highly successful and make good use of the limited resources of information we have available to us, so that these psychologists call the reasoning 'fast and frugal' (e.g. Gigerenzer and Todd, 1999). For instance, the students using the recognition heuristic to work out which city is larger came to the correct conclusion while considering only a small amount of the relevant information. Given that standard views of rationality, according to which people can be rational only by checking the consistency of any target belief with all their other beliefs, seem to be beyond the reach of ordinary thinkers due to limits on human cognition, as discussed in section 2.2, the view that the use of heuristics is rational has gained strength (e.g. Samuels, Stich and Bishop, 2002). Fast and frugal reasoning strategies can lead to true beliefs and make effective use of the available resources, so it seems that people can be rational by using them.

One implication of the view that people can be rational using fast and frugal methods of reasoning is that it highlights a new goal that we can achieve by applying our epistemic concepts: to identify the successful use of fast and frugal reasoning processes. It will be useful to be able to identify whether another person is using a fast and frugal reasoning

strategy in such a manner that they are likely to end up with true beliefs. This will enable us to work out whether that person is a reliable source of information[41], so we know whether to share the beliefs that they express or not to share those beliefs. Here, then, we find that results from psychology can supply information about the goals that can be achieved by our epistemic classifications because they indicate that we can usefully identify true beliefs by identifying successful fast and frugal processes.

Another set of psychological literature that shows how psychology can reveal goals that can be achieved by our epistemic classifications is the literature on implicit biases. In chapter one we found that there is now a large body of literature in social psychology showing that people are biased for members of some groups and against members of other groups. These biases lead people to make positive associations with members of some groups and negative associations with members of other groups, without those associations tracking the truth about features of the individual being assessed other than their group membership[42]. According to these results, people favour white people to black people, young people to old people, white children to black children, non-Arab Muslims to Arab Muslims, and the list goes on (Nosek et al, 2007). In light of the psychological literature on implicit biases, it seems as if people should strive to deliver epistemological theories which exclude the products of implicit biases from being classified positively. An account that suggests that people can know something about another person based on an implicit bias against all members of their group seems to be

[41] Craig (1990) distinguishes between identifying reliable informants and reliable information sources. For the sake of simplicity I will ignore this distinction in this discussion. In my view, an important goal that can be achieved by applying our epistemic concepts is to identify reliable sources of information as well as reliable informants. This means that as long as we can identify a reliable source of information by identifying people using heuristics successfully, this is a goal that can usefully be achieved by our epistemic classifications.

[42] I mean to be neutral in my discussions of truth tracking between the various theories associating epistemic notions with truth tracking. But I take it that a process that does not track the truth should not be viewed positively from an epistemic perspective.

unsatisfactory. A general evaluation of a person that focuses on a single feature rather than a wider assessment of that person's characteristics should not be viewed as epistemically valuable because it will fail to track the truth about many important features of the individual. This means that epistemologists should aim to develop a theory according to which the products of implicit biases are classified negatively, so that people using positive epistemic concepts (knowledge, rationality, justified belief) can achieve the goal of tagging beliefs which are not based on implicit biases[43].

In the examples of fast and fugal reasoning and implicit biases we find that results from psychology can identify goals that can be achieved through our epistemic classifications. These goals simply would not have been recognised as important from the armchair. Once one recognises the nature of implicit biases it almost seems self-evident that it can be valuable to develop an epistemological theory that excludes their products from being classified as knowledge, rationality or justified belief. But psychologists only discovered the nature of implicit biases by recording people's split-second responses to members of different groups—something that would not have been noticed by armchair epistemologists.

More generally, it seems as if epistemologists should not depend upon people's ability to imagine the goals that can be achieved through our epistemic classifications. My argument suggests that what goals can be achieved through our epistemic classifications

[43] Gendler (2008, 2008a) argues that the states discussed in the social psychology literature should be distinguished from beliefs because they are less under our control and less likely to track the truth. She calls them aliefs. What this argument suggests is that if implicit biases produce aliefs, it can be useful for epistemologists to consider results from psychology so that they can recognise that one of the functions of epistemological concepts is to distinguish between beliefs and aliefs, the latter of which fail to track the truth and therefore lack epistemic justification. In cases such as the one described, where someone's actions seem to be determined by weapon bias, it seems we should deny that they are epistemically justified either because they are based on unfounded beliefs or because they are based on aliefs.

will be determined by the nature of the human beings who are evaluated. Because people are able to make good use of fast and frugal methods of reasoning and are prone to implicit biases, these facts should be reflected in our understanding of the goals which can be achieved by our epistemic classifications. If human psychology were different then different goals could be achieved. Armchair philosophers cannot expect to have the best possible understanding of the nature of human psychology. They should recognise that results from psychology can reveal facts about human psychology that they could not have even imagined. This means that philosophers should recognise their own limits and consult results from psychology while they attempt to get the understanding of human psychology required to grasp the goals that can be achieved through our epistemic classifications. They should embrace the idea that results from psychology can be vital to epistemology because they can *reliably* reveal goals to be achieved through epistemic classifications, unlike the imaginations of armchair epistemologists.

These arguments should bring to mind the general arguments for methodological naturalism presented in chapter one. Similar arguments can be provided for both radically naturalistic neo-pragmatism and methodological naturalism in general because the former is an instantiation of the latter. Radically naturalistic naturalism contributes a vitally important proper understanding of human psychology to epistemology, so it is fully consistent with the aims of methodological naturalism. The results from psychology which were shown to be relevant to epistemology in chapter one are now viewed as relevant to epistemology because they reveal goals that can be achieved through our epistemic classifications. But, more generally, the results from psychology should still be viewed as important to epistemology for the same reasons as they were

found to be important in chapter one: because they contribute to an accurate account of knowledge, rationality and justified belief which truly reflects the way that people can revise their beliefs. It is as a consequence of truly reflecting the way that people can revise their beliefs that radically naturalistic neo-pragmatism enables epistemologists to develop accounts which are consistent with people using their epistemic concepts to achieve important goals when evaluating other people's belief revision. With an accurate picture of human belief formation it will be possible to get a handle on what goals can be achieved by classifying some belief forming processes positively and some negatively, and then to achieve these goals.

5. Radically Naturalistic Neo-pragmatism in Action

To better understand the consequences of radically naturalistic neo-pragmatism, let us see how adopting such a methodology can affect the selection and endorsement of an epistemological theory. Take the case of implicit biases. Whether or not one thinks that an important goal to be achieved through our epistemic classifications is to tag only those states which are not the products of implicit biases as knowledge, rationality or justified belief will decide whether or not one finds some accounts in epistemology convincing. For example, as we found in chapter one, some internalists defend the view that people can be justified in their beliefs as long as they can provide justifying reasons, even if those reasons did not form the basis of their beliefs. Conee and Feldman (2004: p. 93) describe the view according to which a belief can be justified if some reasons for accepting it can be given, even if the belief is not properly arrived at and therefore not what they call 'well-founded'. Psychological findings suggesting that people often have beliefs that are not properly arrived at, because they are based on implicit biases, seem to show that people should be reluctant to accept that beliefs can

be justified just because some reasons can be marshalled in support of them. This is because it is easy to imagine someone marshalling reasons in support of a highly prejudiced view that is the product of an implicit bias. For example, a person could take the position that Barack Obama's health care reforms are poor based on an implicit bias against black people. She might then be able to marshal some justifying reasons for this position, such as that the British National Health Service has problems and the US should not move closer to that system. The internalist view of epistemic justification we are discussing would imply that this person could be justified in this belief based on an implicit bias, which will prevent the person from tracking the truth about the real reasons why health care reform might or might not be a good thing, just because she can marshal some reasons to justify the belief. Those who are convinced that the products of implicit biases should not be classified as knowledge, rationality or justified belief will reject this internalist theory of epistemic justification because it fails to ensure that the products of implicit biases, such as this one, are judged to be unjustified.

Those who do recognise that there are important goals that can be achieved through our epistemic classifications that can only be discovered by considering results from psychology will therefore reject the theory of epistemic justification which has the unsatisfactory consequence of suggesting that the products of implicit biases will be justified. They will recognise the important goal of identifying and tagging the products of implicit biases negatively and will only accept an epistemological theory according to which this goal can be achieved by applying our epistemic concepts. As a consequence of this, they will not accept the version of internalism we are discussing.

This desirable consequence is not guaranteed to anyone who adopts a different approach to epistemology, including Weinberg. In fact, it seems likely that the internalist theory

of justification being discussed will be viewed kindly by those who adopt Weinberg's methodology. Recall that Weinberg values dialectical robustness. He thinks that many important goals can be achieved if people can provide reasons for their beliefs and discuss these reasons with other people. As a result, he finds many aspects of internalism attractive[44]. It is likely that he would also find the version of internalism being discussed attractive because if a person can provide some reasons justifying her beliefs then she can discuss these reasons with other people and try to persuade others to share her beliefs. As long as people are able to provide some reasons to others for why they should share the same position, the dialectical robustness constraint that Weinberg endorses will be met. It is also not usually beyond the capacities of ordinary individuals to provide some reasons in support of their beliefs, so those who take results from psychology into consideration to see whether people can meet the dialectical robustness constraint by giving some reasons for their beliefs are likely to conclude that they can—they will not reject the internalist view of justification associating it with giving some reasons on the basis that results from psychology suggest that people cannot give such reasons. This means that those who use Weinberg's methodology and value the promotion of dialectical robustness are highly likely to accept an internalist theory of justification that suggests that the products of implicit biases can be justified. For those who want to avoid this undesirable consequence, it is important to see that some goals that can be achieved through our epistemic classifications can only be understood by

[44] Weinberg (2006) explicitly rejects the strict internalist constraint according to which people have justified beliefs only if they have access to the reasons for their beliefs because this constraint would prevent many of our ordinary inductive beliefs from being justified. This is because we do not have conscious access to the reasons why we have many ordinary inductive beliefs. This does not mean, however, that Weinberg would reject all intermalist constraints on justification. On the contrary, it seems as if an approach on which a person can have justified beliefs if they can present justifying reasons is likely to be attractive to those who see internalism as valuable but doubt whether people can really have access to the reasons driving their belief formation.

considering results from psychology. In other words, it is important to accept radically naturalistic neo-pragmatism.

In sum, then, whether or not a person takes a radically naturalistic neo-pragmatist approach to epistemology will be decisive to their acceptance and endorsement or rejection of specific epistemological theories. And it seems as if people who adopt radically naturalistic neo-pragmatism are more likely to accept a correct epistemological theory because they will fully understand some important reasons for endorsing or rejecting certain theories which others will not understand—i.e. the goals that can usefully be achieved by people adopting the concepts described by those theories.

6. Conflicting Goals: A Potential Problem?

I have argued so far that epistemologists should consider the goals that can be achieved through our epistemic classifications and should consider results from psychology because these can reveal some important goals. This raises a question about what should happen if there is a conflict of goals that can be achieved by our epistemic classifications. How can we decide which of the conflicting goals should be used as a constraint on any epistemological theory? If epistemologists do discover conflicting goals that can be achieved through our epistemic classifications, does this cause a problem for radically naturalistic neo-pragmatism?

There clearly is an issue surrounding conflicting goals. Even if we just consider the cases discussed in section 4 of fast and frugal reasoning strategies and implicit biases it is clear that there are some potential conflicts. When I discussed fast and frugal reasoning processes I argued that one goal that could be achieved by our epistemic classifications is to identify the successful use of these processes that are fast,

automatic, unconscious and frugal processes (i.e. processes that process only a small amount of information), tagging the successful use of the processes with positive epistemic classifications. The problem arises when we consider how implicit biases are similar to fast and frugal reasoning processes. Implicit biases can also be described as fast, automatic, unconscious and frugal processes (they only focus on a small amount of information about an individual being assessed, i.e. the group membership of that individual). However, we do not want to classify implicit biases positively. So it seems that there is a conflict of interest: we are required to heap both positive and negative epistemic appraisal upon fast, automatic, unconscious and frugal processes. Does this leave radically naturalistic neo-pragmatism in trouble?

There are a number of responses that can be made in defence of radically naturalistic neo-pragmatism, to show that it is not really in trouble. Firstly, the example just given actually shows how important it is for epistemologists to consider results from psychology when aiming to understand the goals that can be achieved by our epistemic classifications, rather than showing that we should not consider such results. What the similarities between the successful and unsuccessful processes show is that epistemologists must consider results from psychology *in some detail* so that they can discover the precise conditions under which fast, automatic, unconscious and frugal processes are reliable. Once epistemologists recognise that it can be valuable to identify fast and frugal heuristic reasoning processes and implicit biases they should not be put off by the difficulties that there might be in differentiating between the processes. Instead they should see how important it can be to get their hands dirty, to understand the perhaps subtle differences between the conditions under which fast, automatic, unconscious and frugal processes are reliable and unreliable. The rewards for engaging

in this project will be to understand fully how exactly we can achieve the valuable goals identified, such as giving successful fast and frugal heuristic reasoning positive epistemic status and implicit biases negative epistemic status. This should be enough of a reward to motivate epistemologists to consider results from psychology in some detail.

Secondly, it is important that the phenomenon of conflicting goals is not unique to radically naturalistic neo-pragmatism; it is common to all methodologies in epistemology that involve epistemologists considering the goals that they can achieve through their epistemic classifications. And, in fact, the phenomenon of conflict is common to all philosophy. People are always going to disagree about what counts as a valuable epistemological theory and about what counts as making the right sort of distinctions. Philosophers do not despair and give up on philosophy because they find that there are conflicts that need to be resolved. On the contrary, they engage in debates with the aim of resolving conflicts—that is simply what philosophers do. What is unique to teleological approaches is just that the conflicts are likely to surround the question of which goals it is most valuable for people to achieve through our epistemic classifications. But in many ways this is a desirable sort of conflict because people can easily engage in debates about what goals it is valuable for people to achieve. Things are far more difficult for those epistemologists who are driven by intuition because if there is a clash of intuition it is likely that there will be a deadlock with people having brute responses which are incompatible, which they might not even be able to explain the source of. What is distinctive about teleological approaches to epistemology is not that they bring conflicts with them, then, but just the type of conflict they are likely to bring. And, in my view, the type of conflict that is likely to be found by those taking a teleological approach to epistemology is not a troubling one because it is one that can

resolved by debating what goals we really should be able to achieve through our epistemic classifications.

It might ultimately be necessary to abandon some goals we identify as valuable because they come into conflict with others. For example, one might decide that it is so important to exclude implicit biases from our positive epistemic classifications that we should exclude all fast, automatic, unconscious and frugal processes, forsaking the goal of tagging the products of heuristic strategies positively. But rather than highlighting a problem for radically naturalistic neo-pragmatism, the possibility of these goals conflicting actually shows how valuable it can be to discuss the relative importance of various goals that can be achieved through our epistemic classifications, some of which can only be revealed by considering results from psychology. After all, we would not have been aware that the goal of identifying implicit biases is valuable at all without considering results from psychology.

One possibility that it is worth considering is that there are such subtle differences between the successful and unsuccessful processes that no existing epistemological theory will distinguish between them, although there is some principled distinction to be found. For example, it might be that the best existing account of rationality associates being rational either with using fast, automatic, unconscious and frugal processes, because these are efficient and often effective, or with not using processes matching this description and instead using processes that are slower, controlled, conscious and involve considering a large amount of information. If the best existing account of rationality took either of these forms, then a consequence of endorsing either account would be accepting that heuristics and biases and implicit biases should both be assessed positively or neither should. This account of rationality would not pick up on

the subtle distinction between those fast, automatic, unconscious and frugal processes that lead to true beliefs and those that lead people to adopt prejudiced views of others. In this kind of situation, if no existing theory of rationality is able to do the job of making this subtle distinction, and yet there is a way to make this distinction in a principled way, then in my view existing theories should be overhauled. Epistemologists should strive to develop epistemological theories which are able to make the subtle distinctions we require so that we can achieve the goals we find valuable. They should be highly influenced by results from psychology, so much so that they will reject all existing accounts of knowledge, rationality or justified belief if they suggest that we cannot achieve through our epistemic classifications the valuable goals revealed by results from psychology that we can realistically achieve.

7. A Comparison

So far I have argued the following. Epistemologists should adopt a methodology that draws on teleological approaches described by Craig and Weinberg. They should depart from Craig and Weinberg's methodologies because they should consult results from the sciences to reveal goals that can be achieved and human needs that can be met through our epistemic classifications. Epistemologists should, in my view, only accept an account in epistemology if it suggests that the concepts of knowledge, rationality and justified belief can be used to achieve the valuable goals that can be identified by considering results from psychology. They should be willing to reject all existing accounts if they are inconsistent with people being able to achieve their goals through their use of epistemic concepts. Many of these claims are similar to those made by Michael Bishop and J.D. Trout in defence of their position, strategic reliabilism (e.g.

Bishop and Trout, 2005, 2008). For a final piece of exposition of my position, it will therefore be fruitful to compare it to strategic reliabilism.

According to strategic reliabilism, epistemologists should evaluate cognitive processes, but they should not evaluate them as epistemologists within standard analytic epistemology do. They should not evaluate the processes to see whether they fit with our concepts of and intuitions about knowledge, rationality and justified belief. Instead, epistemologists should evaluate cognitive processes for their epistemic excellence. For strategic reliabilists, a cognitive process is epistemically excellent if it is conducive to the truth and, therefore, reliable. It should be valued if it is robustly reliable, that is, reliable across a wide range of environments. Finally, an epistemically excellent process will make good use of the resources available and it will be valuable when used to deal with significant matters.

Strategic reliabilists describe a number of reasoning strategies that they think meet the criteria of being robustly reliable, making good use of resources and dealing with significant matters. They discuss statistical prediction rules[45]. Statistical prediction rules are mechanical models that are used to predict outcomes. Weights are given to various factors that might affect the outcomes being assessed and models are developed that predict the likelihood of any outcome given the presence of relevant factors. The weightings given to each factor are determined empirically, by assessing the past correlation between the factor and the outcome being assessed. One example that Bishop and Trout discuss is of the use of statistical prediction rules in predicting the success of electroshock treatment (2005: p. 13). It has been found that statistical

[45] The successful use of actuarial prediction rules which was used in chapter one as an example of results from psychology that pose a challenge to armchair epistemological theories (Bishop, 2000) counts as a statistical prediction rule and would be advocated by strategic reliabilists.

prediction rules can be highly successful at predicting the success of electroshock treatment if those applying the rules have access to information about three factors: the marital status of the patient, the length of time of the psychotic distress and the patient's insight into his or her own condition (Wittman, 1941). Statistical prediction rules are able to outperform hospital and psychological staff members when these cues are available.

Bishop and Trout claim that the success of statistical prediction rules applied to the question of electroshock treatment is far from an anomaly and that in many cases the use of statistical prediction rules is more reliable than the use of expert testimony. They therefore suggest that it can be epistemically excellent to apply statistical prediction rules because they are robustly reliable strategies that make good use of a limited resource of only a small amount of information (marital status, length of psychotic distress and patient's insight into condition) and deal with a significant matter (whether electroshock treatment will be effective). Bishop and Trout suggest that the traditional project of epistemology should be replaced by the search for an understanding of the areas of cognitive science (e.g. psychology, AI, etc.), which they call ameliorative psychology, that outline strategies such as statistical prediction rules that can be adopted by people to improve their reasoning.

There are various similarities between strategic reliabilism and radically naturalistic neo-pragmatism. Both views suggest that results from psychology should contribute to an understanding of what belief forming processes are epistemically valuable. Both projects reject conceptual analysis and the use of intuitions as brute evidence in epistemology. Both also suggest that epistemologists should be willing to reject all existing theories in epistemology. I suggested at the end of section 6 that if none of the

existing theories enable us to achieve our epistemic goals then they should all be replaced and strategic reliabilists reject the whole project of standard analytic epistemology and the existing theories produced therein. However, there is one important respect in which my position differs from that of Bishop and Trout.

The important respect in which strategic reliabilism differs from radically naturalistic neo-pragmatism is that it focuses on the reasoner, describing advice that can be taken by individuals to improve their own reasoning,

> Epistemology is but a hollow intellectual exercise if it does not ultimately provide a framework that yields useful reasoning advice. Strategic Reliabilism provides a framework for figuring out how one ought to reason about particular problems (Bishop and Trout, 2005: p. 138).

Strategic reliabilists do not only reject the idea that conceptual analysis is important, they switch the focus away from the usual enterprise of understanding how we should classify each other's beliefs, suggesting that epistemologists should aim to understand how people can best improve their own reasoning. I do not think the only role for epistemology is to provide advice for reasoners about how they can ameliorate their reasoning. This is because epistemologists are interested in concepts that can usefully be applied to other people to assess their reasoning abilities and beliefs. Epistemic concepts are required to answer to various human needs and many of these needs involve evaluating other people. In my view, then, strategic reliabilism's shortcoming is that it focuses on the reasoner and how she can improve her reasoning, failing to describe how epistemology can also supply an account of knowledge, rationality and justified belief that describes how these concepts can be used by the evaluator.

In contrast to strategic reliabilists, those who defend a teleological approach to epistemology, focus on the wider use of epistemic concepts, as instruments that can be

used to evaluate and tag other people and their beliefs/belief forming processes. As a consequence of recognising that epistemic concepts are used by people to evaluate others, they consider the situation of the evaluator. Take, for example, Craig's discussion of the causal theory of knowledge, described in section 1. Craig argues that the causal theory of knowledge cannot describe what knowledge really is because, if knowledge were belief caused by the fact known, then people could not use the concept of knowledge to identify reliable informants. Only if evaluators could use the concept of knowledge as belief caused by the fact known to identify other people as reliable/unreliable informants would the causal theory of knowledge be correct. Here Craig is considering the standpoint of the evaluator and a theory is only acceptable if an evaluator can make good use of the concepts it describes. This is a valuable aspect of a teleological approach to epistemology because our epistemic concepts are used by us to make many important distinctions, so we should be sure that our accounts of knowledge, rationality and justified belief describe them in such a way that they can enable us to make these distinctions.

To understand the importance of taking the standpoint of the evaluator into consideration when developing an epistemological account, consider the following example. Some psychologists are working to identify strategies for people to use to overcome implicit biases. For example, it seems that people can overcome a bias against black people by considering positive role models who are black, such as the actor Denzel Washington, before making evaluations (Kawakami et al., 2000). It seems as if strategic reliabilists will judge these strategies to be epistemically excellent. They are likely to produce true beliefs, because they will produce beliefs free from prejudice. The cognitive processes will be used to deal with a significant problem for the individual,

assuming that the individual using the strategy thinks that overcoming biases is significant. And it seems that at least some of the strategies, such as the one described by Kawakami et al., are relatively simple to use, so they will make good use of cognitive resources. It therefore seems as if at least in some cases strategies to overcome implicit biases should be classified by strategic reliabilists as epistemically excellent.

Things get less clear, however, when considering what the radically naturalistic neo-pragmatist should say about the use of reasoning strategies designed to overcome implicit biases. It seems right to say that people who choose to use an ameliorative strategy to overcome implicit biases whenever possible are making good decisions. But should evaluators really be expected to be able to identify when people are using such a strategy? It seems not. It is unrealistic to expect that one person should be able to work out whether another person has made himself informed about the results from psychology about implicit bias and is using a strategy to overcome the bias[46]. In fact, once one considers the situation of the evaluator it seems that one should not always expect people to use their epistemic concepts to identify when other people are using the reasoning strategies that ameliorative psychology would deem to be epistemically excellent. This means that those beliefs produced by reasoning strategies which ameliorative psychology would classify positively would be more appropriately classified negatively on the radically naturalistic neo-pragmatist account. A realistic goal for epistemic concepts to be used to achieve in the implicit bias case would be to identify situations in which people are likely to be biased and to classify the products of these situations negatively. If epistemic concepts are used in this way then they will not

[46] In some situations an evaluator might be able to directly ask whether another person is using an ameliorative strategy, but in so many others it seems as if the evaluator would not be able to ask. Imagine asking a stranger who is sitting next to you at the football whether she is using an ameliorative strategy to overcome implicit bias when she tells you that a certain player is particularly prone to diving. This is unlikely to go well.

reflect how people can overcome implicit biases by using ameliorative strategies described by psychologists, but they will err on the side of caution and be more likely to ensure that all beliefs based on implicit biases are classified negatively. And given that evaluators cannot be expected to identify when other people are using ameliorative strategies, this seems to be the best outcome we can hope for. Taking a radically naturalistic neo-pragmatic approach to epistemology will therefore produce a different and, in my view, better outcome to taking a strategic reliabilist approach.

In sum, what is distinctive about radically naturalistic neo-pragmatism is that it both considers results from psychology to establish the goals that can be achieved through our epistemic classifications *and* focuses on the situation of the evaluator and the uses that she can make of her epistemic concepts. It provides the opportunity for epistemologists to develop accounts of knowledge, rationality and justified belief that can be used by evaluators to make useful evaluations of the beliefs and belief forming processes of other people. Strategic reliabilism does not do this, and neither do alternative theories about the methods to use in epistemology, so radically naturalistic neo-pragmatism has advantages over all the available alternatives.

Conclusion

Epistemic concepts are required to answer to human needs. By taking a teleological approach to the methods to be used in epistemology, epistemologists are able to tackle this fact head on. They take into consideration the goals that can be achieved through our epistemic classifications, work out which ones are valuable and only accept accounts which describe our epistemic concepts in such a way that they can be used to achieve the important goals.

At the same time, epistemologists adopting a teleological approach to epistemology can avoid various problems with the other methodologies that have been implicitly or explicitly endorsed by epistemologists, such as the problems that come with conceptual analysis and the use of intuitions as evidence in epistemology. Because they do not aim for an understanding of the necessary and sufficient conditions required for the appropriate application of our epistemic concepts, they are not caught in the stagnant process of studying our potentially flawed and not very useful concepts. They cannot be criticised for assuming that concepts take the form of necessary and sufficient conditions; something that has been challenged by results from psychology. Those who adopt a teleological approach to epistemology are able to avoid these criticisms of standard approaches to epistemology while ensuring that there is still a degree of continuity between folk concepts and the conceptions of knowledge, rationality and justified belief and those they describe in their theories. This is because they can ensure that the concepts of knowledge, rationality and justified beliefs as they describe them are useful in the same ways that folk concepts are. As long as they are as useful as folk concepts, it seems that the concepts of knowledge, rationality and justified belief described by those who adopt a teleological approach to epistemology are sufficiently continuous with folk concepts.

We have thus seen that there are various gains that can be made by an epistemologist who adopts a teleological approach to epistemology. Not any teleological approach will do, however. What the examples discussed in this chapter show is that some goals that we can achieve through our epistemic classifications will only be discovered by epistemologists who consider results from psychology. This means that epistemologists should not adopt the teleological approach to epistemology described by Craig because

he is not a methodological naturalist, and therefore does not emphasise the value of using results from psychology while developing an epistemological theory. Epistemologists should also not adopt the teleological approach to epistemology described by Weinberg because, although he accepts that there can be a role for results from psychology in epistemology, he only stresses the importance of using these results to ensure that any account in epistemology describes knowledge, rationality and justified belief in such a way that they are within the reach of ordinary thinkers. He does not stress the importance of using results from psychology to reveal goals to achieve through our epistemic classifications. In contrast, according to radically naturalistic neo-pragmatism, epistemologists should aim to discover all the goals that can usefully be achieved through our epistemic classifications. They should take results from psychology into consideration to develop an idea of the goals to be achieved by epistemic classifications before assessing every epistemological theory to see whether the concept of knowledge (or rationality or justified belief) it describes would enable us to achieve important goals if the theory were correct.

An epistemologist who uses radically naturalistic neo-pragmatism to develop an account of knowledge, rationality or justified belief would be using a methodology consistent with the principles of methodological naturalism because she would take results from psychology into consideration when developing her account. As a result, she would obtain all the advantages of methodological naturalism described in chapter one. Radically naturalistic neo-pragmatism therefore seems to be the viable account of the methods to use in epistemology consistent with methodological naturalism that this thesis has been searching for.

Now we have a clear picture of how epistemologists should do epistemology. They should be methodological naturalists who take a teleological approach to the methods used in epistemology. By adopting a teleological approach, it is possible for methodological naturalists to avoid the various problems with other accounts of the methods to be used in epistemology. And by adopting methodological naturalism, those who would otherwise adopt a teleological approach to epistemology will develop an accurate picture of the goals that can usefully be achieved via our epistemic classifications. Now we are ready to reach the conclusion of this thesis: methodological naturalism is valuable and it can be best implemented by adopting a teleological approach to epistemology, the approach I have called 'radically naturalistic neo-pragmatism'. There is no guarantee that those people who use this methodology to approach the questions of epistemology will converge on a single view of knowledge, rationality or justified belief. Nevertheless, they will be able to engage in a fruitful debate about what goals can be achieved through epistemic classifications while striving to ensure that the outcome of their theorising will be an account that is able to capture the goals that can usefully be achieved by applying our epistemic concepts.

Conclusion

In this thesis I have argued that methodological naturalism is a valuable enterprise. By taking results from the sciences into consideration when developing their epistemological positions, theorists are able to reflect how it is that people, with real cognitive capacities and limitations, are able to have knowledge, be rational and have justified belief. Epistemologists are able to focus on relevant possible facts about human psychology and accommodate these to develop theories that really reflect human belief formation and revision.

This thesis has provided a survey of some of the prominent accounts of the methods to use in epistemology that are consistent with methodological naturalism, in a search for the most promising approach to use. It has evaluated integrative approaches to epistemology, which defend the use of the traditional methods of epistemology. They suggest that conceptual analysis or reflective equilibrium is valuable as long as adjustments are made to their outputs to reflect the cognitive limitations of human beings. In discussion of these positions we have found that they are problematic because they are unable to overcome the difficulties inherent in targeting concepts rather than psychological phenomena and the difficulties involved with using intuitions as a source of epistemological theories.

Certain suggestions have been put toward by those who defend integrative approaches that are supposed to show that it is worthwhile retaining traditional methods within the framework of methodological naturalism. For example, Alvin Goldman argued that the only way that we can ensure that our epistemological theories are truly epistemological is to analyse epistemological concepts. We found, however, that Goldman's position is

inconsistent because he both recognises the need to revise epistemic concepts and thinks it is valuable to engage in conceptual analysis. We found that the benefits gained from engaging in conceptual analysis can be obtained by studying but not analysing concepts, trying to work out what we want to achieve through our epistemic classifications rather than aiming to uncover necessary and sufficient conditions for the application of our concepts. As a result, I have argued that conceptual analysis should not be retained as even a small part of the methodology of epistemology. Stein provided another defence of an integrative approach to epistemology when he suggested that the method of reflective equilibrium will ensure that epistemological theories won't be influenced by untrustworthy intuitions, but only trustworthy ones. We found that it was unlikely and certainly not guaranteed that the method of reflective equilibrium, even with the influence of philosophical and scientific theories, would overcome the influence of all untrustworthy intuitions. I concluded that the attempts to defend integrative approaches to epistemology have failed and we should consider alternative methodologies consistent with methodological naturalism.

One prominent alternative to integrative approaches is Hilary Kornblith's account of how epistemologists should endeavour to understand knowledge. Kornblith suggests that knowledge is a natural kind that can be identified by (a) using our intuitions about obvious cases as evidence about the natural kind and (b) identifying the reliably produced true beliefs implicated in cognitive ethology. He suggested that concepts should not be the target of epistemology and that epistemologists should study the natural phenomena of belief revision rather than our concepts thereof. In chapter four I challenged the idea that knowledge is a natural kind identifiable in the way that Kornblith describes. I argued that intuitions are unlikely to provide an adequate source

of information about what is required for knowledge because even if they can identify obvious cases of knowledge they cannot be relied upon to identify other cases. They will therefore not provide guidance about what belongs in Kornblith's putative natural kind. I also showed that there are problems that come with borrowing the categories of science, as it is implied that epistemologist should do by the claim that knowledge is a natural kind of true belief produced by cognitive equipment performing its proper function and identified by scientists. I argued that we should only accept that scientific classifications are important to epistemology if they enable us to make certain important kinds of valuable classification.

I suggested that epistemologists should not be guided purely by untutored intuition nor should they borrow distinctions from the sciences and instead they ought to recognise that there are certain important goals we aim to achieve by using our epistemological distinctions. Epistemologists, according to this view, should engage in the project of understanding the goals that we aim to achieve through our epistemic classification. They should only accept an account of knowledge if it suggests that we can use the concept of knowledge to achieve certain ends (something similar can be said about our accounts of rationality and justified belief). Our judgements about the goals we have in our epistemic classifications should be used to decide between competing intuitions and to decide which scientific distinctions to borrow from because any theory that fails to answer to human needs (such as between various beliefs about modern subjects or institutions) will be an unsatisfactory epistemological theory. Because our human goals are so important to the process of identifying what is knowledge, knowledge is not a natural kind.

The idea that epistemologists need to investigate the goals people can achieve through their epistemic classifications, which was outlined in response to Kornblith's position, was developed more fully in the final chapter and exposition of what I think is the most promising approach that epistemologists can take to the methods of their discipline. I argued that methodological naturalists can acquire the benefits of studying concepts to ensure the continuity of epistemology with folk concepts (as demanded by Goldman) by taking a teleological approach to epistemological concepts, studying the goals that can usefully be achieved by applying those concepts. By taking this approach, epistemologists are able to ensure that the concepts described in their theories of knowledge, rationality and justified belief come as close as possible to performing the valuable goals that ordinary epistemological concepts perform and more.

I considered two versions of teleological approaches to epistemology, those presented by Edward Craig and Jonathan Weinberg. I highlighted the many advantages of using these methods, including the various ways that they can avoid the criticisms levelled at the methods traditionally used in epistemology. However, I also noted one significant shortcoming of these positions: they failed to give a sufficiently large role to results from the scientific study of human psychology in understanding the goals that can be achieved through our epistemic classifications.

Throughout the thesis I have stressed that an epistemological theory that fails to reflect what science says about how people can successfully revise their beliefs will be poorer as a result. Scientific theories can help epistemologists to identify belief forming processes that seem at least *prima facie* to be appropriate candidates for being classified as knowledge. The example that I have used frequently as a case study is the psychological study of human reasoning. This case illustrates how psychological results

can challenge our assumptions about how we ought to revise our beliefs if we want to get desirable results. It suggests that it can be valuable to revise our beliefs based on little information, using shorthand methods or fast and frugal heuristics, rather than following the standard picture of rationality according to which rationality is associated with checking the consistency of beliefs by applying logic or probability theory. Not only does it challenge the standard conception of rationality, the case also challenges the idea that our epistemological concept of rationality ought to function in such a way that it only identifies beliefs that have been checked for consistency as rational.

Here is where I think it is important to make adjustments to previous teleological approaches to epistemology. Whereas Craig and Weinberg suggest that we should examine the goals that we can achieve through our epistemic classifications, they fail to acknowledge quite how much of a contribution psychological results can make to our understanding of the goals that we can benefit from achieving. Psychological results can reveal how people can successfully/unsuccessfully obtain beliefs and, in the process, they reveal new goals that can be achieved in our epistemic classifications. It is my view, then, that progress can be made in epistemology by adopting a teleological approach to epistemology that takes seriously the possibility that science can reveal new goals that we can achieve through our epistemic classifications.

It would be too much to hope for consensus to be reached by epistemologists about either what goals can be achieved or about what account of knowledge (rationality or justified belief) describes it so that the epistemic concept can be used to achieve the important goals identified by the teleological approach to epistemology. However, by taking the approach outlined in this thesis, epistemologists will be able to avoid the many criticisms of traditional analytic epistemology, they will be able to focus on

debates about what is important about the categories that epistemologists aim to understand, such as knowledge, rationality and justified belief. The project of identifying those psychological results that are relevant could be a long and arduous one, and it might not be that everyone will agree about the significance of particular results. However, as we have seen in this thesis, there are great rewards awaiting those who pursue this task, so some degree of disagreement should not deter anyone.

Bibliography

Bealer, G. (1998) 'Intuition and the Autonomy of Philosophy' In Depaul, M and Ramsey, W. (eds.) *Rethinking Intuition: The psychology of intuition and its role in philosophical inquiry*, Lanham, Maryland: Rowman and Littlefield, pp. 201-239.

Bishop, M. (2000) 'In Praise of Epistemic Irresponsibility: How Lazy and Ignorant Can You Be?' *Synthese*, 122, pp. 179-208.

Bishop, M. A. and Trout, J.D. (2005) *Epistemology and the Psychology of Human Judgment*. New York: Oxford University Press.

Bishop, M. A. and Trout, J.D. (2008) 'Strategic Reliabilism: A Naturalistic Approach to Epistemology' *Philosophy Compass*, vol. 3, 5, pp. 1049-1065.

BonJour, L (1994) 'Against Naturalized Epistemology' *Midwest Studies in Philosophy*, XIX: 283-300.

Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999) 'Can Ignorance Beat the Stock Market?' In Gigerenzer, G., Todd, P. M. & the ABC Research Group, in *Simple Heuristics That Make Us Smart.* New York: Oxford University Press.

Boyd, R. (1988) 'How to be a Moral Realist' In Sayre-McCord, G. (ed.) *Essays on Moral Realism* Itaca, New York: Cornell University Press.

Boyd, R. (1991) 'Realism, Anti-foundationalism and the Enthusiasm for Natural Kinds', *Philosophical Studies,* vol. 61, 1-2, pp. 127-148.

Brandom, R. (2000) *Articulating Reasons: An introduction to inferentialism* Cambridge, Mass: Harvard University Press.

Carruthers, P. (2010) 'Introspection: Divided and Partly Eliminated', *Philosophy and Phenomenological Research* 80, 1, pp. 76-111.

Cherniak, C. (1986) *Minimal Rationality*. Cambridge, MA: The MIT Press.

Churchland, P. (1987) 'Epistemology in the Age of Neuroscience', *Journal of Philosophy* 84, pp.:546-53.

Conee and Feldman (2004) *Evidentialism: Essays in Epistemology.* Oxford: Oxford University Press.

Craig, E. (1990) *Knowledge and the State of Nature* Oxford: Oxford University Press.

Cummins, R. (1998) 'Reflections on Reflective Equilibrium' In Depaul, M and Ramsey, W. (eds.) *Rethinking Intuition*: *The psychology of intuition and its role in philosophical inquiry*, Lanham, Maryland: Rowman and Littlefield, pp. 113-127.

Daniels, N. (1979) 'Wide Reflective Equilibrium and Theory Acceptance in Ethics', *Journal of Philosophy,* 76, 5, pp. 256–82.

Dawes, R. (1994) *House of Cards: Psychology and Psychotherapy Built on Myth*, The Free Press, Macmillan, New York.

Depaul, M and Ramsey, W. (eds.) (1998) *Rethinking Intuition*: *The psychology of intuition and its role in philosophical inquiry*, Lanham, Maryland: Rowman and Littlefield.

Feldman, R. (1999) 'Methodological Naturalism in Epistemology' in *the Blackwell Guide to Epistemology*, Greco J. & Sosa. E (ed.) Malden, Ma: Blackwell, pp. 170-186.

Feldman, R. and E. Conee (1985) 'Evidentialism', *Philosophical Studies* 48, pp. 15–34.

Ferguson, M. J. (2007) 'The Automaticity of Evaluation' in Bargh, John, A. (ed.) *Social Psychology and the Unconscious: The Automaticity of Higher Mental Processes*, (New York, NY, US: Psychology Press), pp. 226.

Fodor, J. (1998) *Concepts: Where Cognitive Science Went Wrong* Oxford: Oxford University Press.

Gendler, Tamar Szabo (2008) 'Alief and Belief', *Journal of Philosophy,* Vol. 105: pp. 634-663.

Gendler, Tamar Szabo (2008a) 'Alief in Action', *Mind and Language,* Vol. 23: pp. 552-585.

Gettier, E. L. (1963) 'Is Justified True Belief Knowledge?', *Analysis*, Vol. 23, pp. 121–23.

Gigerenzer, G. (2007) *Gut Feelings* London: Penguin.

Gigerenzer, G. & Selten, R. (eds.) (2001) *Bounded rationality: The adaptive toolbox* (Cambridge, MA, US: The MIT Press).

Gigerenzer, G. & Todd, P. (1999) *Simple Heuristics That Make Us Smart* Oxford: Oxford University Press.

Goldman, A. (1967) 'A Causal Theory of Knowing', *Journal of Philosophy*, 64, 12, pp. 357-372.

Goldman, A. (1976) 'Discrimination and Perceptual Knowledge', *Journal of Philosophy,* 73, pp. 771-791.

Goldman, A. (1986) *Epistemology and Cognition*, Harvard University Press, Cambridge.

Goldman, A. (1993) 'Epistemic Folkways and Scientific Epistemology', *Philosophical Issues,* 3, pp. 271-285.

Goldman, A. (2007) 'Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status', *Grazer Philosophische Studien*, 2007, 74, pp. 1-26.

Goldman, A. & Pust, J. (1998) 'Philosophical Theory and Intuitional Evidence' In Depaul, M. and Ramsey, W. (eds.) *Rethinking Intuition: The psychology of intuition and its role in philosophical inquiry*, Lanham, Maryland: Rowman and Littlefield, pp. 179-197 reprinted in Goldman (ed.) (2002) *Pathways to Knowledge: Private and public*, New York: Oxford University Press, pp. 73-95.

Goldstein, D. and G. Gigerenzer (1999) 'The Recognition Heuristic: How Ignorance Makes us Smart' In G. Gigerenzer, P. Todd, and the ABC group (eds): In Press, *Simple Heuristics that Make Us Smart*, Oxford University Press, New York.

Goldstein, D. G. & Gigerenzer, G. (2002) 'Models of Ecological Rationality: The Recognition Heuristic', *Psychological Review*, 109, 1, pp. 75-90.

Goodman, N. (1955) *Fact, Fiction and Forecast* Bobbs-Merrill, Indianapolis and New York.

Greenwald, A.G., McGhee, D.E. & Schwartz, J.L.K. (1998), 'Measuring individual differences in implicit cognition: The implicit association test', *Journal of Personality and Social Psychology*, 74, 6, pp. 1464-1480.

Greenwald, A. G., Smith, C. T., Sriram, N,, Bar-Anan, Y., & Nosek, B. A. (2009), 'Race attitude measures predicted vote in the 2008 U.S. Presidential election', *Analysis of Social Issues and Public Policy,* 9, pp. 241-253.

Haack S. (1993) *Evidence and Inquiry: Towards Reconstruction in Epistemology* Oxford: Blackwell.

Jackson, F. (1998) *From Metaphysics to Ethics. A Defence of Conceptual Analysis.* Oxford: Oxford University Press, Clarendon Press.

Kahneman, D. and A. Tversky (1982) 'On the Psychology of Prediction', *Psychological Review* 80, pp. 237–51. Reprinted in D. Kahneman, P. Slovic, and A. Tversky (eds), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, pp. 48–68.

Kaplan, M. (1994) 'Epistemology Denatured' *Midwest Studies in Philosophy*, 19, pp. 350-365.

Kawakami, K., Dovidio, J. F., Moll, J., Hermsen, S., and Russin, A. (2000) 'Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on

stereotype activation', *Journal of Personality and Social Psychology*, vol. 78, 5, pp. 871-888.

Kim, J (1988) 'What is Naturalized Epistemology?' In Tomberlin, J.E. (ed.), *Philosophical Perspectives 2* Asascadero, CA: Ridgeview Publishing Co., pp. 381-406.

Kitcher, P (1992) 'The Naturalists Return', *Philosophical Review*, 101: 53-114.

Kornblith, H. (1994) 'Introduction: What is naturalized epistemology' In Kornblith, H. (ed.) *Naturalizing Epistemology* 2nd Edition, Cambridge: MIT Press.

Kornblith, H. (1998) 'The Role of Intuition in Philosophical Inquiry: An account with no unnatural ingredients' In Depaul, M.R. and Ramsey, W. (eds.), *Rethinking Intuition: The psychology of intuition and its role in philosophical inquiry*, Lanham, Maryland: Rowman and Littlefield.

Kornblith, H. (2002) *Knowledge and its Place in Nature* Oxford: Oxford University Press.

Kornblith, H. (2007) 'Naturalism and Intuitions', *Grazer Philosophische Studien*, 2007, 74, pp. 27-49.

Kornblith, H. (2011) 'Why Should We Care About the Concept of Knowledge?', *Episteme*, vol. 8, 1, pp 38-52

Kripke, S. (1972) 'Naming and Necessity', In *Semantics of Natural Language*, Davidson, D. and Harman, G. (eds.) Dordrecht; Boston: Reidel.

Kusch, M. (2009) 'Testimony and the Value of Knowledge' In Haddock, A., Millar, A. & Pritchard, D. (eds.) *Epistemic value* Oxford: Oxford University Press, pp. 60-94.

Lehrer, K. (1990) *Theory of Knowledge* Boulder, CO: Westview Press.

Meehl, P. (1954) *Clinical Versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence*, University of Minnesota Press, Minneapolis.

Meissner, C. A. & Brigham, J. C. (2001) 'Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review', *Psychology, Public Policy, and Law,* vol. 7, 1, pp. 3-35.

Nagel, J. (2007) 'Epistemic Intuitions', *Philosophy Compass* 2/6, pp. 792–819.

Nagel, J. (forthcoming) 'Intuitions and Experiments: A Defence of the Case Method in Epistemology', forthcoming in *Philosophy and Phenomenological Research.*

Neta, R. (2007) 'How to Naturalize Epistemology' In D. Pritchard & V. Hendricks (eds.) *New Waves in Epistemology*, London: Palgrave MacMillan.

Nichols, S., Stich, S. & Weinberg, J.M. (2003) 'Meta-Skepticism: Meditations in Ethno-Epistemology' In S. Luper (ed.), *The Skeptics* Aldershot, U.K.: Ashgate Publishing, pp. 227-247.

Nisbett, R. & Wilson, T. (1977) 'Telling More Than We Can Know: Verbal Reports on Mental Processes', *Psychological Review*, 84, pp. 231-95.

Nosek, B. A., Smith, F.L., Hansen, J.J., Devos, T., Lindner, N.M., Ranganath, K.A., Tucker Smith, C., Olsen, K.R., Chugh, D., Greenwald, A. G. & Banaji, M.R. (2007) 'Pervasiveness and correlates of implicit attitudes and stereotypes', *European Review of Social Psychology,* 18, pp. 36-88.

Papineau, D., "Naturalism", *The Stanford Encyclopaedia of Philosophy (Spring 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2009/entries/naturalism/>.

Pappas, G., (2008) 'Internalist vs. Externalist Conceptions of Epistemic Justification', *The Stanford Encyclopaedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/justep-intext/>.

Pinker, S. (1997) *How the Mind Works* New York: Norton.

Putnam, H. (1975) 'The Meaning of "Meaning"' In *Philosophical Papers* vol. 2, Cambridge: Cambridge University Press.

Quine, W.V.O. (1969) 'Epistemology Naturalized' in *Ontological Relativity and Other Essays*, New York: Columbia University Press, pp. 69-90.

Quine, W.V.O. (1969) *Ontological Relativity and Other Essays*, New York: Columbia University Press.

Rawls, J. (1972) *A Theory of Justice* Oxford: Oxford University Press.

Rosch, E., & Mervis, C. B. (1975) 'Family resemblances: Studies in the internal structure of categories', *Cognitive Psychology*, 7, pp. 573–605.

Samuels, R., Stich, S. & Bishop, M., (2002) 'Ending the Rationality Wars: How to Make Disputes About Human Rationality Disappear' in Renée Elio (ed.), *Common Sense, Reasoning, and Rationality* Oxford: Oxford University Press, pp. 236-268.

Schnall, S., Benton, J., & Harvey, S. (2008) 'With a Clean Conscience: Cleanliness reduces the severity of moral judgments', *Psychological Science,* 19, pp. 1219-1222.

Siegel, H. (1984) 'Empirical Psychology, Naturalized Epistemology and First Philosophy', *Philosophy of Science*, 51, pp. 667-676.

Smith, E. E., & Medin, D. (1981) *Categories and Concepts* Cambridge, MA: Harvard University Press.

Sosa, E. (1998) 'Minimal Intuition' In Depaul, M and Ramsey, W. (eds.) *Rethinking Intuition*: *The psychology of intuition and its role in philosophical inquiry*, Lanham, Maryland: Rowman and Littlefield, pp. 257-269.

Sosa, E. (2009) 'A Defence of the Use of Intuitions in Philosophy' In Murphy, D. & Bishop M. (eds.) *Stich and His Critics,* Malden, MA: Blackwell pp. 101-112.

Stein, E. (1996) *Without Good Reason: The rationality debate in philosophy and cognitive science* New York: Oxford University Press.

Stich, S. (1990) *The Fragmentation of Reason*, Cambridge, MA: MIT Press.

Stich, S. (1992) 'What is a Theory of Mental Representation', *Mind,* vol. 101, 402, pp. 243-261.

Swain, S., Alexander, J. & Weinberg, J. M. (2008) 'The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp*'*, *Philosophy and Phenomenologial Research*, 76, pp. 138-155.

Tversky, A., & Kahneman, D. (1974) 'Judgment under Uncertainty: Heuristics and Biases', *Science* 185, 1124–31.

Tversky, A. and Kahneman, D. (1982) 'Judgments of and by Representativeness' In Kahneman, D., Slovic, P. & Tversky, A. (eds) *Judgment under uncertainty: Heuristics and biases* Cambridge: Cambridge University Press.

Weatherson, B. (2003) 'What Good Are Counterexamples?', *Philosophical Studies*, 115, pp. 1-31.

Weinberg, J. (2006) 'What's Epistemology For?: The Case for Neo-pragmatism in Normative Metaepistemology' In Hetherington, S. (ed.) *Epistemology Futures*, Oxford: Oxford University Press.

Weinberg, J.M., Nichols, S. & Stich, S. (2001) 'Normativity & Epistemic Intuitions', *Philosophical Topics*, 29, 1 & 2, pp. 429-460.

Williamson, T. (2007) *The Philosophy of Philosophy* Oxford: Blackwell.

Wittman, P. (1941) 'A Scale for Measuring Prognosis in Schizophrenic Patients', *Elgin Papers*, 4, pp. 20-33.