



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

- This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.
- A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.
- This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.
- The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.
- When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

THE NATURE AND VALUE OF
KNOWLEDGE

Epistemic Environmentalism

SHANE RYAN

Submitted for the degree of PhD by research
The University of Edinburgh
2013

I have read and understood the University of Edinburgh guidelines on plagiarism. This thesis was composed by me and is entirely my own work except where I indicate otherwise by use of quotes and references. No part of it has been submitted for any other degree or professional qualification.

Shane Ryan

Acknowledgements

I thank my primary supervisor Duncan Pritchard. I am very grateful to Duncan for encouraging me to present at the 2010 *Episteme* conference which ultimately led me to my thesis topic. More generally Duncan's positivity and support has helped make the journey from the conception of my thesis topic to the completion of my thesis much lighter and less daunting than it would likely otherwise have been. I am also appreciative of the example he has given me as a philosopher. In particular, the clarity of thought and discrimination that shine through his writing, which helped me understand so much of my topic, has been a model to me.

I also thank Jesper Kallestrup, who became my secondary supervisor at the beginning of my third year. I have appreciated his support of my work.

I am grateful to the Department of Philosophy at the University of Edinburgh for awarding me a *Career Development Scholarship*.

There are many fellow PhD candidates at the University of Edinburgh who have contributed to the development of my work through their feedback at presentations and in conversation. In particular I thank Eusebio Waweru, Orestis Palermos, and Tony Bolos.

A special thank you goes to Andrea Kruse at Ruhr University Bochum whose enthusiastic and critical feedback on my work encouraged me enormously.

I am grateful to a number of people who entered and enriched my life during the course of my PhD, especially Patrizia Pecl, Olle Blomberg, Olga Krikun, David Peng and Sophia Hong, Naira Davlashyan, and Zarina Zaynutdinova. Thanks too to the squash and football players for some glorious games.

Finally, thank you to my parents.

Abstract

My thesis examines the nature and value of knowledge and normative implications of its value. With this in mind I examine Greco's account of knowledge in detail and consider whether it convinces. I argue against the account on a number of fronts; in particular I argue against Greco's treatment of the Barney and Jenny cases. In doing so I draw on the dialectic in the literature and go beyond it by showing how his treatment of those cases is such as to raise problems for his treatment of other cases. More specifically I argue that Greco's treatment of the Barney case is such as to threaten his treatment of standard Gettier cases and his treatment of the Jenny case threatens his treatment of the Careless Math Student case. I also consider an alternative virtue epistemic approach offered by Pritchard which I reject. In attempting to overcome the challenges that the Barney and Jenny cases pose I articulate an alternative account according to which what I call "epistemic grace" is a requirement of knowledge. It is via this epistemic grace requirement that I also account for the value of knowledge. Recognition of the value of knowledge serves as the basis for the articulation of the notion of epistemic environmentalism. With epistemic environmentalism in view, trust is analysed and its significance to the gaining of knowledge, albeit knowledge of a certain kind, is considered. Finally, the normative implications of epistemic environmentalism are laid out in a framework to show how findings in epistemic value theory relate to approaches that can provide a basis for justifying intervention or non-intervention in the assisting of the attaining or holding of epistemic goods of value.

Contents

Introduction.....	1
Chapter I: The Value of Knowledge.....	6
§1.0.0 <i>Introduction</i>	6
§1.1.0 <i>The Value of Knowledge: Intuitions and Problems</i>	7
§1.1.1 <i>Value</i>	8
§1.2.0 <i>Meno</i>	11
§1.3.0 <i>Greco's Virtue Approach</i>	13
§1.4.0 <i>Distinctive Value; A Counter Argument</i>	15
§1.4.1 <i>Cognitive Achievement without Knowledge? - JB</i>	16
§1.4.2 <i>Possible Response: Justification – Just for True Belief?</i>	21
§1.5.0 <i>The Value of Knowing Pointless Truths; A Counter Argument</i>	24
§1.5.1 <i>Knowledge of Pointless Truths</i>	25
§1.5.2 <i>A Possible Defence</i>	28
§1.6.0 <i>Conclusion</i>	29
Chapter II: Greco's Account of the Nature of Knowledge.....	31
§2.0.0 <i>Introduction</i>	31
§2.1.0 <i>Barney Cases on Greco's Account: An Introduction</i>	32
§2.1.1 <i>The Barney Case</i>	33
§2.1.2 <i>Greco's handling of the Barney Case</i>	34
§2.1.3 <i>Greco's Handling of the Roddy Case</i>	36
§2.1.4 <i>The Barney and the Roddy Cases Taken Together</i>	37
§2.1.5 <i>Greco: Possible Responses</i>	39
§2.1.6 <i>Epistemic Twin Earth</i>	42
§2.2.0 <i>Testimonial Knowledge on Greco's Account: The Dialectic</i>	44
§2.2.1 <i>A New Challenge</i>	49
§2.3.0 <i>Pritchard's Alternative Account</i>	54
§2.4.0 <i>Conclusion</i>	59
Chapter III: Epistemic Grace.....	61
§3.0.0 <i>Introduction</i>	61

§3.0.1 Chapter Motivation.....	62
§3.1.0 Testimony.....	62
§3.1.1 The Discourse on Testimonial Knowledge.....	64
§3.1.2 Virtuous Testimonial Belief: A Problem.....	67
§3.1.3 The Credulist Response.....	68
§3.1.4 The Humean Response.....	71
§3.1.5 Traits and Abilities.....	75
§3.1.6 Preliminary Conclusion.....	82
§3.2.0 Epistemic Grace: Revisiting the Gettier Problem.....	83
§3.2.1 Epistemic Grace: Gettier Cases.....	87
§3.2.2 Dealing with Gettier Cases.....	90
§3.3.0 A Possible Worry.....	92
§3.4.0 Conclusion.....	92
Chapter IV: Valuing Virtue and Grace.....	94
§4.0.0 Introduction.....	94
§4.1.0 The Value of Knowledge and that which Falls Short of Knowledge.....	95
§4.1.1 The Final Value of Knowledge?.....	102
§4.1.2 The Distinctive Value of Knowledge?.....	105
§4.2.0 Epistemic Value and the Epistemic Environment.....	106
§4.3.0 Conclusion.....	110
Chapter V: Trust.....	111
§5.0.0 Introduction.....	111
§5.1.0 An Analysis of Trust: Introduction.....	112
§5.1.1 The Nature of Trust: Distinctions.....	113
§5.1.2 The Nature of Trust: Requirements.....	116
§5.1.3 The Nature of Trust: The Goodwill View.....	118
§5.2.0 Trust: Expert, Epistemic Authorities and the Layperson.....	122
§5.2.1 Identifying Experts.....	122
§5.2.2 Epistemic Dependence.....	123
§5.3.0 Conclusion.....	127
Chapter VI: Epistemic Interventionism.....	129
§6.0.0 Introduction.....	129
§6.0.1 Background.....	129
§6.1.0 The Question.....	130
§6.2.0 The Status of Final Value Claims.....	134
§6.3.0 Possible Answers.....	136
§6.3.1 Paternalism.....	136

<i>§6.3.2 Soft Paternalism</i>	141
<i>§6.3.3 Libertarian Paternalism</i>	143
<i>§6.3.4 Legal Moralism</i>	145
<i>§6.3.5 Liberalism</i>	149
<i>§6.4.0 Conclusion</i>	150
Thesis Conclusion	151
Bibliography	158

Introduction

Why knowledge? Why has a whole branch of philosophy developed which has as one of its most central tasks, if not the most central task, the study of knowledge? One response might be to suppose there to have been a misstep along the way and that the focus on knowledge is actually unwarranted. That so much consideration and writing has been given over to trying to explain the nature of knowledge, rather than nearby alternative epistemic standings, and considerably less has been said about why knowledge rather than these alternatives is deemed worthy of such attention hasn't helped. Although this state of affairs is now changing, the question – why knowledge?, is having the effect of raising doubts about the focus of epistemology. Yet we have some comfort from the knowledge that successive generations of philosophers have concerned themselves with the nature of knowledge. Rather than think there has been some misstep along the way that has led to the focus on knowledge, an alternative possibility is that knowledge's draw, inherent in its nature, has gone unarticulated even by those investigating it, or at least, perhaps, until recently. In order to go some way towards answering the question – why knowledge?, my thesis investigates the value of knowledge, and does so by way of an examination of its nature.

A question that the emerging debate on the value of knowledge raises but which is yet to receive much exploration, is what are the normative implications of the value of knowledge, whatever that value may be. If knowledge is found to be valuable in a way that makes sense of the attention it has received, then in an assessment of the worth of social arrangements, the effect of those social arrangements on the attainment and preservation of knowledge may be a relevant consideration. A closely related question is aside from having implications for assessment of the worth of social arrangements, should that value, if found, prompt us to alter social

arrangements. Or, put differently, if social arrangements can be altered so as to effect a greater attainment of knowledge, then should those arrangements be so altered?

In my first chapter I set out the three value problems of knowledge, the solutions of which, if there are solutions, stand to provide a basis for making sense of the focus on knowledge in philosophy. Following Pritchard I examine Greco's (2010) account of knowledge from which it seems to follow that knowledge is distinctively valuable vis-à-vis other epistemic goods that fall short of knowledge. I set out Greco's argument for the value of knowledge, which falls out of his account of the nature of knowledge. I explicate Greco's argument that knowledge is valuable as a cognitive achievement and that achievements are finally valuable in virtue of being constitutive of the good life, while making clear that for Greco the value enjoyed by knowledge as a cognitive achievement is not supposed to be shared by standings which fall short of knowledge. I argue, however, that Greco's account as it currently stands does not show that knowledge is distinctively valuable. I do this by showing that justified belief is also plausibly conceived of as a cognitive achievement. Separately, I examine Greco's commitment to the claim that every instance of knowledge is an instance of something that has pro tanto final value. I challenge the claim by way of intuitions regarding cases of knowledge of pointless truths.

In Chapter II I examine challenges to Greco's account of the nature of knowledge, which draw on Pritchard's (2010) criticism that Greco's account is in one respect too strong. The challenge examined has potential implications for his account of the value of knowledge, given that his account of the value of knowledge falls out of his account of the nature of knowledge.

I assess how Greco's view, that knowledge is true belief because of cognitive ability, fares in accounting for an apparently typical case of testimonial knowledge, the Jenny case. In assessing Greco's view on this point I consider criticism from Pritchard, and Greco's response to that criticism. I examine Greco's original motivation for moving to his agent reliabilist account of knowledge. I argue that the Careless Math Student case, a case that Greco cites as providing impetus for such a

move, can be revamped in such a way so as to be structurally analogous with the testimony case, as described by Pritchard. The result, I argue, is that either Greco has to give up on the possibility of testimonial knowledge in cases currently thought of as typical or defend a counterintuitive knowledge ascription with regard to the revamped Careless Math Student case. I also examine the challenge Barney-type cases pose for Greco's account. With the support of an example I argue that Greco's treatment of the Barney case is such as to threaten the attractive way in which he has previously dealt with standard Gettier cases on his account.

Having identified arguments against the account that Greco provides in the first two chapters of my thesis, I return to trying to find an answer to the question with which we began by examining the nature and value of knowledge. In Chapter III I offer my own solution to overcoming the two challenges articulated in Chapter II. The account I advance not only handles the Jenny case but also testimonial knowledge in young children. It does so by drawing on a broader notion of virtuous belief than is employed by Greco; a notion according to which belief may be virtuous if it is from an epistemically virtuous trait or from a virtuous cognitive ability. I further argue that whether a belief is virtuous depends on factors such as the kind of agent believing and what the particular epistemic duties of the agent are. I then make the case for an unorthodox reading of Gettier-type cases, a reading that serves as inspiration for the articulation of an epistemic grace requirement for knowledge. While I claim that virtuous belief is a requirement for knowledge, I argue that a virtuous belief may or may not enjoy epistemic grace, the way the virtuous belief has been formed may or may not actually be a right way of forming a belief given the environment.

After providing an account of the nature of knowledge, the task in Chapter IV is to set out the value of knowledge and in particular examine the value of knowledge in comparison with that which falls short of knowledge. I argue that knowledge, having epistemic grace as a requirement, is more valuable and differently valuable than that which falls short of knowledge. I argue that there is a special value in gaining a good which requires a virtuous contribution from an agent but which cannot be secured by that virtuous contribution alone. Drawing on what I write in Chapter I, I argue that

although knowledge is differently valuable than that which falls short of knowledge, it doesn't enjoy a different kind of value than that enjoyed by achievements. The result is that while knowledge is special on my account, epistemic goods that fall short of knowledge are also finally valuable. Further drawing on my first chapter, I make the case that the claim that knowledge is finally valuable should be understood as being a claim that it is in the nature of knowledge to be finally valuable. Having shown how my account of the value of knowledge handles the value problems, I go on to consider what the broader normative implications are of the value of knowledge and the epistemic goods that fall short of knowledge. This brings me to the articulation of epistemic environmentalism – an approach according to which the value of epistemic goods are such as to, *ceteris paribus*, give us reason to alter or preserve social arrangements so as to attain or continue to hold valuable epistemic goods.

In Chapter V I examine kinds of relations that have a bearing on the quality of the epistemic environment. A common feature of these relations is that they require trust in order to have the potential to lead to knowledge. I provide an analysis of trust. In doing so, I distinguish it from reliance and identify goodwill and competence conditions. I examine Hardwig's (1985) claim that the non-expert is dependent on the expert for knowledge within the expert's domain. He discusses cases of scientific malpractice, suggesting that experts can't simply be relied upon. In order for the non-expert to gain knowledge in an expert domain according to Hardwig, the non-expert must trust the expert.

Having discussed how factors in the epistemic environment may affect the attainment of epistemic goods, I lay out approaches, with regard to social arrangements, that provide a basis of theoretical support for epistemic interventionism and that provide a basis of theoretical opposition to such interventionism. These theoretical approaches include hard paternalism, libertarian paternalism, soft paternalism, legal moralism, a liberal regulatory approach and political liberalism. In laying out a wide range of theoretical options, I am laying out the kind of commitments that may be made in any forthcoming debate that places

epistemic goods among goods, such as liberty and equality, traditionally central to political discourse.

Chapter I: The Value of Knowledge

§1.0.0 Introduction

In this chapter I argue that Greco's account of the value of knowledge is both too narrow in a certain respect and too broad in another. Before arguing for that claim, I make explicit some basic assumptions and distinctions with regard to the use of intuitions and kinds of values relevant in the discussion of the value of epistemic goods. I begin my discussion of the value of epistemic goods proper, in the traditional way, with consideration of the value problem as it is presented in Plato's *Meno*. While I highlight possible defences of the Socratic position that could be developed further, I make the case that Greco's account has some immediate theoretical advantages that make it worthy of consideration. By arguing that knowledge is valuable, and is valuable in a way in which standings that fall short of knowledge are not, Greco appears to solve various value problems that have been identified in the value of epistemic goods debate; that his account of the value of knowledge falls out of his account of the nature of knowledge adds to the theoretical attractiveness of his account.

I set out Greco's argument for the value of knowledge as a cognitive achievement and how his argument is taken to support the claim that the value of knowledge as a cognitive achievement is not one, in terms of degree or kind, shared by epistemic standings which fall short of knowledge. I make the case, however, that justified belief is also plausibly conceived of as a cognitive achievement. Developing this argument, I make the claim that Greco's account can best be understood as making an implicit assumption that justified belief is valuable only in so far as it helps get true belief; but I argue that such an assumption is mistaken. In this respect, I argue that his account of the value of knowledge is too narrow. I also consider a challenge to the claim that knowledge has pro tanto final value. I do so by drawing on

intuitions regarding a case of knowledge of a pointless truth. The case has a particular structure such that it cannot so easily be explained away by defenders of the claim that knowledge has pro tanto value as previous cases of knowledge of pointless truths have been. As a consequence, I argue, we get the result that on Greco's account instances of knowledge are attributed an implausible value. It is in this respect that his account of the value of knowledge is too broad.

§1.1.0 The Value of Knowledge: Intuitions and Problems

In this section, I make explicit some not very controversial methodological assumptions and other assumptions found in the discourse. Let's turn to methodological assumptions. First, consideration of intuitions is a good place to begin an assessment of something's value. Second, it counts as a mark in favour of an account of the value of epistemic goods if that account can make sense of our intuitions. Making sense of our intuitions can mean cashing out why our intuitions are right or diagnosing why they are misleading, or why some intuitions are right and others are misleading. Third, a counterintuitive account, by virtue of its claims or conclusion, has more to do to persuade us given that philosophical intuition that something is the case is taken to be prima facie support for that something being the case.

There are a number of purported intuitions pertaining to the value of knowledge that play a role in shaping the debate about the value of epistemic goods generally. There is a widely shared intuition that knowledge is valuable – I take it that just about every philosopher working in the area shares this intuition. The debate isn't focused on the value of knowledge in isolation, but rather tends to be on the value of knowledge vis-à-vis some or any epistemic standing that falls short of knowledge.

Duncan Pritchard (2010: 5-8) outlines three related problems concerning the value of knowledge:

1. The primary problem: why is knowledge more valuable than mere true belief?;

2. The secondary problem: why is knowledge more valuable than that which falls just short of knowledge?;
3. The tertiary problem: why is there a difference in kind rather than a mere difference of degrees between the value of knowledge and whatever falls short of knowledge?

The problems are problems in that knowledge is intuitively more valuable in the ways assumed in the problems but current accounts of the value of knowledge face challenges in either cashing out those intuitions or explaining them away in a satisfactory way. Furthermore, the relation between the problems is matryoshky¹ like in that a solution to the secondary problem promises to contain a solution to the primary problem and a solution to the tertiary problem in turn promises to contain a solution to both the secondary and primary problems.

While the intuition that knowledge is more valuable than true belief is broadly acknowledged, it's not clear whether other such comparative intuitions are shared. Miranda Fricker (2009: 127-128), considering the claim that knowledge has a value that is different in kind from the value of whatever falls short of knowledge, writes that the claim corresponds “to no natural philosophical intuition or question”. Similarly John Greco (2010: 97), responding to value problems raised by Jonathan Kvanvig, writes that “there is no pre-theoretical reason to think that knowledge is more valuable than its parts taken separately”. I think it is probably safe to take what he says here to imply that he thinks that there are no pre-theoretical intuitions that suggest that knowledge has such a comparative value. I am not going to try to show that either is intuitive; rather, by recording these challenges to claims about what our intuitions are I just want to make clear that where the burden in the dialectic lies has been disputed.

§1.1.1 Value

¹ “Matryoshky” is the name for traditional Russian nested dolls which form a set. A smaller doll is contained within a larger doll; the largest doll contains all the other dolls of the set.

Before getting to particular accounts of the values of epistemic goods, it will be helpful to set out standardly accepted kinds of value. Being clear on kinds of value also has special relevance in this chapter given that a significant portion of it examines the argument that Greco's account of the value of knowledge explains the distinctive value of knowledge, which, as the language used already suggests, requires distinguishing the putative kind of value that knowledge has from the putative kind of value that standings that fall short of knowledge have.

Definitions of the different kinds of value are as follows:

Extrinsic value: X is valuable on the basis of **a property external** to X;

Intrinsic value: X is valuable on the basis of **a property internal** to X;

Instrumental value: X is valuable as **a means** to something else;

Final value: X is valuable **for** its own sake.

Letting X be knowledge, we can say that knowledge may be regarded as being either extrinsically or intrinsically valuable, or as being instrumentally or finally valuable.²

What does it mean to say something has extrinsic value? It means that the thing has value in virtue of a property external to the thing in question. For example, the cup of water has value in virtue of a property not intrinsic to the water, the property of quenching thirst. To say that something has intrinsic value is to say that something is valuable in itself. More specifically, an intrinsically valuable good is valuable simply in virtue of a property internal to that good. Pleasure might be thought to be one such good.³

Alternatively a good may be instrumentally or finally valuable. An obvious example of a good with instrumental value is money; money is valuable as a means to getting things, such as nice clothes. To say that knowledge is instrumentally valuable is to

² Extrinsic and intrinsic value are standardly paired together as counterparts, as are instrumental and final value paired together. For more on this see Zimmerman (2010).

³ Claiming something to have intrinsic value, however, involves positing the existence of non-natural properties very different from natural properties like shape and mass. Zimmerman (2010) observes that the metaphysics of intrinsic value, including its very possibility, have increasingly come to be questioned since Moore's distinction between natural and non-natural properties.

say that knowledge is valuable in so far as it can help us get something else. Less familiar is final value. Rather than being about value in virtue of an internal property, a good that is finally valuable is valuable for its own sake either in virtue of relational properties or intrinsic properties. The first book produced by the first printing press is an example of something that is finally valuable. (Pritchard, 2010: 30). It is the relational property that the first book has to the first printing press that is significant for its value. Note that we are interested here in what is finally valuable rather than what is treated as finally valuable. Valuing a good or an object for its own sake does not imply that the good is finally valuable. A miser values money for its own sake but that does not mean that money is finally valuable. Money is not rightly valued for its own sake. To say that knowledge, for example, is finally valuable is to say that knowledge is valuable for its own sake and that its value doesn't reduce to instrumental value.⁴

Accounts of the value of knowledge focus on the prima facie or pro tanto value of knowledge rather than attempting to provide either an all things considered account of the value of knowledge or an assessment of the worth of individual instances of knowledge such as, for example, knowledge that saved a person's life. An all things considered account would say what value knowledge has after external factors that influence overall value are taken into consideration.

By saying that knowledge has prima facie value x , one is saying that on the face of things knowledge has this value and therefore it is appropriate to treat an individual instance of knowledge as having this value unless there is reason to think it doesn't have this value. By saying that knowledge has pro tanto value x one is saying that knowledge always has this value but that this value may be outweighed by other factors in all things considered accounts of its value.⁵

⁴ See Mark Schroeder (2008) for a brief discussion of how "intrinsic value" is sometimes used in a way such that final value is subsumed by the term.

⁵ An alternative way to explain the difference is to do so in terms of defeasibility. A good that is prima facie valuable may in a particular instance have that value undercut such that the particular instance has no value; similarly, a good that is prima facie valuable may in a particular instance have that value overridden such that, although the good remains valuable to some extent, it is not all things considered valuable. In contrast, a good that is pro tanto valuable may not have its value undercut.

§1.2.0 *Meno*

Discussion of the primary value problem arises in Plato's *Meno*. The discussion is preceded by an exchange between Socrates and Meno in which Socrates claims that the good man must be useful and one way of him being so is by correctly guiding others in their affairs. (Plato, 1963: 380). The value problem arises as Socrates seeks to explain his claim that it appears to be wrong to assume that the good man needs knowledge to do good.^{[6][7]}

Socrates begins his investigation into distinguishing the value of knowledge from true opinion by way of an example.^{[8][9]} He asks Meno what difference there would be in the correctness of guidance given by someone who knows the way to Larissa and someone who has a true opinion about which is the road there but who has never travelled that road before. Meno accepts that the guide who has a true opinion, so long as it remains, would be just as good a guide as the man who knows. (Plato, 1963: 381).

Socrates, considering why one should be preferred to the other, claims that the superiority of knowledge lies in its keeping a true opinion in place; knowledge does

⁶ Benjamin Jowett (Plato, 2008: 51) translates Socrates as saying it is wrong, whereas W. K. C. Guthrie (Plato, 1963: 380) translates Socrates as saying that it appears to be wrong.

⁷ Ethics, knowledge, and inquiry are each discussed in the *Meno*, but a reoccurring topic is virtue.

⁸ The discussion of the distinction is brief and what is meant can be interpreted in different ways. Interpreting the *Meno* is generally a knotty task as the dialogue quickly traverses a range of topics, makes numerous references – the significance of which are obscure, and many of the comments made in the dialogue seem best understood as displays of Platonic irony. For example, near the beginning of the dialogue Socrates tells Meno that he, Socrates, is a forgetful person and towards the end of the dialogue Socrates asks Meno to convince Anytus about what Meno believes to be true following their discussion of virtue. (Plato, 1963: 354, 384). These comments are ironic because during the dialogue the case is made that knowledge comes about through a process of recollection and that virtue can't be taught. (Plato, 1963: 371, 380). Furthermore, Dominic Scott (2006: 163) tells us that Anytus "was one of two people most active in bringing about Socrates' trial and execution". A joint reading of the *Meno* and Robin Waterfield's (2010) *Why Socrates Died* creates an impression that the *Meno* is closely connected to Socrates' role and responsibility in Athenian society.

⁹ Note that "true opinion" here is often read as true belief, see Pritchard (2008) for an example; when this is so the question of comparative value that is being raised relates to the primary problem. (Pritchard, 2008). It would be a mistake to simply assume that Socrates' use of what is translated as "knowledge" and "true opinion" or "true belief" fits neatly with how we understand those terms today. For example, Scott (2006: 179) interprets Socrates' use of, what is often translated as, "knowledge" as playing the role of understanding why something is the way it is. Nevertheless, the Socratic discussion still serves as a rich starting point for contemporary discussions of the value of knowledge.

so by tying a true belief down with explanatory reasoning.¹⁰ One way of thinking about this is that when we gain knowledge we gain a kind of epistemic foothold in reality or the world that won't give way as a mere true belief will. As Plato writes; "[t]rue opinions are a fine thing and do all sorts of good so long as they stay in their place, but they will not stay long. They run away from a man's mind; so they are not worth much until you tether them by working out the reason."¹¹ (Plato, 1963: 381). Returning to the original example, if you know the way to Larissa then your guidance of others to their destination will be sure. If, however, you merely have a true belief, then that true belief may easily be lost, perhaps by evidence that suggests your belief is mistaken.

On the one hand, however, stably held true beliefs are not exclusively true beliefs that are known. A dogmatic true belief may be stably held without us taking that true belief to be knowledge. In this case we get the counterintuitive result that such dogmatic beliefs have the same value as knowledge. On the other hand, knowledge may be lost in the face of an apparent though not actual defeater, and possibly, in some cases at least, more easily so than a true belief. Kvanvig (2003: 15-16) describes a case in which my mathematical knowledge that p would be lost if a renowned mathematician sincerely testified that not p . Not knowing that there had been any such testimony and that testimony in fact being false, I could retain my true belief that p but lose my knowledge that p . It seems plausible to think that we do have certain beliefs and knowledge that are of a kind such that my knowledge is more unstable in the Socratic sense than my true belief. Scientific and technical knowledge that I amassed in the past through expert testimony, perhaps indirectly in school, but which I encounter little evidence for or against now seem to be especially like this. If this is right then some knowledge would be less valuable than some mere true beliefs.¹²

¹⁰ More specifically, Socrates claims that the explanatory reasoning or the working out of the reason for something is accomplished by the process of recollection. (Plato, 1963: 381).

¹¹ Socrates uses the statues of Daedalus, which, because they're so life-like, will run away if not tied down as a metaphor for true belief. It would be interesting, in the broader context of interpreting the *Meno*, to unpick why this particular metaphor is chosen. Daedalus is a mythical master craftsman, some of whose creations, such as Icarus' wings and the labyrinth of the Minotaur, led to tragedy.

¹² Fricker (2009) discusses, in the context of discussing the value of knowledge, both dogmatic true beliefs and knowledge that is more easily lost than true beliefs. The latter would be particularly

One way of interpreting the claim of knowledge's stability, is as being the claim that stability will generally be present in cases of knowledge, whereas this is not so in the case of true beliefs. This could yield the conclusion that knowledge is generally more valuable than true belief. But granting that true beliefs that have greater stability than knowledge are more valuable than knowledge seems counter-intuition; and not only that, such a response wouldn't leave us well-placed to address the second and third value problems; it looks like we would need a different explanation as to why knowledge is more valuable than justified true belief. There is therefore some motivation to look beyond the *Meno* for an account that better makes sense of our intuitions.

An alternative line of argument that might be developed, based on Socrates' emphasis on stability conferring extra value on knowledge, is to suppose that if a true belief is not truly grounded or bound then it is not knowledge. The thought is that when something is really known then there just aren't new considerations about the world that can shake it loose; it will remain stable. Such an account of knowledge resembles that of Descartes'; knowledge is something secure, it can't be called into doubt. On this view, there may be dogmatically held true beliefs that will be just as stable as true beliefs that are known; therefore some true beliefs will be just as valuable as knowledge, but there won't be any true beliefs that enjoy greater value than knowledge. This seems, however, to be setting the bar too high for what can be knowledge. More will be said on this when I turn to discuss testimonial knowledge in Chapter 3.

§1.3.0 Greco's Virtue Approach

Greco endorses the view that a good theory of knowledge will say something about both the nature of knowledge as well as the value of knowledge, and that

troubling if it could be shown that knowledge more often has this feature in comparison to true belief, as then the Socratic account of the value of knowledge would be untenable; knowledge just wouldn't be more stable than true belief. In the broader discussion of the value of knowledge, Fricker herself defends a *Meno* type response according to which knowledge is generally more valuable than true belief as knowledge generally yields resilient true belief.

understanding the nature of knowledge will help us understand the value of knowledge. (Greco, 2010: 71). Indeed, his account of the value of knowledge falls out of his account of the nature of knowledge. (Greco, 2010: 97). What's more, his account of the value of knowledge is one which, if correct, explains not only why knowledge is more valuable than true belief, but also why knowledge is more valuable than that which falls short of knowledge and why knowledge is distinctively valuable from that which falls short of knowledge.

According to Greco, S knows that p, if and only if S believes truly **because** of S's cognitive ability.¹³ For example, S's believing the truth that 28×9 is 252 is explained by S's believing from ability, more precisely a reliable cognitive ability. For Greco, a cognitive ability is by its nature reliable and is a process grounded in the cognitive character of an agent. An upshot of this account, it is claimed, is that no separate condition is needed for dealing with Gettier-type cases. (Greco, 2010: 75). That Greco claims that knowledge is **exclusively** accounted for as a true belief because of "epistemically virtuous belief forming processes", means that the kind of virtue epistemology which he endorses is what Pritchard (2010: 24) has termed "robust virtue epistemology".

How does his account of the nature of knowledge help us to understand the value of knowledge? Greco (2009: 318) argues that knowledge is valuable in a way that mere true belief is not because it is a cognitive success that is creditable to the agent while mere true belief is not. Knowledge is creditable in that it is a cognitive success that is because of the exercise of one's cognitive abilities. Greco characterises success from ability as being an achievement, and knowledge as being a cognitive success from a cognitive ability as being a cognitive achievement. He contrasts someone gaining true beliefs because of their cognitive abilities with getting things right by "blind chance, dumb luck, or something else". (Greco, 2009: 318). If I luckily happen to have a true belief, then we would neither suppose my belief to be a case of knowledge nor a cognitive achievement creditable to my cognitive abilities. Particularly pertinent to one of my arguments later is his position on justified though

¹³ Greco (2009: 318) sees intellectual abilities as a species of intellectual virtue.

false belief. He writes that while knowledge is a success from cognitive ability and therefore a cognitive achievement, justified though false belief, along with any other subset of the constituents of knowledge, including justified true belief, is not a success from cognitive ability and can't have the value that a cognitive success from cognitive ability has.¹⁴

Referencing Aristotle, Greco (2010: 97-98) holds that achievements are constitutive of human flourishing and so are finally valuable; and knowledge, as an achievement, therefore is finally valuable. More recently, and more modestly, Greco (2011) has simply held achievements to be finally valuable, valuable for their own sake, without drawing on an Aristotelian account of human flourishing.¹⁵ Given his achievement account of knowledge, Greco is well placed to answer the tertiary value problem and so solve each of the other value problems as well.

Borrowing from Pritchard (2010: 31), Greco's argument is laid out below:

- (P1) Achievements are successes that are because of ability (Achievement thesis);
- (P2) Knowledge is a cognitive success that is because of cognitive ability (Robust Virtue Epistemology);
- (C1) So, knowledge is a cognitive achievement (Knowledge as Achievement thesis);
- (P3) Achievements are finally valuable (Value of Achievements thesis);
- (C2) So, knowledge has final value.

§1.4.0 Distinctive Value; A Counter Argument

¹⁴ He claims that "virtuous belief that is not true" is not "intrinsically valuable, or constitutive of what has intrinsic or final value, in just the way that knowledge is." (Greco, 2010: 99). In the context of his overall account, what he writes here is frustratingly vague. He seems to be leaving open the possibility that virtuous belief that is not true might also have the same kind of value as knowledge albeit in some other way. However, Greco (2011) clearly endorses the knowledge as achievement argument as providing a solution to the tertiary value problem, i.e. showing that knowledge is distinctively valuable vis-à-vis that which falls short of knowledge.

¹⁵ Greco (2009, 2010) previously characterised knowledge as having intrinsic value, and later as having both intrinsic and final value.

I make the case that even granting the assumptions and conclusions of Greco's argument, we have reason to reject the claim that knowledge should be thought of as distinctively valuable. More precisely, I challenge his denial of justified belief being a kind of cognitive achievement.^[16]^[17] I do so based on the plausibility of the argument below:

- (P1) Achievements are successes that are because of ability (Achievement thesis);
- (P2*) Justified belief is a cognitive success that is because of cognitive ability;
- (C1*) So, justified belief is a cognitive achievement;
- (P3) Achievements are finally valuable (Value of Achievements thesis);
- (C2*) So, justified belief has final value;
- (P4*) Knowledge has final value;
- (C3*) Knowledge and justified belief have the same kind of value.

§1.4.1 Cognitive Achievement without Knowledge? - JB

Greco needs there to be no cognitive achievement creditable to the cognitive agent that is not knowledge but which falls short of knowledge, and for knowledge to be such an achievement; otherwise his claim that knowledge is distinctively valuable is undermined. My objection is that accounting for the value of knowledge as an achievement does not show that knowledge is distinctively valuable. If other epistemic standings that fall short of knowledge are also achievements then the achievement account can't show knowledge to have distinctive value. Both justified true belief and justified belief are plausibly achievements. My argument is intended to show both that it does not follow from Greco's argument that knowledge is distinctively valuable and that the account of the value of knowledge makes an implicit functionalist or teleological assumption. Understanding the assumption can in turn help us to understand why justified belief is not regarded as an achievement;

¹⁶ It's important to note that I think of "justified belief" as being interchangeable with "epistemically virtuous belief". This becomes especially relevant in Chapter III.

¹⁷ Another way of challenging the conclusion that knowledge is distinctively valuable is to argue that true belief is also finally valuable, albeit not in virtue of being an achievement. I discuss the value of true belief in Chapter IV.

mere justified belief cannot be a cognitive success because a justified belief that isn't true isn't successful.

As has been noted, Greco's account can say why knowledge is more valuable than mere true belief; mere true belief can't plausibly be thought of as any kind of achievement creditable to the cognitive abilities of the believer.¹⁸ But does Greco's account help with justified beliefs? Shouldn't justified beliefs be thought of as a kind of achievement?

Greco's account of the value of knowledge is general in the sense that, as we are told, knowledge is finally valuable because it falls into a category, the category of achievement, and that which is an achievement is finally valuable. The account is not of knowledge's value per se. A dialectical worry with such an approach, though this may reflect how things are and so be warranted, is that it leaves open the possibility that there are cognitive achievements other than knowledge. A few candidates of what might be counted as cognitive achievements, or what Kvanvig (2005: 286) describes as "successful cognition", are: "tenets that are empirically adequate", "justified assumptions", "understanding" and "responsible inquiry". These can plausibly be regarded as achievements and hence be valued for the same reason and in the same sense as knowledge; if this is so, then knowledge is not the only epistemic good that is an achievement, and therefore knowledge does not have distinctive value vis-à-vis other epistemic goods. What is important for my purposes is to show that the claim that knowledge has distinctive value should be understood as claiming that knowledge has distinctive value in relation to standings that fall short of knowledge, not in relation to epistemic standings generally.¹⁹

It is important to note here that I am not interested in the claim that knowledge is distinctively **epistemically valuable**. Greco's account of the value of knowledge, though not Kvanvig's, is not analysing value in a strictly epistemic way. For example,

¹⁸ In Chapter IV I consider whether true belief is also finally valuable, though not as an achievement.

¹⁹ One might wonder, however, if knowledge has the same value as all the goods that Kvanvig lists, then why it is that knowledge receives so much attention from epistemologists. Again this is something I return to in Chapter IV.

knowledge is said to be valuable as an achievement but an achievement, even a cognitive achievement, is not something that is inherently epistemic in nature. That Greco's account of the value of knowledge does not restrict itself to the epistemic is unsurprising for a virtue account, this is given that the exercise of virtue generally is taking to be constitutive of the good life.

Greco accepts as plausible that the successful exercise of intellectual courage is "intrinsically good, and constitutive of the best intellectual life". He seems to do likewise when it comes to wisdom and understanding. He writes that; "[o]n the view that results, there is a plurality of intellectual virtues, and their successful exercise gives rise to a plurality of epistemic goods. The best intellectual life – intellectual flourishing, so to speak – is rich with all of these." (Greco, 2010: 98). Remember, an achievement according to Greco is a success from cognitive ability, and a cognitive ability is a species of virtue. The implication is that these other things are also achievements. He is not saying that knowledge is the only epistemic good with final value. The best way then to understand how his account can hope to show that knowledge is distinctively valuable is to understand the claim as being restricted to the value of knowledge in relation to standings that fall short of knowledge.

If it can be demonstrated that Greco's account doesn't show that knowledge enjoys distinctive value in relation to justified belief, then the same will hold for justified true belief; demonstrating this will also indicate that there is a gap in his account. If such a gap is indicated then I believe it will not be easy for Greco to fill this gap. Intuitively, it even seems plausible to think that getting justified beliefs regardless of their truth is an achievement. If we do regard justified beliefs as achievements then they meet the criteria for the kind of value that knowledge has on Greco's account.

My opponent might simply deny such an intuition and say that a belief that is not true but is held because of the cognitive ability of the believer, and so is justified, is not a cognitive success and therefore cannot be an achievement. Having a justified belief shouldn't be thought of as a cognitive success but rather, in so far as it can be a

cognitive success, it is only a part of the cognitive success that knowledge constitutes; a footballer who takes a good shot but who doesn't score doesn't have success.

But notice that to keep the example analogous to what's going on when there is a justified belief, and not merely a belief, I have to talk about a "good shot". In what sense is taking a good shot not a success and not an achievement? It's good, and it's good rather than bad or turning out to be lucky, because of the footballing ability of the player.

Even if my opponent holds fast to the view that having a justified belief should not be considered a discrete success, the onus is on her to explain why knowledge and true belief, but not justified belief, is such a success. She can't just assume that true belief is the goal and that anything that falls short of that necessarily can't constitute a cognitive success. I take it that I've given grounds to worry about the capacity of Greco's account to give a satisfactory explanation as to why cognitive success because of ability excludes standings that fall short of knowledge. Next, I consider what moral should be drawn from this worry.

Greco's account tries to explain the value of knowledge by saying that knowledge is valuable as an achievement. One might think that this doesn't work because achievement as a category isn't sufficiently demanding to ultimately exclude justified belief. Believing this one might seek to retain the form that Greco's argument takes but substitute some other category that knowledge falls into in place of achievement. The argument form being that knowledge is valuable as it is an instance of kind X and instances of kind X have value Y. For this to work the category should plausibly be the kind of thing that is finally valuable or at least of distinctive and superior value vis-à-vis whatever value justified belief has, as well as any other standings that fall short of knowledge. The approach of finding a substitute for achievements raises a number of issues but the point that is salient for my purposes is that it misses a deeper worry regarding Greco's account.

The denial of justified belief as a cognitive success seems best explained as being the result of an implicit functionalist or teleological assumption.²⁰ The purpose of justification/cognitive ability is just to get something else and as such justified beliefs or beliefs formed because of cognitive ability should never be thought of as cognitive successes in their own right. The something else can't be just anything else. Greco certainly thinks that if ability is the cause of the belief one has being right then there is an achievement and may, as previously mentioned, think likewise when it comes to understanding and wisdom.²¹ Seeing as I've already argued that it is best to understand the claim about the distinctive value of knowledge as being in relation to states which fall short of knowledge, it seems right to say that for Greco, in the restricted set of knowledge and that which falls short of knowledge, justification/cognitive ability is just for getting true belief; if they don't get that then there is no chance for an achievement. Put another way, in the restricted set that we're confining ourselves to, the only kind of cognitive success is one which at least contains true belief. If justification/cognitive abilities can't hit that target, then there can be no cognitive success and therefore no achievement.

The problem with running the value of knowledge argument by way of the Value of Achievement thesis, or any alternative to cognitive achievement, in order dialectically to get the claim that knowledge is distinctively valuable, is that, as things stand, what should be counted as success remains unclear or at least contentious. Without a premise as to what counts as a success, whether it be an implicit teleological assumption or otherwise, what falls into the category of achievement, if we keep the basic form of Greco's argument, will also remain contentious. For example, if S is just aiming at having a justified belief, then if S succeeds in his aim because of his ability, the result is that it seems appropriate to say that S is exhibiting a kind of achievement.

²⁰ Fellow virtue epistemologist, Ernest Sosa (1991: 225, cited in David 2001: 152), is explicit in his assumption of a "teleological conception of intellectual virtue, the relevant end being a proper relation to truth." For Sosa, an intellectual virtue is a quality that helps maximise truths while minimising errors.

²¹ Neither understanding nor wisdom need be thought of simply as more complex forms of knowledge. For example, understanding is plausibly conceived as not requiring factivity.

The claim that there is an implicit functionalist assumption is not being put forward here as a reason to claim that the account can't show that knowledge is distinctively valuable. An argument is required from a defender of this type of account, however, given that it seems plausible to think of justified belief as also being a cognitive success and being a cognitive success that is because of cognitive ability.

It might be thought that what I have shown is not that there is a good case that all justified beliefs are achievements but rather that only some are, and that all I need to do to make trouble for Greco's account is to show that some are. For example, if someone exercises impeccable reasoning in forming a belief, then they will be justified in believing as they do. The case of impeccable reasoning might be thought of as a case of an achievement; though, as not all justified beliefs are cases of impeccable reasoning, this would not imply that all justified beliefs are achievements. However, if impeccable reasoning is accepted as an achievement, has the case been made that knowledge does not have distinctive value vis-à-vis a standing which falls short of knowledge, justified belief?

My worry that the case has not been made relates to the fact that impeccable reasoning isn't simply a justified belief, impeccable reasoning includes something that goes beyond what is necessary for a belief to be justified – and also for being a necessary condition for knowledge. Impeccable reasoning does fall short of knowledge in the sense that it needn't be the case that the belief formed on its basis is true. Greco, however, already seems to be onboard with there being a plurality of epistemic goods having final value and he could claim that the part that goes beyond what is required for justification is what makes it seem plausible that it is of the same kind of value as knowledge. Someone can believe that some of the goods that have final value contain a justificatory element but still believe that the value of knowledge is of a different kind than that of standings which fall short of knowledge, including justified belief. To put it another way: even if A contains B, C contains B; A has value X, C has value X; B does not necessarily have value X.

§1.4.2 Possible Response: Justification – Just for True Belief?

In this section I consider two possible responses to my objection that there is a gap in Greco's argument. First I will consider a biological functionalist argument in support of the claim that true beliefs are cognitive successes and that justified beliefs are not. The second argument that I consider is a teleological argument about the nature of justification. The claim to be examined is that justification is only valuable in so far as it helps to get true beliefs and, as a result, if a belief is justified but is not true, then it cannot be thought of as a success.

Greco might try to justify the claim that justified beliefs should not be thought of as discrete cognitive successes by way of a biological functionalist argument. The argument might be that we need true beliefs to interact with our environment in such a way as to ensure at the least that we survive and pass on our genes or whatever counts as minimal evolutionary success. Whether those beliefs are justified or not does not ultimately matter, success is measured at least in terms of survival and procreation. Success is ensured when our beliefs are true. When we are gaining true beliefs our cognition is working well, that is, we are assisted with survival and procreation. Therefore, true beliefs are cognitive successes.

If we were to accept such a framing of cognitive success, however, then it would be implausible to suppose that cognitive successes for cognitive agents are restricted to having true beliefs and that all true beliefs should be considered cognitive successes. Plausibly, some false beliefs, say that contribute to our well-being, could be counted as cognitive successes in that they assist with evolutionary success. For my purposes here, however, we should only concern ourselves with whether it is appropriate to conceive of cognitive success as including true belief and excluding justified belief. This is our concern as cognitive success as including true belief is Greco's positive claim, and justified belief, a standing which falls short of knowledge, has been proposed as offering a counterexample with regard to Greco's claim about the distinctive value of knowledge.

Should all true beliefs be regarded as cognitive successes if we understand cognitive success in the biological functionalist way outlined above? Some true beliefs may actually undermine our chances of evolutionary success. For example, it is conceivable that a particular true belief gained by an agent could trigger an agent committing suicide or losing their mind. Therefore, those individual true beliefs cannot credibly be counted as cognitive successes if cognitive success is determined by evolutionary success.

A possible response is to say that cognitive success or failure should not be judged according to the veracity of individual beliefs in isolation, but according to the truth or falsity of a belief and the contribution it makes in that regard to other beliefs. The former scope of assessment could be misleading as an individual belief ordinarily contributes to the holding of other beliefs and if we were to simply judge an individual belief a success or failure in virtue of its veracity then we might miss out on the aspect just mentioned. But Greco does judge cognitive success on the basis of the simple truth or falsity of a belief and so doing is central to his claims about the nature of knowledge. He can't claim that a false belief that results in a greater number of true beliefs is a cognitive success because, for Greco, knowledge is a cognitive success that is because of cognitive ability, and of course he doesn't want to say that we can know something that is false.²²

Alternatively Greco might want to draw on a teleological view of the value of justification. Lynch (2004, p. 50) describes such a view thus: “[o]nce again, the key point is that the value of believing what is justified is parasitic on the value of believing what is true. Having justified beliefs is good because justified beliefs are likely to be true.”

That justified beliefs should not be supposed to be discrete cognitive successes finds support in the claim that justification is desirable only because truth is desirable. Justified beliefs are more likely to be true than unjustified beliefs, therefore by valuing true beliefs, we derivatively value the means of securing true beliefs,

²² Even if Greco sought to appeal to the nature of true belief as a kind, there seems nothing that links true belief in any necessary way with evolutionary success.

justified beliefs.²³ If this is right, then Greco can argue that justified beliefs should not be deemed successes in isolation from consideration of the truth value of the proposition believed; rather it's having true beliefs, what was desired all along, which is the determinant of cognitive successes. To put it differently, if we want A just because we want B then it might seem strange to say that we were successful when we got A and not B.²⁴

We do, however, have the intuition that justified beliefs are valuable and seemingly independent of whether they get us true beliefs. There's no obvious sense in which the teleological claim is necessary.²⁵ What's more the claim about the value of knowledge is not a claim as to the epistemic value of knowledge; rather it is a claim about the value of an epistemic good, which is owing to that epistemic good falling into a broader, non-epistemic, category of goods. One may grant that the epistemic value of justified belief is derivative of the value of truth and continue to hold the plausible view that justified belief meets the criterion for achievement and as such has the value of achievement.

§1.5.0 The Value of Knowing Pointless Truths; A Counter Argument

In this part of my chapter I offer an argument against the claim that knowledge has pro tanto final value. More specifically, I challenge the Value of Knowledge Thesis by arguing that knowledge of pointless truths does not have pro tanto value or so little pro tanto value so as to be insignificant.²⁶ There is an implicit assumption in this argument that the claim that knowledge is an achievement, Greco's claim,

²³ The claim that justified beliefs are more likely to be true than unjustified beliefs is highly plausible though one could imagine an evil demon world in which it is not the case.

²⁴ For more on the value of justification being tied to the value of truth see David (2005).

²⁵ There is far from being a consensus on whether it is appropriate to say that belief has any aim in an interesting sense at all and what that might be, even if it does. (Owens: 2003). Chignell (2010) writes that: "common candidate here... is knowledge itself (see Williamson 2002, Pritchard 2007...), but some authors claim that justification (Adler 2002) and/or "virtue" in our beliefs is the aim (Zagzebski 2004), while still others plump for a more structurally complex aim such as "understanding" (Kvanvig 2004)." (The years of publication that I have for these works of Williamson and Kvanvig are 2000 and 2003 respectively, rather than the years that Chignell lists.)

²⁶ My focus here is on showing that knowledge of pointless truths, or pointless achievements, have no pro tanto final value or none that is significant. Elsewhere it has been likewise been argued that immoral achievements have no value. For more on this see Ridge (2011: 21).

implies that each instance of knowledge is an achievement. Later in this section I consider a response defending Greco's claim by, in part, denying this implicit assumption.

Recall Greco's argument runs as follows:

- (P1) Achievements are successes that are because of ability (Achievement thesis);
- (P2) Knowledge is a cognitive success that is because of cognitive ability (Robust Virtue Epistemology);
- (C1) So, knowledge is a cognitive achievement (Knowledge as Achievement thesis);
- (P3) Achievements are finally valuable (Value of Achievements thesis);
- (C2) So, knowledge has final value.

My counter argument is:

- (P1) Achievements are successes that are because of ability (Achievement thesis);
- (P2) Knowledge is a cognitive success that is because of cognitive ability (Robust Virtue Epistemology);
- (C1) So, knowledge is a cognitive achievement (Knowledge as Achievement thesis);
- (P3*) Knowledge of pointless truths is not finally valuable (contra the Value of Achievements thesis);
- (C2*) So, not all knowledge has final value.

§1.5.1 Knowledge of Pointless Truths

If knowing is always an achievement and all achievements are finally valuable then all knowledge is finally valuable. A consequence of taking knowledge to be so valuable is that unqualified it implies that every instance of knowledge is finally valuable. If every instance of knowledge has final value then knowledge of so-called pointless truths is finally valuable. An example of knowledge of a pointless truth is

knowledge of the number of blades of grass in a garden.²⁷ But knowledge of pointless truths intuitively isn't valuable, let alone finally valuable.

One obvious response that may seem to accommodate this intuition is to say knowledge has pro tanto final value. To say this is to say that the value of knowledge may be outweighed when other factors are taken into account. With this in mind the defender of the Value of Achievement Thesis might claim knowledge of the number of blades of grass in a garden is valuable, it's just that the time and effort required to know the truth is such that the value of the knowledge is outweighed such that it does not have all things considered final value but rather has pro tanto final value.

But on closer inspection this response is unsatisfying. This defence of the Value of Achievement Thesis works by pointing out that the costs or disvalue of gaining knowledge of a pointless truth may be such as to outweigh the value of gaining knowledge of that truth, but knowledge of pointless truths needn't be costly to gain. Knowledge of some such truths may be gained cheaply. The problem is raised once more as to whether we consider such knowledge finally valuable in the first place at all and if so why. Consider the following case:

Pierre sits in a café by the window looking out onto a relatively busy Parisian side-street. He decides to gain knowledge by counting the number of people who pass by his table on the street outside between every two sips of his coffee. He comes to know that five people passed between his first sip and his second sip, seven people passed between his second sip and his third sip, etc. Let us further add that there is no opportunity cost worth considering, he couldn't have been doing anything else that is valuable.²⁸

If knowledge is finally valuable then it seems that if there is nothing else Pierre could be doing that is valuable then he should count between sips to clock up more and more value. That there may be nothing else Pierre could be doing that is valuable

²⁷ The examples of pointless truths in the literature usually involve some count (blades of grass, strands of hair, and grains of sand) that reaches a very high number.

²⁸ To deal with Pritchard's (2010) argument that for knowledge to be a cognitive achievement then a skill must be exercised, we can add that a special counting skill is being used.

seems a legitimate stipulation.²⁹ The important point that I hope strikes the reader is that even if he has nothing much to do, gaining knowledge of pointless truths doesn't seem to constitute something worth doing, let alone something that is finally valuable.

The defender of the Value of Achievements Thesis might claim that even here, that although knowledge is valuable as an achievement, other factors make it such that it turns out that all things considered this knowledge is not valuable. But this is a case of easy to gain knowledge; if sipping and counting are going to be enough to outweigh the value of knowledge then it seems that the value knowledge has simply as knowledge, the value basic to knowledge generally, rather than whatever different particular value may accrue to different instances of knowledge, is non-existent or perhaps risibly low.³⁰ The value that knowledge has as knowledge will at best be the pro tanto value of the least valuable knowledge, this will at best be a value common to all knowledge, one that is stripped of value contributing factors purely extrinsic to knowledge. I say "at best" as it is conceivable that even the pro tanto value of the least valuable knowledge cannot be stripped of purely extrinsic value conferring factors and so would be an overestimate of the common denominator of the value of knowledge.

Alternatively the defender of the Value of Achievement Thesis might argue that the story of the way Pierre gains knowledge is an irrelevant distraction. The frame the defender may offer to support their position is one within which one considers whether, intuitively, having an extra instance of knowledge is valuable; imagining

²⁹ It might still be objected that if Pierre is counting between sups then he is missing out on doing something more valuable. For example, relaxing is valuable and he is missing out on that. I think the relaxing example doesn't work because it's not something that can be necessarily done at will. Suppose Pierre is waiting for Jean-Paul and he's worried that he might be waiting in the wrong café; he just can't relax. The opportunity cost objection might gain some traction if the case were made that granting that Pierre can sit at the café and count people between sups of coffee necessitates that he can do something else that's more valuable instead. However, it's not obvious what that opportunity cost would be and that it could do the work of convincing us that the value that knowledge has as an achievement is not risibly low or non-existent.

³⁰ One argument that might be advanced is that all knowledge, even Pierre's knowledge, has value in virtue of its possible usefulness. It's possible an evil-doer might threaten to kill Pierre's family unless he can tell her how many people passed between sups, in which case knowing would seem very valuable. This argument based on the possible usefulness of knowledge seems weak, however, when we consider that the converse outlandish scenario might alternatively be the case; the evil-doer might kill Pierre's family if he knows how many people were passing between sups.

two worlds, one in which S has n amount of knowledge and another in which S has n+1 amount of knowledge. The extra item of knowledge is the knowledge of how many people passed between sups one and two of the coffee. The intuition that the knowledge in the Pierre case has no value may now seem a little weaker.

Ultimately, however, it still seems that knowing that 5 people passed between sups one and two of coffee either has no value or perhaps miniscule value. If knowledge has, as the Pierre case suggests, no value, or perhaps some miniscule value, then it is hard to see why knowledge should be protected and/or promoted.³¹

§1.5.2 A Possible Defence

As pointed out in §1.5.0, the argument of that section, which was expanded upon in the subsequent section, depends on an implicit assumption; that P3 of Greco's argument, achievements are finally valuable, implies that individual instances of achievements are finally valuable, and, given the other premises of the argument, that Greco is committed to defending the claim that each individual instance of knowledge, an achievement, has pro tanto final value. It seems natural to draw from Greco's argument that if there is a particular achievement, then, according to the argument, the particular achievement will have final value. Although I take such a reading to be natural, this reading may be denied. In making such a denial, the effect may be that cases such as the Pierre case are rendered irrelevant and the premises as originally stated, though not as they were initially explained, are defensible.

A way to deny the implicit assumption and so protect the argument as it stands might be to claim that the argument doesn't involve a universal claim about the value of knowledge, but rather a general claim, as we saw Fricker makes, about the value of knowledge. One needn't claim that every individual instance is valuable, rather one may say that knowledge, along with other goods that are achievements, are, owing to their nature generally valuable in the way described in the argument. In

³¹ Similarly, such a level of value would by itself provide sparse or no theoretical motivation for the focus of attention that knowledge receives in epistemology.

other words, claiming that it is in the nature of a good to have a certain value needn't imply that each instance of that good has that value. Analogously, it may be claimed that it is the nature of a tiger to be fierce; in saying that, however, the possibility of a non-fierce tiger isn't excluded. Rather, it's just that such a tiger, given the nature of tigers, would be exceptional as a tiger. (Pritchard: 2010b).³²

It seems, however, that the conclusion of the argument, C2, should apply to anything that meets the criteria set out in the premises of the argument. Greco can't deny that Pierre has knowledge and denying any of the other premises would require a major revision of his argument. The argument set out is Pritchard's representation of his position and so it wouldn't be costly for him to argue that there is some slippage from "for any" claims, which surely P1 and P2 are, to "for some" claims he might say that P3 and C2 should be.³³ The cost that would come from such a move, however, is that the value of knowledge would no longer simply fall out of the nature of knowledge. We would require some explanation as to why some instances of knowledge are finally valuable and others are not.

§1.6.0 Conclusion

In this chapter, having set out the broader context of the debate, I argue that Greco's account of the value of knowledge as it stands is both too narrow and too broad. I argue that it's too narrow in that on Greco's account the case is made that the distinctive value of knowledge is owing to knowledge being an achievement and that which falls short of knowledge not being so. I argue, however, that intuitively an epistemically virtuous belief is also an achievement; it does seem the kind of thing that is a cognitive success that is because of ability. I argue that Greco's account is too broad in that on a reading of his account that seems natural it turns out that any instance of knowledge has at least pro tanto final value. I construct a case and argue that although as the case is described the agent therein has knowledge, it is

³² Pritchard (2010b) credits John Turri with the example.

³³ This would also require changing C2 so as to make the revised argument valid. The new C2 would have to be something like "if some of the achievements referred to in P3 are knowledge, then some knowledge has final value." This revision would leave the burden of showing that the antecedent is the case still to be addressed.

implausible that the agent's knowledge has pro tanto final value. In the following chapter I turn my attention to Greco's account of the nature of knowledge, out of which his account of the value of knowledge falls.

Chapter II: Greco's Account of the Nature of Knowledge

§2.0.0 Introduction

In this chapter I examine Greco's handling of Barney-type cases. I argue that his handling of the Barney case is such as to have the consequence that his handling of Roddy type cases is unsatisfactory. More specifically, I argue that the justification Greco gives for denying that abilities are exercised in the Barney case is equally relevant in Roddy type cases. On this basis I make the case that either Greco must revise his grounds for treating the Barney case as he does or he must revise his treatment of Roddy type cases.

Next I outline Kallestrup and Pritchard's Epistemic Twin Earth case. The case is aimed at showing that an independent anti-luck condition is necessary for knowledge. I consider how Greco might respond to the case and whether such a response is successful.

Having completed my examination of the charge that Greco's account is too weak, I turn my assessment to the charge that Greco's account is too strong. I assess how Greco's view, that knowledge is true belief because of cognitive ability, fares in accounting for an apparently typical case of testimonial knowledge, the Jenny case. In assessing Greco's view on this point I consider criticism from Pritchard, and Greco's response to that criticism. I argue that the Careless Math Student case, a case that Greco cites as providing impetus for the move to an agent reliabilist account of knowledge, can be revamped in such a way so as to be the same as the testimony case, as described by Pritchard, in terms of the kinds of contribution made by the putative knower in the testimony case. The result, I argue, is that there is pressure on

Greco either to give up on the possibility of testimonial knowledge in cases currently thought of as typical or to defend a counterintuitive knowledge ascription with regard to the revamped Careless Math Student case.

Having I believe shown there to be significant marks against Greco's account of the nature of knowledge in this chapter, I assess whether Pritchard's competing account of the nature of knowledge, an anti-luck virtue epistemic account, should be considered a superior alternative. I examine how Pritchard's account handles the Barney case and the Jenny case, which are argued to be particularly problematic for Greco's view. I next turn to Greco's charge that Pritchard's anti-luck condition is ad hoc and Pritchard's response that the condition is proper given that it reflects one of our master intuitions about knowledge. Because the anti-luck condition rules out knowledge involving particular, problematic types of luck rather than luck generally, the claim that the condition is ad hoc is difficult to dislodge and counts against anti-luck virtue epistemology.

§2.1.0 Barney Cases on Greco's Account: An Introduction

In this section I set out Pritchard's (2010) charge that Greco's robust virtue epistemology problematically rules in Barney-type cases as being cases of knowledge. I examine Greco's response that for one thing abilities are environment-relative, but I argue that this response is unsatisfactory on the basis that it is at odds with Greco's treatment of the Roddy case, a standard Gettier-type case, in which he grants that there is an exercise of ability, albeit not one that explains the truth of the belief of the protagonist. The reason for finding this response unsatisfactory is, I argue, that there is no salient feature of either type of case such that that feature provides grounds to grant that there is an exercising of ability in one type of case but not the other. To illustrate my point I describe two cases which are similar to one another and have relevant analogies to the Roddy case and the Barney case respectively. I consider how Greco might respond to the argument. My purpose in this section is to show that as things stand Greco's account of the nature of knowledge exhibits a weakness in its handling of the Barney case when considered in

conjunction with its handling of the Roddy case. While what I write here touches on other issues, such as the kind of account of abilities we should have, I focus my attention on doing what has been described above.

§2.1.1 The Barney Case

First, let us turn to the Barney case as described by Pritchard (2012: 251):³⁴

Using his reliable perceptual faculties, Barney non-inferentially forms a true belief that the object in front of him is a barn. Barney is indeed looking at a barn. Unbeknownst to Barney, however, he is in an epistemically unfriendly environment when it comes to making observations of this sort, since most objects that look like barns in these parts are in fact barn façades.

What's of note in this case is that, unlike in standard Gettier-type cases, there appears to be a true belief because of the exercise of cognitive ability.³⁵ There is no so-called lucky intervention involved which results in the agent's belief turning out to be true. Rather, luck is said to be present in that had the agent come to believe on the same basis in modally close cases, then the agent would have gained a false belief.³⁶ For Pritchard, while Barney-type cases aren't cases of knowledge, they remain cases of achievement. In support of this Pritchard (2010b) writes that we would describe an archer, Archie, hitting his intended target because of ability as an achievement even if environmental factors analogous to those present in the Barney case were such that in modally close cases he would have missed. Such environmental factors could for example be that Archie just happened to take aim at the one target that didn't have a forcefield around it.

³⁴ The original barn façade case first appeared in a paper by Alvin Goldman (1976). Goldman credits the example to Carl Ginet.

³⁵ It should be noted that I am here representing Greco's account as being one on which an exercise of ability rather than ability is necessary for knowledge. It seems perfectly plausible that there may be occasions on which an agent does have the relevant ability which, if exercised, has the potential to gain the agent knowledge, but which, on those occasions, is not in fact exercised.

³⁶ Pritchard (2010: 36) describes Barney-type cases as involving environmental epistemic luck.

Greco holds a robust virtue epistemic account of knowledge whereby knowledge is true belief because of cognitive ability.³⁷ Pritchard (2010: 35-36) argues that Greco's robust virtue epistemology problematically rules in Barney-type cases as being cases of knowledge. More specifically, he argues that Greco's criterion for knowledge is met, there is a true belief because of cognitive ability; and yet the standard intuition among epistemologists is that protagonists in Barney-type cases don't know, rather such a protagonist has a justified but lucky true belief.

§2.1.2 Greco's handling of the Barney Case

An early response from Greco (2009b: 21) to the Barney case, and implicitly cases of that type generally, is to argue that no such environment-relative cognitive ability is exercised.³⁸ For Greco (2010: 10) an environment-relative ability to Φ implies a reliable environment-relative ability to Φ . Cognitive ability plays the same kind of role as justification does in more traditional epistemic accounts, so incommodiously for Greco, it seems that he might have to say, if using more traditional epistemological language, that Barney-type cases are not cases of justified true belief but rather are cases of mere true belief.

But what does Greco mean by an environment-relative ability? Greco and Turri (2011) write that on the former's account abilities are always environment-relative. We can't exercise our abilities in just any conditions. Greco gives the example of Tiger Woods's ability to sink putts being conditional on the environment being a certain way. We wouldn't say he has the ability to sink putts in hurricane-force winds. Elsewhere Greco (2010) claims that abilities are condition-relative as well as environment-relative. For example, whether a baseball player has the ability to bat

³⁷ As mentioned previously, to say that an account is a robust virtue epistemic account is just to say that it is an account according to which knowledge is **exclusively** accounted for as a true belief brought about by a virtue or an ability condition.

³⁸ It seems fair to say that this is still the general thrust of his response, but, as we shall see, he argues that further description of Barney-type cases are required.

well will depend not just on things like the weather, but also conditions such as him not having sand in his eyes.³⁹

Greco (2010: 76-77) foresees a species of the generality problem hampering assessment of whether there is an exercise of ability in the Barney case. If we ask whether Barney has an ability relative to an environment, with the environment fixed at one particular level of generality, that of Barn Façade County, then the answer is no. If we ask whether Barney has an ability relative to different levels of generality, for example the particular farm upon which the one real barn stands and no barn façades stand or across the globe upon which the number of real barns vastly outnumber the number of barn façades, then the answer is yes. This consideration raises the question of how we are to fix the appropriate level of generality of the environment to which we must judge whether an ability is relative.

Greco (2010: 78) claims that the level of generality of an environment, when determining whether or not an ability is present, should be fixed “according to the interests and purposes of relevant practical reasoning.” By way of illustration, Greco offers an example of how the level of generality can be fixed by a “practical reasoning context”. Suppose S says that R “has the ability to hit baseballs”; what is being claimed about R’s ability will differ depending on whether R is a baseball executive discussing whether to trade R or S is a Little League coach discussing what role to give a new seven year-old player; the conditions and environments that the abilities of each player are relative to can be expected to differ enormously. (Greco 2010: 78). The Barney case, Greco (2010: 79) writes, is underdescribed in that it is not clear what practical interests are in play, and, as discussed, we need to know that in order to be in a position to determine the generality of the environment to which Barney may or may not have a relative ability.⁴⁰ So without knowing this we are not

³⁹ It seems more appropriate just to talk about conditions; after all, surely environmental factors relevant to whether an ability can be exercised are simply conditions of a kind. Greco (2010:77) writes that the two may overlap but that for his purposes “environment” should be thought of as “sets of relatively stable circumstances and “conditions” as sets of shifting circumstances within an environment.”

⁴⁰ Greco, drawing on Edward Craig’s (1990) *Knowledge in the State of Nature*, also claims that “the concept of knowledge is used to flag good information and sources of information for use in practical reasoning.” (Greco, 2010: 78). It’s not clear to me, but it looks like he is saying that knowledge

in a position to judge whether Barney can be said to be exercising the relevant environment-relative ability, and ultimately whether he knows.⁴¹ Greco (2010: 80) does offer a further description of the Barney case in which he adds that Barney is in the area to calculate property taxes owing and that barns and barn façades are liable to different rates of tax. Greco's analysis is that Barney does not possess the relevant environment-relative ability, and therefore there is no exercise of the relevant cognitive ability.

§2.1.3 Greco's Handling of the Roddy Case

What is crucial on Greco's (2010: 101) account of knowledge, and for his knowledge as achievement claim, is that a belief is true because of ability rather than there simply being a coincidence of the exercise of ability and true belief. So there must be an environment-relative ability that is exercised and the exercising of that ability must explain why the belief that is gained is true.

To see a motivation for Greco's claim that knowledge is true belief **because** of cognitive ability, consider the Roddy case from Chisholm (1977: 105), which has been adapted by Pritchard (2012: 251):

Using his reliable perceptual faculties, Roddy non-inferentially forms a belief that there is a sheep in the field before him. His belief is also true. Unbeknownst to Roddy, however, the truth of his belief is completely unconnected to the manner in which he acquired this belief since the object he is looking at in the field is not a sheep at all, but rather a sheep-shaped object which is obscuring from view the real sheep hidden behind.

Intuitively Roddy doesn't know. Although a justified true belief is present, or, in Greco's terminology, a true belief and the exercise of cognitive ability are present, the true belief is not because of the exercise of cognitive ability and therefore it is not

conforms to the same sort of pattern as ability does, in that whether someone has a cognitive ability that is such that knowledge may be gained because of its exercising will depend on practical considerations. And when it comes to cognitive abilities specifically in relation to knowledge we can identify what some of those considerations are; for example, flagging good information and sources of information.

⁴¹ Of course having an environment-relative ability does not imply the exercising of that ability, although the exercising of an environment-relative ability does imply the possession of that ability.

a case of knowledge. Roddy seems not to know precisely because the truth of his belief lacks the appropriate causal connection to the exercise of his cognitive ability. It's a theoretical plus for Greco's (2009b: 19) account that its diagnosis of why Roddy doesn't know appears exactly right.⁴² Furthermore, this is how Greco (2009b: 19-21) purports to handle standard Gettier-type cases generally.

§2.1.4 The Barney and the Roddy Cases Taken Together

As we have seen, Greco's treatment of Barney-type cases differs from his treatment of the Roddy case and standard Gettier-type cases generally. The treatment of the two types of cases differs in that Greco claims there to be an exercise of cognitive ability in standard Gettier-type cases and not in Barney-type cases. In his treatment of the Roddy case and standard Gettier-type cases generally, one of the most appealing aspects of his robust virtue epistemic account is on display. There seems to be no principled reason, however, for Greco to claim there to be an exercise of cognitive ability in the Roddy case, but not in Barney-type cases. To put the issue into sharper relief, why think that Roddy has an environment-relative ability in his case, according to which there is a sheep shaped object in his vicinity, but that Barney doesn't have an environment-relative ability in his case, according to which there are Barn facades in his vicinity?

To see what's at issue I consider two cases intended to be analogous to the Roddy case and Barney case respectively. First, however, some background details of the two cases need to be spelt out.

Tony lives in part of the Amazon Rainforest. Together with Tony live hundreds of x-type birds. Sometimes Tony goes up to a tree, listens carefully and identifies the presence of an x-type bird in the tree by its singing. In fact when there is an x-type bird in the tree and it is singing and Tony is at the foot of the tree listening

⁴² The later Greco would presumably want to know what practical interests are in play in the Roddy case in order to determine whether Roddy is indeed exercising an environment-relative ability. There seems something costly about always having to identify practical interests in order to be in a position to determine whether an agent is justified in their belief, but I won't explore this thought any further here.

carefully, Tony can reliably identify an x-type bird as being in the tree. Let's say that Tony is his tribe's tax inspector. "Householders" are liable to varying rates of tax depending on the number of trees on their properties that have at least some x-type birds in them when Tony is doing his rounds.⁴³

For the first time ever some y-type birds, a cousin species of x-type birds, have ended up in Tony's part of the rainforest. More precisely, five y-type birds have gotten lost during the annual migration of y-type birds and the five spend a day in Tony's part of the rainforest before moving on.

Lost Birds, case one:

On one occasion that day, Tony ventures up to the foot of a tree, listens carefully, hears singing, and forms the true belief that there is an x-type bird in the tree. However, the singing that Tony hears is in fact that of a y-type bird, it just so happens that there is also an x-type bird in the same tree.

As in the Roddy case, in this Lost Birds case there is a true belief and there appears to be an exercise of cognitive ability. And, as to the Roddy case, to this Lost Birds case there is the response intuition that the protagonist doesn't know.

Lost Birds, case two:

On a separate occasion that day, Tony ventures up to the foot of a tree, listens carefully, hears singing, and forms the true belief that there is an x-type bird in the tree. The singing that Tony hears is in fact that of an x-type bird, it just so happens that in each of the five surrounding trees there is a y-type bird singing.

As in the Barney case, in the second Lost Birds case there is a true belief. And, as to the Barney case, to this case there is the response intuition that the protagonist doesn't know. My hope is that the cases taken together also show that intuitively it would be very odd to think that Tony has an environment-relative ability in one but not the other, or that Tony is exercising an environment-relative ability in one but

⁴³ I realise the romantic image of a man in the forest taking time to listen to birdsongs has been ruined by allowing talk of property and taxes to intrude, but I do so in order to meet Greco's requirement of describing the practical interests in play and this set of practical interests mirrors a set, described previously, with which Greco himself fills out the Barney case.

not the other; I only need the latter to suggest a problem with Greco's handling of the cases.

It's possible, however, to press the point further in a way that reveals something interesting about Roddy type cases at least and, if correct, implies that a defence of Greco's current way of handling the two cases differently is hopeless. Suppose there were just one x-type bird and one y-type bird in Tony's vicinity and they both happened to be in the same tree. Let's say that the hundreds of other x-type birds and the four other y-type birds are off at the other end of Tony's part of the rainforest. Now if the y-type bird but not the x-type bird is singing then it's a Roddy type case loosely construed. If it's the x-type bird but not the y-type bird that is singing, then it is a Barney type case. The way in which Tony forms his belief is such that he could have easily been wrong. If the foregoing is accepted, then Greco's claim that there is an exercise of ability in Roddy type cases but not in Barney-type cases is untenable. After all now a loosely construed Roddy type case, on Greco's analysis, should be seen as like the Barney type case, except that in the Roddy case there is no causal connection between the "unreliable cognitive faculty"/reliable ability to the truth of the belief in the Roddy case whereas there is in the Barney case.

§2.1.5 Greco: Possible Responses

How might Greco respond? Rather than alter how he handles the two cases, Greco might argue that at least one of the cases as described should itself be altered. Noting that Barney-type cases require that the basis of belief of the protagonist must be such that the protagonist could have easily been wrong, Greco might argue that Roddy type cases should be understood as cases in which it is necessarily such that the basis of the protagonist's belief is so that the belief could have not easily not been appropriately connected to the truth, even if in the actual case it isn't. In other words, Roddy cases should be such that despite the protagonist's belief not being appropriately connected to the truth in the actual world, nearby possible worlds should be such that coming to believe on the same basis the protagonist does have a belief that is appropriately connected.

It's hard to see how this could be right in the case as described. Presumably in nearby worlds in which the sheep isn't, from Roddy's line of vision, standing directly behind the sheep shaped object but rather is standing a little to the side of the sheep shaped object, Roddy will form the false belief that there are two sheep in the field. It seems that the only case in which we might say that Roddy forms a belief that there is a sheep in the field that is a true belief because of Roddy's ability will precisely be an obvious Barney type case; that is, a case in which the sheep just so happens to be, from Roddy's line of vision, standing directly in front of the sheep shaped object, but where Roddy's basis for belief is such that he could have easily formed a false belief. When we think of it like this, the Roddy case, as originally described, seems to imply that his belief was formed in the case in such a way that he could have easily been wrong. Of course in general forming beliefs about the number of sheep in a field on the basis of perception is a good way of forming beliefs but then the same can be said about forming beliefs about the presence of barns.

It looks like what I've written about the Roddy case is applicable to other standard Gettier cases.⁴⁴ Consider the Edmund case as described by Pritchard (2012: 250) which is closely based on an original Gettier case:

Edmund forms a belief that Jones owns a Ford on excellent grounds. He then validly infers that either Jones owns a Ford or Smith is in Barcelona, and accordingly forms a belief in this entailed proposition solely on the basis of his grounds for believing the entailing proposition and the relevant deduction. As it happens, the entailing proposition is false; the entailed proposition, however, is true since it just so happens (and unbeknownst to Jones) that Smith is in Barcelona.

Again, once the protagonist's belief being true is not explained by the protagonist's basis for belief in the particular case, then it will be vulnerable to the same charge. Although Edmund's basis for belief may generally be a good one, in this case it's

⁴⁴ Interestingly, it appears that what I've written does not apply to lottery type cases. Whether one should consider lottery type cases as cases of justified true belief is, however, debatable. See Smith (2010).

not a good basis for believing the target proposition. It seems that Edmund's belief is only true because of the presence of intervening luck.

An interesting point to be taken from the cases is that if there is a requirement to assess whether an environment-relative ability is present in a particular case, then standard Gettier-type cases are going to be just as likely as receiving a negative verdict as Barney-type cases. What's worse for Greco's position, such a way of looking at the cases makes standard Gettier-type cases appear to be epistemically worse than Barney-type cases, though perhaps this is fitting given that Barney-type cases aren't quite so widely rejected as a case of knowledge. The claim then that there is an exercise of ability in the first type of case but not in the latter is all the harder to defend.

The reason Greco generally has to resist saying that there is an exercise of ability in Barney-type cases, is to avoid the counterintuitive result that such cases are cases of knowledge. However, if one has a different analysis of knowledge, one that does not claim that knowledge is true belief because of the exercise of ability, then one can accept that there is true belief because of the exercise of ability in such cases without that committing one to saying that such cases are cases of knowledge. For example, if one were to claim that for an analysis of knowledge both an ability condition and an anti-luck condition are needed, rather than just an ability condition, then one could say that there is an exercise of ability in the Barney case without being committed to saying that it is a case of knowledge.

But isn't what Greco writes about abilities being environment-relative plausible? While it is plausible to deny that, say, Tiger Woods can exercise a reliable ability to sink putts in hurricane-force winds, the claim that ability implies reliability seems less intuitive. Plausibly a long distance Olympic gold medalist's career best time is down to her ability; after all it's not as if non-athletes could luckily replicate her time, and yet it's precisely the kind of instance of an exercise of ability that is not reliable. It's not reliable in that even though we think the time is down to her ability we don't necessarily think that her ability is such that she can reliably reproduce it. While

what a good alternative account of ability would look like requires spelling out, my point here is that, apart from the difficulties it leads Greco into outlined in this section, there are independent, intuitive grounds to be wary of Greco's account.⁴⁵

§2.1.6 Epistemic Twin Earth

A more refined version of Pritchard's (2010) original objection to Greco's position on the Barney case comes from Kallestrup and Pritchard (2011). The claim is that Greco's ability condition is not sufficient to exclude all cases of knowledge undermining luck. Even if an ability is exercised, the authors claim that that is not necessarily sufficient to rule out the agent's basis for belief in a particular case being such that she could have easily been wrong. While they accept that abilities are relative to environments, they deny that they are relative to modal environments in the way that they would need to be for robust virtue epistemology to be true.

Kallestrup and Pritchard ask us to imagine two cases. The agents in the two cases are virtue theoretic duplicates, each has the same abilities. In each case, in each of their respective worlds, the agent believes truly that there is water before her because of the exercise of their perceptual abilities. Thus the local environments are held fixed. The global environment is such that forming beliefs about what is and isn't water on the basis of perceptual ability reliably yields correct results. Thus the global environments are held fixed. Given all this we might think that both agents know.

The twist is that what is modally close in one of the worlds is not what is modally close in the other. In the first world, closeby possible worlds are such that the agent's basis for belief is such that she couldn't have easily been wrong. In the second world, closeby possible worlds are such that the agent's basis for belief is such that she could have easily been wrong. Intuitively, in the second case, despite the agent believing truly because of ability, the agent doesn't know. As the two agents are virtue theoretic duplicates, the thought is that the two cases taken together show that an ability condition alone isn't sufficient for knowledge.

⁴⁵ For alternative accounts of ability see Maier (2011).

In recent work Greco (2011: 227) has simply insisted that in Barn Façade cases, and presumably cases like epistemic twin earth, the agent lacks an environment-relative ability. This seems an unsatisfying response as the epistemic twin earth cases have been constructed so as involve virtue duplicates. Greco might claim that given the difference in nearby possible worlds in the epistemic twin earth case, the agent can't be said to be exercising an environment-relative ability, and so can't be a virtue duplicate of the agent in the other world. But such a response looks odd when we consider other cases of the exercising of abilities. Presumably we would think that Archie is exercising the same ability to hit bull's-eyes both in weather conditions in which the actual world and nearby worlds are normal and a world in which actual weather conditions are normal but in most nearby possible worlds a gust of wind would have blown the arrow off target.

However, Greco's defence does not rest on simply denying that there could be an exercising of an environment-relative ability alone. He also claims that judgements about whether an ability has been exercised depend on interests and purposes. Here the thought is that epistemic abilities are such that what we want from them is that they can discriminate between real barns and fake barns. But this is precisely different from the kind of interests and purposes that attach to hitting bull's-eyes in archery. The thought seems to be that what's going on in nearby possible worlds isn't significant when it comes to judging whether or not someone has the ability to hit bull's-eyes but that it does matter when it comes to the exercise of cognitive abilities.

This surely is too demanding. We wouldn't know much of anything if we not only had to have true beliefs because of the exercise of ability, but our abilities had to be such that they ensured or even reliably produced safe true beliefs. Consider the first world in Kallestrup and Pritchard's example in which nearby possible worlds are such that the agent's basis for belief is such that she couldn't have easily been wrong. It's not her basis for belief that makes it such that she couldn't easily been wrong. After all each of the agents in the two worlds described by Pritchard and Kallestrup are supposed to be virtue theoretic duplicates. If Greco wants to maintain this line of

defence he'll have to deny that the agent in the first world knows. The claim that cognitive abilities, the exercise of which are conducive to knowledge, must be such that they also produce safety looks to be setting the bar too high for knowledge.

Alternatively Greco may claim that true belief is not explained by ability in cases in which a belief is unsafe. Even if there is an exercise of ability and it produces a true belief, Greco might claim that what is salient in explaining the true belief is luck. Such a response seems unsatisfying. To say that Barney's true belief in the actual world is not down to the exercising of ability seems simply to privilege Greco's own account of knowledge. If someone were to ask why Barney has the true belief that there's a barn before him, it would seem natural to say that it is because he sees a barn before him.

Admittedly, the objection considered to Greco's account of the nature of knowledge is not a knock down argument against that account but rather demonstrates a mark against that account.

§2.2.0 Testimonial Knowledge on Greco's Account: The Dialectic

It will be useful in this section to trace the course of the dialectic with regard to how cases of testimonial knowledge are handled on Greco's account of the nature of knowledge. Pritchard (2010:41-42) argues that handling apparently typical cases of testimonial knowledge poses a problem for Greco's robust virtue epistemic account of the nature of knowledge. Consider the following case:

Our protagonist, whom we will call "Jenny", arrives at the train station in Chicago and, wishing to obtain directions to the Sears Tower, approaches the first adult passer-by that she sees. Suppose further that the person that she asks has first-hand knowledge of the area and gives her the directions that she requires. Intuitively, any true belief that Jenny forms on this basis would ordinarily be counted as knowledge. (Pritchard, 2010: 40).⁴⁶

⁴⁶ A couple of things about the case should be noted here. First, Jennifer Lackey (2007: 352) originally raised this case in the form of a counterexample to Greco's account, albeit drawing a different conclusion than the one Pritchard draws. Second, the Sears Tower is now officially known as the Willis Tower.

The thought is that the Jenny case poses a problem for the robust virtue epistemic account because Jenny does not seem to gain a true belief in the appropriate way, i.e. because of her ability, and as such her true belief doesn't seem creditable to her cognitive agency. Pritchard (2010: 40) distinguishes "of credit", what he claims can be said of Jenny with regard to her true belief, from "primarily creditable", what he claims Greco needs to be able to say of Jenny with regard to her true belief given his agent reliabilist account of the nature of knowledge.

Pritchard (2010: 41) further describes the case in a way that seems natural. He writes that it is right to think that Jenny's true belief is of some credit to her. After all, having arrived in the city and not knowing where one of the city's famous sights is, asking someone how to get to that sight is the right kind of thing to do; and it's not as if she asked a young child, she asked an adult, and presumably not one who obviously looked like a tourist. Furthermore, Jenny wouldn't just believe the person she asked no matter what. Had the testifier given Jenny directions while say trying to suppress sniggers, then presumably Jenny wouldn't have believed him. As well as how the testimony is delivered, what testimony is delivered, or in other words the content of the testimony, may also influence Jenny's belief. Presumably if Jenny were told, upon asking how to get to the Sears Tower, to take the next train back to New York, then she wouldn't believe the testifier. If all this is right then it seems correct to say that Jenny's true belief is of credit to her.

But this "of credit" does not suffice. Pritchard (2010: 40) claims that what the robust virtue epistemologist needs when he writes that knowledge is creditable to an agent is that the true belief is primarily creditable to the cognitive agent. Kallestrup and Pritchard (2012: 87) write that saying that a true belief is primarily creditable to an agent means that "the most salient factor in her cognitive success is her cognitive ability." Plausibly, cases of testimonial knowledge are not cases in which the recipient of testimony has a true belief in this way. Even if there is vagueness about the exact degree required, it seems right to think that the need for primary credit is no more than Greco's account of the nature of knowledge implies. The characterisation

of knowledge is after all true belief that is because of cognitive ability, where for Greco the cognitive agent truly believing and exercising cognitive ability must be one and the same.

Greco's (2011: 228) response at first glance may appear to challenge the view that what fell into the "of credit" category is not sufficient to do the job. He claims that on his view he can grant that lots of important work may be done elsewhere while remaining in a position to say that someone knows because of their ability. He provides the following case, intended to be analogous to the Jenny case, in support of this claim:

Ted receives a brilliant, almost impossible pass, and then scores an easy goal as a result. In the case we are imagining, it is the athletic abilities of the passer that stand out. The pass was brilliant, its reception easy. (Greco, 2011: 228).

The idea here is that just as it may be easy to score a goal following a good pass, when one receives good testimony it may be relatively easy to gain a true belief. In both the football case and the testimonial knowledge case, even if good work elsewhere makes success easy, success, scoring a goal or gaining a true belief, may still be because of ability.

Recently Kallestrup and Pritchard (2012: 90) have explicitly distinguished the Jenny case from one in which the assumptions discussed about Jenny's contribution to her gaining knowledge are explicitly built in. I'll refer to the latter as the Jenny* case. They claim that retreating to the view that Jenny's true belief need only be of credit to her cognitive ability will "come at the expense of preventing robust virtue epistemology from offering their distinctive response to the Gettier problem"; robust virtue epistemologists will no longer be in a position to offer the diagnosis they do of cases like the Roddy case, which is one of the most attractive features of the robust virtue epistemic view. (Kallestrup and Pritchard, 2012: 90).

But this claim looks mistaken. Greco's claim is that the **truth** of an agent's belief is creditable to his cognitive agency in cases of knowledge. In the Roddy case the truth

of the agent's belief is precisely not down to the agent's ability. In fact the agent's belief that there is a sheep in the field is unconnected with the fact that there is a sheep in the field. Kallestrup and Pritchard could develop how the case might be thought to be analogous. They might say that Roddy isn't looking at a house and thinking that there are sheep in there; he is looking in the kind of place that is such that if there are sheep around, then that's often where sheep are. But an important disanalogy with the Jenny case is that the Roddy case, along with standard Gettier cases generally, is naturally described as a case whereby what brought about the protagonist's belief, seeing the sheep shaped object, is unconnected from the truth of his belief.

There is an independent reason why we should reject the strategy suggested by Greco's Ted case. The claim that knowledge is true belief that is because of ability paints a certain picture of knowledge, one in which it is an agent's own abilities that explains the agent's having the particular true belief that she has and why that agent is creditable for her true belief. It is intuitive to think that it is Jenny's informant, rather than Jenny's cognitive ability, that is salient in explaining her true belief of how to get to the tower and so is the one who is primarily creditable for her true belief. As such, the picture of knowledge Greco presents jars somewhat with an ascription of knowledge in the Jenny case. A related cause for unease with Greco's account follows from consideration of just how much of our knowledge is testimonial knowledge, assuming we think that there is something unsatisfactory about an account which is awkward in its handling of a central type of case rather than a peripheral type of case.⁴⁷

The Ted case, however, is perhaps somewhat misleading. Greco's treatment of testimonial knowledge can be understood in a different way, one which makes more sense of Greco's position that knowledge is true belief because of cognitive ability. Consequently, true beliefs in cases of testimonial knowledge come out as being more than just of credit to the agent.

⁴⁷ I discuss the Far-Reaching Dependence Thesis in the next chapter.

Greco conceives of an ability as being necessarily reliable, therefore when we say that the recipient of testimony is exercising ability in the ways described, it is a reliable ability that Greco has in mind. Greco therefore has a reliability condition built into his ability condition, and he has no requirement for reliability to be built into his account beyond that. As a result, when Greco (2010: 81) writes that in order for an agent to gain testimonial knowledge that agent must be “a reliable receiver of testimony”, what he has in mind is that the agent must be exercising a cognitive ability that reliably produces true testimonial based beliefs.

Kallestrup and Pritchard (2012) also offer an argument against this treatment of testimony. Drawing on two cases which they describe, they argue that the kind of contributions that the recipient of testimony can make, which are made explicit in the Jenny* case, “cannot possibly ground testimonial knowledge”. (Kallestrup and Pritchard, 2012: 93). A premise in their argument is that often whether someone has testimonial knowledge will depend on background epistemic environmental conditions being met. More precisely, the claim is that often testimonial knowledge can only be gained if the relevant testimonial offering takes place in an environment in which there is “community-wide epistemic monitoring and policing” of testimony.⁴⁸ (Kallestrup and Pritchard, 2012: 94). The thought being that such practice will normally weed out the bad testifiers, so facilitating the gaining of testimonial knowledge.

The cases which Kallestrup and Pritchard (2012: 94-95) draw on in making their argument that the typical contributions of a testimonial recipient can’t normally ground testimonial knowledge run as follows:⁴⁹

Case 1:

H is in G where most testifiers are reliable testifiers. Nonetheless, H has an ability to discriminate between reliable and unreliable testifiers and this ability is regularly exercised in her acquisition of testimonial knowledge. Testimonial exchanges in G are regularly monitored and policed in the appropriate way. H acquires knowledge upon accepting the testimony of a reliable testifier.

⁴⁸ For more on this see Goldberg (2011), who Kallestrup and Pritchard cite.

⁴⁹ I’ve slightly adapted the two cases for presentational purposes.

Case 2:

H, unbeknownst to her, is no longer in G but is in fact in B. In B too most testifiers are reliable testifiers. Unlike B, however, testimonial exchanges are not monitored and policed in the appropriate way; in fact monitoring and policing is such that H is mostly exposed to unreliable speakers; let's say purposely so. Despite H exercising the same ability as in G, H has no idea that she is being mostly exposed to unreliable testifiers. Suppose, however, that a reliable testifier slips through the net and that H acquires a true belief upon hearing the testimony of that testifier.

A first thing to note is that H is supposed to be a virtue theoretic duplicate who is exercising the same ability in both cases. A second thing to note is that the Case 2 has the structure of a Barney type case. Kallestrup and Pritchard argue that H does not acquire knowledge in the second case. The reason they offer is that knowledge requires safety. In the first case the basis for H's belief is such that she couldn't have easily come to believe a falsehood; there is appropriate community-wide epistemic monitoring and policing. In contrast, in the second case, given the kind of monitoring and policing taking place, H's belief is unsafe. By believing on that basis she could have easily formed a false belief. Such a diagnosis, they claim, however, is not open to the robust virtue epistemologist who must explain the difference as being down to a difference in the abilities of the agent. The challenge for the robust virtue epistemologist arises from the two cases being presented as cases in which there is no difference as to the virtue of the protagonist, though intuitively one is a case of knowledge and the other isn't.

A defender of Greco's virtue epistemic account may deny that there is virtue duplication in the two cases, or may at least deny that in both cases the protagonist is in a position to exercise the appropriate cognitive abilities. Such a defender seems likely to repeat the same defence made against Barney-type cases; she will likely claim that abilities are environment-relative. She will however also face the same challenges to those defences.

§2.2.1 A New Challenge

My purpose in what follows is to be in a position to assess Greco's handling of the Jenny case by way of examining what is involved in the exercise of an ability in such a case. In order to do this I examine Greco's motivation for moving to an agent reliabilist view, including a case that can be developed so as to be structurally similar to the Jenny* case but which intuitively isn't a case of knowledge.

A motivation for the characterisation of knowledge as true beliefs because of reliable ability (agent reliabilism) and not simple process reliabilism is to avoid counterexamples to which the latter form of reliabilism is thought to be vulnerable. More specifically, Greco (2010) has argued that agent reliabilism avoids possible cases in which strange and fleeting processes are the cause of true beliefs reliably being produced. In such cases the thought is that we should accept that the relevant true beliefs are produced by a reliable process and yet we shouldn't accept that they are cases of knowledge. Greco (2010: 149) provides the following cases, the second case being of particular interest for my purposes:

For an example of a **strange cognitive process**, consider the case of the Serendipitous Brain Lesion. Suppose that S has a rare brain lesion, one effect of which is to reliably cause the true belief that one has a brain lesion. Even if the process is perfectly reliable, it seems wrong that one can come to have knowledge that one has a brain lesion on this basis.

For an example of a **fleeting process**, consider the case of the Careless Math Student. Suppose that S is taking a math test and adopts a correct algorithm for solving a problem. But suppose that S has no understanding that the algorithm is the correct one to use for this problem. Rather, S chooses it on a whim, but could just as well have chosen one that is incorrect. By hypothesis, the algorithm is the right one, and so using it to solve the problem constitutes a reliable process. It seems wrong to say that S thereby knows the answer to the problem, however.

Neither of the above mentioned cases are cases in which the agent gains true beliefs because of ability. Indeed Greco (2010: 150) claims that cognitive virtues or abilities are such that they just can't be strange or fleeting in the relevant sense. The moral Greco draws from such cases is that the sort of reliable processes that are knowledge

conducive are ones which are grounded in the agent's cognitive abilities.⁵⁰ Greco (2010: 150) holds virtue, of which ability is a species, to be such that it "must be stable across close possible worlds". An important point to emphasize here is that these cases and what Greco has to say about them clearly illustrate that Greco doesn't believe that an agent getting something reliably right logically implies that the agent is doing so because of an ability.

Now Greco's motivation for agent reliabilism should be clear. Having made the moves he's made, can he still describe apparently typical cases of testimonial knowledge as indeed cases of testimonial knowledge? To answer this question I think it will be instructive to look more closely at the Careless Math Student case.

First, let's be clear that by saying that "S has no understanding that the algorithm is the correct one to use for this problem", Greco should not be taken as suggesting here that some kind of internalist condition for knowledge is required, rather this is just to indicate that S stumbles across the algorithm. For the case to do the work it's intended to, i.e. show that simple process reliabilism provides the wrong analysis of knowledge and that an ability condition, with ability implying reliability, is needed, then it must be the case that S believes the answer, even if he has just chosen the algorithm on a whim; perhaps he thinks he's choosing according to a finely honed mathematical intuition but is in fact, as stated, choosing on a whim. Applying the algorithm, he believes the resulting answer he arrives at and that answer is arrived at in a reliable way – by using the correct algorithm. Despite S having a reliably formed belief as to the answer, he does not know the answer.

The Careless Math Student case can be revamped so as to be analogous to the Jenny* case, Pritchard's further description of the Jenny case, such that the contribution of the abilities of each to their respective beliefs is the same. First, even if the student chooses the algorithm on a whim, we can stipulate that he is somewhat discriminating; for example, he's not going to believe that just any algorithm will do, he'll be able to see if an algorithm is obviously suspect. Jenny is similarly

⁵⁰ One might challenge the claim that knowledge can't be gained from so called strange and fleeting processes. For such a challenge, see Goldman (2011).

discriminating; just as the student can spot an obviously suspect algorithm, so too Jenny has some idea what a suspect testifier looks like. Second, the student may give up on an algorithm on the basis of how it delivers the answer, say that when he sees what's involved in applying the algorithm he may in some cases no longer believe it's the right one, just as Jenny may not believe what the testifier says on the basis of how the testimony is delivered. Third, just as the student may discard an algorithm if it yields an obviously incorrect answer, so too Jenny may disregard testimony if it's obviously not correct.

One response to the revamped Careless Math Student case, given its symmetry with the Jenny* case and a concern to get the right intuitive result in the Jenny* case, might be to say that the student knows. Even if the student exercises like abilities as Jenny, however, there is still plenty of room for him to end up believing truly a particular algorithm and the result it yields, without us having the intuition that he knows. His abilities only limit the range of what he might believe, they don't explain why he has a true belief. The same goes for Jenny.⁵¹ When put like this, in neither the Jenny* case nor the revamped Careless Math Student case, is there a true belief because of ability.⁵²

Jenny will have to be doing more than the above described if she is to be said to know on Greco's account. As discussed, Greco holds that for an agent to gain testimonial knowledge that agent must be a reliable receiver of testimony, and Greco is committed to cashing out reliability here in terms of cognitive ability. Agent reliabilism involves a restricted form of reliability, not all reliably formed beliefs are because of cognitive ability.

⁵¹ Think of a contestant on a quiz show answering a question with a multiple choice of answers who can rule out some answers on the basis of his ability. This ability to filter out and perhaps even have positive reasons to believe some of the answers needn't be sufficient to know. Say if he rules out three of six possible answers, on the basis of his cognitive ability he may have no grounds for choosing one particular answer from the remaining three. If he does pick one of those three, believes the answer to be right, and it happens to be right, we wouldn't say that he knows.

⁵² Notice that the same kinds of contributions could be made by the protagonist in the temp case without us saying that they are sufficient for knowledge.

One way of thinking about the case is to say that the way Jenny forms her belief is such that, if she does indeed know, she couldn't have easily been wrong. But how can we say that this is down to an ability on Jenny's part? If the testifier had have given her slightly different and false directions, then Jenny would have been none the wiser. Greco is left in a position of needing to say that if Jenny, or indeed any testimonial recipient, knows, then her abilities are sufficiently honed to say that it is her abilities that will ensure that she reliably gets true beliefs. To put it another way, her abilities to, say, determine which testifiers to believe and whether to continue believing on the basis of the delivery and the content of testimony, must be such to ensure that she reliably gets true testimonial based beliefs.⁵³

Given Greco's theoretical commitments, it can't **just** be the case that a testimonial recipient is in a "friendly" testimonial environment, say where testifiers are sincere and accurate, whereby true testimonial based beliefs are easy to come by. An environment favourable to gaining true beliefs doesn't necessarily imply a corresponding environment-relative ability to produce the true beliefs; just consider the Serendipitous Brain Lesion and Careless Math Student cases. What might be said in defence of Greco's position is that if the testimonial environment is a relatively friendly one, then exercising ability so as to be a reliable receiver of testimony may not be particularly onerous. Testimonial knowledge may be seen as similar to perceptual knowledge, where our perceptual environment appears to be so friendly that true beliefs are, so to speak, practically handed to us on a plate, but where, despite the friendliness of the environment, in cases of testimonial knowledge we can be said to be exercising ability. But whether our testimonial environment is friendly in this way is not obvious, a case would need to be made that it is so.

Whether an agent needs to be a reliable receiver of testimony on the basis of the exercising of her cognitive ability in order to have testimonial knowledge even if the environment is friendly in the way described is also open to question. An obvious

⁵³ Of course, just one such ability, say the ability to determine which testifiers to believe, might be sufficiently honed to make it the case that the recipient of testimony is on the basis of cognitive ability a reliable receiver of testimony. It seems more implausible, however, that a reliable receiver of testimony is just so because of the exercise of only one such ability rather than because of the exercise of an array of abilities.

worry is that of young children, who we think do gain testimonial knowledge though it seems counterintuitive to say that they are reliable receivers of testimony because of the exercising of their cognitive abilities.⁵⁴ Even in cases that don't involve children it's not obvious that there is such a requirement.

Consider the following case:

John believes two putative experts about health matters. His doctor tells him truly that he has a nut allergy and that he should avoid eating nuts. His new age health guru tells him falsely that he has a lactose allergy. John believes them both. In fact John forms lots of his beliefs about his health on the basis of what his new age health guru falsely says.

John, in his reception of medical testimony, does not exercise an ability such that he is a reliable recipient of testimony, although he does exercise the kind of abilities detailed in the Jenny* case. Yet intuitively when he believes what his doctor tells him, he gains knowledge.

§2.3.0 Pritchard's Alternative Account

In this section I examine a rival to Greco's own account. I consider whether Pritchard's competing account of the nature of knowledge, an anti-luck virtue epistemic account, is a more attractive alternative. I examine Pritchard's analysis of knowledge and how it handles the two cases, the Barney case and the Jenny case, that were argued to be particularly problematic for Greco's view. I next turn to Greco's criticism that the anti-luck condition is ad hoc and to Pritchard's response to the charge, that the condition isn't ad hoc when you see it as a recognition of one of our master intuitions about knowledge. I'll argue that this response is ultimately unsatisfying given that it is accepted that certain kinds of luck are benign with regard to knowledge and that as such the condition Pritchard proposes is not a straightforward recognition of a pre-theoretical intuition. In that the anti-luck condition only rules out particular problematic cases of luck, the claim that the condition is ad hoc appears difficult to dislodge against Pritchard's anti-luck virtue epistemology. I give consideration to a different diagnosis of our knowledge

⁵⁴ I consider testimonial knowledge in young children in detail in the next chapter.

intuitions that if correct would serve to support an account of the nature of knowledge that is similar in spirit to Pritchard's own account.

I believe that so far I've shown difficulties for Greco's claim that knowledge is true belief because of the exercise of ability. There are competing accounts of the nature of knowledge in the neighbourhood. One such account is Pritchard's anti-luck virtue epistemology. Robust virtue epistemology has been favoured over anti-luck virtue epistemology partly because it has been thought that robust virtue epistemology can adequately account for knowledge without need for a separate anti-luck condition.

Both the Barney and Jenny cases have been discussed as problem cases for Greco's robust virtue reliabilist account of knowledge. The thought is that the two cases taken together show Greco's account to be both too strong, it rules out seemingly typical cases of testimonial knowledge, and too weak, it does not rule out cases of environmental luck. An alternative to a virtue theoretic approach has been anti-luck epistemology, but Pritchard argues that such an account is prey to the kind of cases that motivated the virtue approach in the first place. Pritchard's diagnosis of the dialectic is that the two different approaches flounder because they fail to appropriately incorporate one of what Pritchard claims to be the two "master intuitions" about knowledge into their account. These intuitions are that knowledge is from ability and that knowledge is not lucky in certain respects. Once these two intuitions about knowledge are recognised, recognised as setting distinct bars for what can count as knowledge, and an account is articulated that reflects the importance of these two intuitions, then, Pritchard argues, we will be in a position to have a successful account of knowledge.

Pritchard's (2012: 273) own particular analysis of knowledge, informed by these two intuitions, is that:

“S knows that p if and only if S’s safe true belief that p is the product of her relevant cognitive abilities (such that her safe cognitive success is to a significant degree creditable to her cognitive agency).”⁵⁵

That S’s true belief must be safe reflects the demands of the anti-luck intuition, and that the true belief must be the product of S’s relevant cognitive abilities reflects the demands of the ability intuition. Given that a single ability condition isn’t burdened with having to do all the explanatory work as to how someone can be said to know, it’s unnecessary for Pritchard to require that knowledge be true belief because of ability.⁵⁶ Pritchard is in a position to offer a weakened alternative ability condition than the one that Greco offers.

Pritchard’s account comfortably deals with the Barney case. Barney doesn’t know because although Barney has a true belief that is to a significant degree creditable to his cognitive agency, Barney’s belief doesn’t satisfy the anti-luck condition. The way in which he formed his belief is such that he could have easily been wrong; Barney’s belief is unsafe. When it comes to the Jenny case, Pritchard’s account benefits from being in a position to avail of a weakened ability condition; Pritchard’s ability condition is more permissive in its treatment of seemingly typical cases of testimonial knowledge than Greco’s own ability condition. It looks like Pritchard’s account of the nature of knowledge will get better results than Greco’s account in terms of judging putative cases of testimonial knowledge that are intuitively cases of testimonial knowledge as indeed being cases of testimonial knowledge. A worry, however, is that while Pritchard’s account fares better than Greco’s account with regard to cases of testimonial knowledge, Pritchard’s formulation, that cognitive success be to a significant degree creditable to cognitive agency, is such as to still set the bar too high. If one thinks that a young child can gain testimonial knowledge then it doesn’t seem plausible to think that in doing so her cognitive success is to a significant degree creditable to her cognitive agency.

⁵⁵ In my discussion of Pritchard’s analysis I focus on the aspects that are of interest for the purposes of this chapter, rather than providing a general overview of his position.

⁵⁶ It’s possible to see Pritchard’s formulation as offering a **single albeit complex condition** for knowledge, but for pragmatic explanatory reasons in my discussion of the formulation I will refer to **conditions** which reflect components of Pritchard’s formulation.

In considering Pritchard's alternative account, albeit presumably without being privy to its most recent articulation, Greco (2011: 227) argues that it compares unfavourably to his own account. He writes that Pritchard's account draws on "independent ability and anti-luck conditions" and that the motivation for the two conditions "is to handle two different kinds of cases". (Greco, 2011: 227). He charges that they lack a unifying theoretical motivation. As such Greco reasserts the charge that the inclusion of an anti-luck condition is ad hoc. The second charge against Pritchard's account is that anti-luck virtue epistemology must adopt a revisionist approach to the question of knowledge's value, and this is considered to be a mark against the account.⁵⁷ Greco (2011: 230) claims his account, in contrast to Pritchard's account, to have theoretical elegance on its side in both respects.

An obvious response to the first charge for Pritchard is to say that it's not ad hoc to have an account that reflects central intuitions we have about the target of that account. It's not clear why having different conditions doing different jobs by itself should be regarded as in anyway problematic. Even on the tripartite account of knowledge the truth condition is ruling out some potential cases that the justification condition is not and vice-versa. A way of reading the ad hoc charge that looks far more troubling for Pritchard's account is to regard it as aimed at the inclusion of Pritchard's particular anti-luck condition. The thought is that the inclusion of the anti-luck condition is being justified on the basis of reflecting an anti-luck intuition, but Pritchard's anti-luck condition doesn't simply reflect an anti-luck intuition; Pritchard's anti-luck condition deals with a particular species of luck, after all not all luck is knowledge undermining. When we think of it like this, that really the anti-luck condition isn't simply the articulation of an intuition that knowledge can't be lucky but rather it is the articulation of a condition that knowledge can't be lucky in certain respects, then thinking of this as a pre-theoretical intuition looks much more questionable.

⁵⁷ I restrict myself to examining the first charge in this chapter. It should be noted however that Pritchard (2010: 87-88) acknowledges that the account he offers is revisionist.

The worry is that we have the intuition in certain types of cases, standard Gettier cases and Barney cases, that the subject described doesn't know. As a result we see that we need to adjust our account of knowledge so as to rule out these types of cases being counted as knowledge. Now an opponent may claim that what an anti-luck approach does is finds a feature of these types of cases and adopts as a condition that cases with that feature cannot be knowledge.⁵⁸ But the anti-luck approach hasn't identified a feature that forms its own kind; it's not as if any true belief with the feature of being lucky is ruled out as potential knowledge. As already mentioned some species of luck are knowledge undermining but some species aren't; epistemic environmental luck is an example of the former while evidential luck is an example of the latter.⁵⁹ When it's put like this the inclusion of an anti-luck condition in an account of the nature of knowledge does look ad hoc. When that pre-theoretical intuition needs to be so refined as to capture environmental and intervening luck, but not benign varieties of luck, i.e. ones that don't undermine knowledge, then the ad hoc charge becomes difficult to dislodge.⁶⁰

An alternative course that might be open to Pritchard's account is not to justify the inclusion of his particular anti-luck condition on the basis of it reflecting an anti-luck intuition, but rather on the basis of it reflecting the intuition that a true belief from ability must be safe in order for it to count as knowledge. This way of doing it wouldn't leave a mismatch between the intuition and the condition. Presumably thinking of there being an anti-luck intuition especially made sense prior to thinking of there being an ability intuition; it helped explain why knowledge can't be mere true belief. However, given the identification of an ability intuition, i.e. that in order

⁵⁸ The particular notion of safety that Pritchard defends is not simply one that if accepted as being a requirement for knowledge will rule out Gettier-type cases, but it is one that he argues is better than rival alternatives.

⁵⁹ Pritchard (2005: 136) characterises evidential luck as luck "that the agent acquires the evidence that she has in favour of her belief". He illustrates this kind of epistemic luck by way of a case described by Peter Unger (1968: 159), though not identified by Unger as illustrating this kind of luck, in which, to simplify somewhat, a man passes by his employer's door and, just as he is doing so, happens to overhear that he is to be fired. The thought is that although the protagonist acquires his evidence by luck, he just happened to overhear that he is to be fired, this kind of luck is not knowledge undermining.

⁶⁰ As has already been discussed there are significant marks against Greco's account. As things stand, if one account is to be favoured over another it looks like it will be on points rather than via some knockdown argument.

to know one's true belief must be in some relevant respect down to one's abilities, it looks as if the anti-luck intuition hasn't been joined by a wholly distinct intuition, but rather one that can be taken to capture some of what the anti-luck intuition was supposed to capture. Of course, Greco's agent reliabilism has no specific anti-luck condition to deal with standard Gettier cases. Now it can be argued that what's left-over which the ability intuition doesn't encompass when articulated is an intuition that in order for a true belief from ability to count as knowledge it must be safe, or more simply a safety intuition about knowledge.⁶¹

Of course revising one's claims about what our pre-theoretical intuitions are leaves one particularly vulnerable to the charge that one's revised claims are somehow the product of inappropriate theoretical spillover. The worry is that if our pre-theoretical intuitions can be easily misidentified, then the claim that they should guide our analyses starts to lose credibility.

Whatever about these concerns with Pritchard's account of the nature of knowledge, it's clear that the account does give us the right result in Barney-type cases and if the account can be shown to give us the right result in the case of testimonial knowledge in young children, then it is in a better position than Greco's account in terms of yielding results that match our response intuitions.

§2.4.0 Conclusion

In this chapter I discussed the Barney case and the Jenny case, the handling of each pose significant challenges to Greco's account of the nature of knowledge given his existing theoretical commitments. I argued that his treatment of the Barney case is problematic. More specifically, I argued that he is committed to saying that in the Barney case there is not an exercise of the relevant cognitive ability, but that that claim appears unsustainable if he is to maintain his current analysis of standard Gettier-type cases. Next I considered a separate challenge to his treatment of Barney-

⁶¹ Of course an argument would be required as to why Pritchard's particular account of safety best meets a safety intuition. As noted, Pritchard does argue for the superiority of his own account of safety over rival accounts.

type cases based on examples involving virtue theoretic duplicates. I then turned to how his account deals with cases of testimonial knowledge. In doing so I laid out the course of the dialectic starting with the challenge, posed by Pritchard, to making the case that in standard cases of testimonial knowledge the agent is primarily creditable for the truth of her belief. I detailed Greco's response and the further response of Kallestrup and Pritchard. I then argued that there is a tension between Greco's treatment of testimonial knowledge and his motivation for moving away from a process reliabilist account of knowledge to an agent reliabilist account of knowledge. Finally I examined Pritchard's anti-luck virtue epistemic account of knowledge as an alternative to Greco's robust virtue epistemic account of knowledge. In doing so I considered its theoretical advantages and disadvantages; one disadvantage being that the charge that the inclusion of the anti-luck is ad hoc. I discussed one potential way of responding to the charge.

Chapter III: Epistemic Grace

§3.0.0 Introduction

In this chapter I begin by following on from where I left off in Chapter II. Having in that chapter considered the challenges both Greco's and Pritchard's rival accounts of the nature of knowledge face, in this chapter I argue for how I think those challenges can be overcome, and in so doing I set out an alternative account of the nature of knowledge. In my attempt to overcome the two challenges outlined I seek to build on the foundations of the virtue theories of knowledge offered by Greco and Pritchard rather than embarking on a wholesale departure from those accounts.

The first challenge is to show that such an account yields the answer that plausible cases of testimonial knowledge are indeed cases of testimonial knowledge. Although on Pritchard's account the required contribution from the cognitive abilities of an agent is somewhat less demanding than the contribution required on Greco's account, as with Greco's account, however it problematically rules out some cases as being cases of testimonial knowledge, plausible cases of young children knowing on the basis of testimony. The second challenge is dealing satisfactorily with Barney-type cases. An account should yield the answer that Barney doesn't know. While Pritchard's account clearly does this, it comes at the price of doing so in a seemingly ad hoc way. Ideally an alternative account will yield the same answer to Barney-type cases, but do so in a way that is more convincingly motivated.

I argue that the challenge of articulating an account of knowledge on which plausible cases of testimonial knowledge come out as indeed being cases of testimonial knowledge can be overcome primarily by rejecting the claim that the potential knower is required to believe truly because of the exercise of a cognitive ability. Rather, I argue that what is required is that belief be because of the exercise of virtue.

I make my case in such a way as to try to reconcile the phenomenology of testimonial knowledge with a virtue epistemic account of knowledge.

I attempt to overcome the challenge posed by Barney-type cases initially based on a non-standard reading of Gettier cases according to which it's bad epistemic luck that is knowledge undermining. Building on this reading, I argue that thinking of knowledge as requiring what I call epistemic grace both yields the right results in standard Gettier-type cases and Barney-type cases, and does so in a well motivated way.

§3.0.1 Chapter Motivation

Apart from attempting to overcome the challenges discussed in Chapter II, this chapter has a broader significance in my thesis. Providing an alternative account of the nature of knowledge puts me in a position in Chapter IV to return to the criticisms articulated against Greco's account of the value of knowledge in Chapter I and attempt to provide a better account of the value of knowledge based on my own account of the nature of knowledge. Furthermore, my explanation of how testimonial knowledge is possible is such that it lays groundwork for understanding factors that are relevant to the quality of the epistemic environment.

§3.1.0 Testimony

If the moral from the Jenny case is that Greco's requirement, that an agent believes truly because of the exercise of ability, is too strong, if it's accepted that Pritchard's claim that the prospective knower must make a significant contribution from ability to their believing truly is likewise too strong, and if we want a solution that fits within the framework of virtue epistemology, then we need to articulate an account on which either simply less of a contribution is required of the agent or a more radical solution is found. The first strategy, requiring less of a contribution from the agent, looks doomed if we accept that Greco's account is in certain respects both too strong and too weak. It is the second strategy I adopt in accounting for testimonial knowledge and for the nature of knowledge more broadly. More specifically, I argue

that rather than requiring for knowledge that an agent have a belief that is true because of the exercise of a cognitive ability, what is required is that an agent have a belief that is because of the exercising of virtue. As noted previously, Greco (2009: 318) sees intellectual abilities as a species of intellectual virtue, though it's clear that he regards knowledge as requiring true belief because of ability. My strategy may, therefore, be thought of as defending a position according to which it is a broader category of contribution that is a requirement for knowledge.

Before arguing for my proposed way of overcoming the first challenge, I set out the context in which the debate about testimonial justification is taking place in order to help motivate the position I ultimately defend. In doing so I describe what we should want a theory of testimonial justification to do. This includes giving an account of testimonial knowledge in a way which accords with the corresponding phenomenology, and explaining how a testimonial belief may be a virtuous belief to differing degrees, while at the same time being in a position to explain testimonial knowledge in young children. I argue that while a credulist from the outset is well placed to account for how virtuous testimonial belief accords with the corresponding phenomenology, she is not well placed to explain how a testimonial belief may be a virtuous belief to differing degrees without having to forego an explanation as to how testimonial knowledge in young children is possible. A Humean, meanwhile, looks to be in a good position to explain how a testimonial belief may be a virtuous belief to differing degrees, but appears to be unable to do so in a way that accords with the relevant phenomenology. A promising avenue that appears to put the Humean in a position to do the latter, drawing on dual process theory, is such that it will hinder an explanation of how testimonial knowledge in young children is possible.

The position I ultimately defend is that it is the exercising of epistemic virtue that is a requirement for testimonial knowledge, and knowledge generally. I argue that young children have epistemically virtuous testimonial belief, if their belief is because of the exercise of virtue, where virtue is construed more broadly than mere cognitive ability. Furthermore I argue that whether or not an agent can rightly be said to be exercising virtue depends on factors specific to that agent. The important point is that

believing because of trust may in a particular situation be epistemically virtuous for one agent, say a young child, but the same belief on the same basis in the same situation may not be virtuous in another agent, say a police investigator.

In this half of my chapter my focus is on making an argument for a particular account of the testimonial justification required for knowledge. I build on this account of justification to provide an account of the nature of knowledge, an account that I argue can satisfactorily provide the right results for both plausible cases of testimonial knowledge and Barney-type cases.

§3.1.1 The Discourse on Testimonial Knowledge

I take testimony to be "... the assertion of a declarative sentence by a speaker to a hearer or to an audience." (Adler, 2010). I regard this as an approximation of what testimony is, but one that suffices for my purposes.⁶² I further take it that not only is testimonial knowledge possible but that we are dependent on testimony for much of our knowledge. Following Adler (2010), I describe this as the Far-Reaching Dependence Thesis.

At first glance, the mere idea that we can gain knowledge from testimony might seem implausible; let alone the claim that we are dependent on testimony for much of our knowledge. A Cartesian way of thinking about knowledge lends itself to this thought. To elaborate, it may be thought that if even our perceptual beliefs are vulnerable to sceptical worries, then our beliefs dependent on the assertions of others are all the more vulnerable to sceptical worries. Knowledge, for the Cartesian, is reason based conviction which there can be no alternative reason for doubting; as doubts are removed, certainty increases. (Lex Newman, 2010). On this way of thinking of knowledge it's very difficult to see how one can gain true beliefs with knowledge conducive justification from standard cases of testimony. If one attains

⁶² Nevertheless it should be noted that testimony needn't be in the form of spoken word and "audience" here shouldn't be thought of as restricted to hearers. Testimony can, for example, be given in written form and sign language, and, might, be given by bodily movements with communicative potential such as pointing, winking, etc., in the appropriate contexts as Lackey (2008: 25-26) argues.

knowledge, one's knowledge will be equipped to withstand any sceptical challenge; if one knows, one will be certain. (Lex Newman, 2010).⁶³

The approach commonly taken in contemporary epistemology to the examination of the nature of knowledge is notably distinct from the methodology of doubt employed by Descartes for the same purpose. First, fallibilistic accounts of knowledge are widely accepted within the field of epistemology, fallibilistic in that we can know *p* even if our belief that *p* could, in a certain sense, have been wrong but isn't. More specifically, the justification of our belief may be such that it is possible that we could be wrong. (Pritchard, 2006). Second, and relatedly, it's standard practice when examining the structure of knowledge in contemporary epistemology to assume that we do know more or less what we ordinarily take ourselves to know. Third, it's also standard practice to extrapolate from such cases of knowledge as well as hypothetical cases in which we would take ourselves to know, to build an account of the nature of knowledge.⁶⁴

This is not to say that there is a denial of the sceptical problem, rather responding to the sceptic is treated as an enterprise distinct from that of providing an account of the nature of knowledge. C. A. J. Coady (1992: 3) for example distinguishes what he calls positive epistemology from negative epistemology. The former investigates the structure of knowledge or the body of beliefs that can reasonably be thought of as knowledge, and leaves sceptical worries aside; while the latter is concerned with the theoretical problems raised by scepticism.⁶⁵ Greco (2010: 5), in like vein, distinguishes what he calls "the project of explanation", which seeks to explain

⁶³ For an early treatment of how testimonial belief fits in with certainty and what we would now call entitlement, see the writings of the much overlooked John Henry Newman (2007) in *An Essay in Aid of a Grammar of Assent*. Anthony Kenny (2007: 30), in his history of Western philosophy, writes that his is "... arguably the best treatment of the topics of belief and certainty between Hume and Wittgenstein."

⁶⁴ Relevant to this discussion is Chisholm's (1973) *The Problem of the Criterion*, in which it is argued to be problematic to build an account of knowledge both from what we take the extension of knowledge to be without already drawing on what we take the intension of knowledge to be and similarly to start from what we take the intension of knowledge to be without already drawing on what we take the extension of knowledge to be.

⁶⁵ Confusingly, "positive epistemology" and "negative epistemology" are used in a different sense in some of the literature. For example, an alternative use of "positive epistemology" is one that describes an epistemology that assumes the negation of sceptical premises and considers why the original premises should be denied. (Lycan, 2001: 44.)

“what knowledge is and how knowledge is possible”, and “the project of vindication”, which is the project “of showing that we have knowledge”.

Claiming that what follows falls within the realm of positive epistemology or the project of explanation doesn't quite yet allow us to set sceptical worries with regard to testimonial knowledge to one side. After all, it might be thought that much of what we take ourselves to know doesn't include, or at least not to any significant degree, testimonial beliefs in propositions. But given a little consideration this will surely seem mistaken.

My knowledge that Mount Everest is the highest mountain in the world, that smoking can cause cancer, that there exist other planets in our solar system, that the French Revolution occurred, are all based on the testimony of others.⁶⁶ Not only is much of my knowledge of the broader world and its past based on testimony, but much of my knowledge of important facts about my own life and the lives of those around me is also based on testimony; my knowledge, for example, of my date of birth, the occupations of numerous friends and family members, and that, say, two of my friends holidayed in Italy last year are each based on testimony. I take it that most of us are in a similar epistemic position and that therefore the much of what we take ourselves to know includes knowledge that comes from testimony. These examples suggest not only that much of what we take ourselves to know includes some testimonial knowledge but that the Far-Reaching Dependence Thesis is true.

So far I have acknowledged that there might be resistance to treating testimonial knowledge as possible. I have explained that I will follow the contemporary epistemological approach of setting sceptical worries to one side and assume that we do know much of what we take ourselves to know. I have also countered the worry that, even if the contemporary epistemological approach is accepted, it might be objected that testimonial knowledge does not account for a significant portion of what we ordinarily take ourselves to know.

⁶⁶ History as a subject is particularly dependent on testimony. Testimony seems to be the basic raw material or data of history as a discipline.

§3.1.2 Virtuous Testimonial Belief: A Problem

Over the course of the next few sections I articulate a response to the Jenny case and offer a way of explaining testimonial justification in young children in what I take to be typical cases. I do so by developing that response located in the broader testimonial debate with regard to testimonial justification. I set out to achieve this by considering the merits of two rival theories of testimonial justification on the basis of how they fare with regard to three challenges each of which I argue a good theory of testimonial justification should be able to overcome. The challenges, set out in question form, are as follows:

1. How can we account for testimonial justification in a way that accords with the phenomenology of testimonial reception?
2. What makes some testimonial beliefs more justified than other testimonial beliefs?
3. How can testimonial justification in typical cases in young children be explained?

Relevant to addressing the first challenge is what Adler (2010) calls the Uniformity Claim; in standard cases hearers regularly simply accept a testifier's assertions as true. I take this to be a claim about the phenomenology of testimonial reception as well as being an empirical claim; in other words I take the Uniformity Claim as also describing how testimonial reception is commonly experienced. A further claim that I take to be relevant to addressing the first question is that the testimony believed in such cases is ordinarily justified. The motivation for the latter claim is that ultimately we want an account of how testimonial justification accords with the phenomenology of testimonial reception in such a way that we can explain how the Far-Reaching Dependence Thesis is the case. If both claims are indeed true then this would suggest that testimony carries its own justification; or, to put it another way, the fact that something is testified to ordinarily makes believing what has been testified justified, rather than testimony being justified at the expense of some reasoning process, say inference, on the part of the believer. Just as something visually seeming to be the

case, ordinarily makes believing what seems to be the case justified, so to with testimony that something seeming the case ordinarily make believing that testimony justified.

It does however seem intuitive to think that one recipient of testimony may be justified in her testimonial belief; while a fellow recipient may be less justified in believing what is said. To see this, consider the following case:

Jana tells her mother, who is a wise old judge, and her young child a fantastic tale of how, in the space of a fortnight, she lost most of her fortune, say when the stocks she had invested in crashed, and then regained a roughly equivalent fortune, say when other stocks she had invested in soared. Although her tale is fantastic both the judge and the child believe her tale, the content of which is known by Jana and so is in fact true.

Even if they both believe the testimony, we should want an account of testimonial justification to allow for the possibility that the judge is more justified in her belief than the child is in his belief. Such a consideration suggests that the degree of justification enjoyed by a testimonial belief may vary from person to person.⁶⁷ Finally, we should want an account of testimonial justification to explain how testimonial based beliefs are justified in intuitive cases of testimonial knowledge. Such intuitive cases of testimonial knowledge, as discussed elsewhere, include testimonial knowledge in young children.

§3.1.3 The Credulist Response

The works of Thomas Reid and David Hume have inspired rival contemporary accounts of testimonial justification. The Reidian or credulist view is that the Almighty intended us to be social beings and so made us with a “propensity to speak the truth” and with a disposition to believe what we are told by others. (Reid, 1983: 94-95). Given such a propensity and disposition, testimony is likely to be believed and testimony believed is likely to be true, making testimonial belief justified.

⁶⁷ Similarly, we can imagine that the judge’s testimonial belief may be more justified in one situation and less so in one in which she has less experience and is less wise.

The contemporary credulist argument doesn't make reference to God, but claims that adherence to the norm of assertion makes testimonial claims justified. Accounts claiming to articulate the norm of assertion set out the conditions in which it is regarded as appropriate to assert something. One account of the norm of assertion, is that the speaker can properly assert *p*, only if the speaker knows *p*. (Adler, 2010, citing Unger, 1975, and Williamson, 2000). If this is right, then it looks like we're on our way to answering the question as to what makes testimonial belief justified in a way that fits with the Uniformity Claim.

If there is widespread adherence to the knowledge norm of assertion, then if someone asserts *p*, then an audience is justified in believing *p*. Now we look to be in a good position to explain the Far-Reaching Dependence Thesis while continuing to hold the Uniformity Claim. Simply believing what is said won't prevent the recipient of testimony from having a justified belief.

Now, however, we don't seem to be in a good position to answer the second question. Returning to the case of the judge and the child, if testimonial belief is justified because of the norm of assertion then it's not obvious how the judge might be more justified than the child. Relatedly, claiming that knowledge may be gained simply by believing what's testified seems to set the bar too low for knowledge. Elizabeth Fricker (1994: 126) puts the worry about accounts of testimonial justification based on such a claim starkly when she writes that they are "... an epistemic charter for the gullible and indiscriminating." Fricker (1994: 145) recommends that "... a hearer should always engage in some assessment of the speaker for trustworthiness. To believe what is asserted without doing so is to believe blindly, uncritically. This is gullibility." (Both quotes are cited in Adler, 2010).

The credulist may warn that requiring such an assessment of testimonial justification discords with the Uniformity Claim, and risks ultimately leaving us unable to explain the Far-Reaching Dependence Thesis. The credulist may claim that without foregoing an explanation of the Far-Reaching Dependence Thesis, the credulist can

adjust their account in an attempt to address the concerns raised. The credulist may do so by claiming that there should be “counterfactual sensitivity” to possible defeaters. (Adler, 2010). If I ask for directions to the Sears Tower and would have believed the testifier, had he, gleam in eye, unsuccessfully attempted to muffle sniggers while giving directions, then I would be insensitive to what such behaviour normally indicates and not justified in my testimonial belief.

The natural question to ask is whether this move, which claims that counterfactual sensitivity rather than assessment is required, does not too go too far and discords with the uniformity principle. Miranda Fricker (2007: 65-66), criticises this move on the basis that we are left with no explanation of how an agent would go from simply accepting testimony to, when relevant, critically attending to justification relevant aspects of the testimony. In order for there to be counterfactual sensitivity the worry is that there needs to be a kind of continuous assessment or monitoring; the thought being that only this can explain the possibility of noting a defeater and responding accordingly. But the credulist needn't be committed to such continuous assessment or monitoring in order for there to be counterfactual sensitivity. A person may notice that another person's eyebrow has been shaved without requiring of the first person that he always be engaged in monitoring the eyebrows of the other person.

The picture of testimonial justification that emerges from this move is of testimonial justification being more complex than the original credulist idea of testimony carrying its own justification. Rather the picture we get is that testimony enjoys default or prima facie justification but that believing testimony alone isn't sufficient for justification, at least for knowledge conducive justification; the agent's belief must have the outlined counterfactual sensitivity.

Even with these significant adjustments to the credulist account of testimonial justification, it's not obvious that we are in a position to adequately address the second challenge of accounting for differing degrees of justification of testimonial belief, although we can account for one person having testimonial justification from a piece of testimony and another not having testimonial justification from the same

piece of testimony. The addition of a simple counterfactual sensitivity condition does not seem sufficiently dynamic to capture the way in which we'd imagine there to be a wide range of factors of differing weights adding to and detracting from the degree of appropriate justification a belief in a piece of testimony enjoys. In other words, we're not better placed, at this point, to say how the wise old judge might be justified in her testimonial belief to the n^{th} degree while another recipient of testimony may be less justified in his belief in the same piece of testimony. More significantly, however, the counterfactual sensitivity move precludes a promising basis for explaining testimonial knowledge in young children. Plausibly many young children aren't going to possess the counterfactual sensitivity described when listening to the words of their caregivers.

§3.1.4 The Humean Response

Hume (1999: 170) also holds that much of our knowledge comes about via testimony.⁶⁸ In "On Miracles", he writes that belief in testimony is not due to an a priori connection "between testimony and reality" but is based on "the usual conformity of facts to the reports of witnesses". (Hume, 1999: 171). There is no necessary connection between what is testified to and how the world is, but as it happens what is testified to usually does reflect how the world is. The level of justification that testimony enjoys is determined inductively; the more regularly facts and witness reports conform, the greater the justificatory weight of the testimony.

Hume doesn't want to just say that induction shows there to usually be conformity between testimony and reality and that therefore we can just take testimony to be justified; rather he discusses considerations that have a bearing on the justificatory force of testimony. A benefit of this more fine-grained approach is that if we are justified in believing competing claims, two or more claims that can't both or all be true, then we may be able to determine the claim of greater justificatory force;

⁶⁸ Faulkner (1998), Pritchard (2004: 327, footnote 6) and Gelfert (2010) each question how Hume is standardly interpreted in this debate. How he has been standardly interpreted and challenges to that interpretation are not the concern of my chapter. Mindful of these worries, however, in representing his views I stick to a close reading of his work and describe the view I find articulated there "Humean" rather than "reductionist" or "inferentialist".

something that we would be unable to do if we supposed testimony to be justified only on the basis of it being testimony. Similarly, if we have testimonial justification for believing *p* and we have non-testimonial justification for believing not *p*, then a more fine-grained account of testimonial justification may allow us to determine whether *p* or not *p* enjoys greater justification.

The considerations that Hume (1999: 171) regards as significant when determining the justificatory force of testimony include; the presence of contrary testimony; the character of the testifiers; the number of testifiers; the manner in which the testimony is delivered; and the interests the testifiers have in the testimony being affirmed. Another factor is the extraordinariness of the testimony for the recipient of the testimony. The more extraordinary the testimony to the experience of the recipient, the less the testimony should be credited; if, however, the testimony being false would be more extraordinary than the testimony being true, then the testimony should be believed to a degree justified by the remainder, the weight of evidence for *p* after the weight of evidence against *p* has been subtracted. (Hume, 1999: 172). Hume's list is not exhaustive; he supposes there to be many such particulars. (Hume, 1999: 171).

That such particulars contribute to our evidence and so impact on the justificatory force of testimony is supported by induction. Believing ultimately that, excepting in cases of perceptual knowledge and memorial knowledge, we are dependent on experience for knowledge, and therefore also presumably for knowledge conducive justification, he holds that belief should be attuned to the experienced frequency of conjunctions of events in accordance with his evidential calculus. (Hume, 1999: 110, 170). The wise man's degree of belief will reflect this variability of evidential weight; he "proportions his belief to the evidence".⁶⁹ (Hume, 1999: 170).

It appears that Hume is well-placed to explain how the judge can enjoy greater testimonial justification than the child; having more experience to draw on, she can

⁶⁹ In his own work concerning human understanding, John Locke (1975: 697) makes a similar claim. He writes that the mark of a lover of truth is that he does not entertain "any Proposition with greater assurance than the Proofs it is built upon will warrant."

potentially be more justified than the child. However, it seems natural to think that his account of testimonial knowledge discords with the Uniformity Claim. Requiring an agent to check with her experience in order to enjoy testimonial justification, at least consciously, isn't consistent with the claim that in standard cases of testimonial reception we simply accept testimony, and that we can have justified testimonial beliefs in those same cases.⁷⁰

Considering what Hume writes about testimony, it's not obvious why Hume should be taken to be committed to thinking that such checking is necessary. As long as agents do proportion their belief to the evidence, it's hard to see why actually checking with their experiences or not should matter for Hume. Of course the belief couldn't just be formed in a way that makes it so that it is only luckily proportioned to the evidence. A issue that arises is that even if it is correct that it doesn't matter to Hume's assessment of the force of an agent's doxastic justification whether they've actually checked or not, what does matter is that it would be unclear how it might happen that the agent's experience might inform their belief in the way that Hume thinks it should in the absence of actually checking. This is a problem given that we want to explain the Far-Reaching Dependence Thesis in a way that fits with the Uniformity claim.

Psychological research suggests that it may be possible to explain how an agent's experience might inform their belief in a way that would be consistent with the uniformity claim. Dual process theorists argue for the claim that "there are two different modes of processing", sometimes described as system 1 and system 2.⁷¹ (Evans, 2008: 256). System 1 processes are generally agreed to be "rapid, parallel and automatic in nature", while system 2 is taken to be "slow and sequential" and evolutionarily recent in comparison with system 1. (Evans, 2003: 454). In a more recent work Evans adds that there seems to be widespread agreement that system 1 is

⁷⁰ Hume (1999: 171) might challenge the uniformity thesis, he writes that frequently testimony is met with hesitation. However, I'll leave exploration of that possible challenge to one side as I believe it is not immediately crucial to the articulation of a promising view of testimonial justification.

⁷¹ Different authors have labelled the modes differently. Other labels for dual processes include experiential and rational, heuristic and systematic, intuitive and analytic, holistic and analytic. (Evans, 2008: 257).

unconscious and system 2 is effortful. (Evans, 2008: 270). Given what we know about the two systems, though admittedly in the absence of any definitive experimental results, it seems plausible to expect that system 1 picks up on cues of differing strengths that support and cues that undermine belief in testimony. It further seems plausible that experience of testimonial reception can contribute to an agent being sensitive and appropriately responsive to cues and so enjoying greater justification when they believe a piece of testimony. It would follow from the relevant work being done by system 1 that an agent's experience of standard cases of testimonial reception would conform to the Uniformity Claim. If this is right then we have an explanation of testimonial justification that accords with the Uniformity Claim and goes some way towards explaining what makes some testimonial beliefs more justified than others. To see that this seems supported by the system 1/system 2 distinction, consider how Frankish (2009) describes the workings of the two systems thus:

[M]ost of our behaviour is generated without the involvement of **personal** reasoning (system 2). Think about the actions involved in such everyday activities as driving a car, holding a conversation, or playing sports. These are intelligent actions, which are responsive to our beliefs and desires (think of how beliefs about the rules of the game shape the actions of a football player), and a great deal of complex mental processing must be involved in generating them. Yet, typically, they are performed spontaneously with no prior conscious thought or mental effort.⁷²

Nevertheless, even if there is a way for experience to act as an input into testimonial justification in a way that is consistent with the Uniformity Claim and the Far-Reaching Dependence Thesis, such an approach surely isn't quite right. First, it appears that such an approach brings us no closer to explaining testimonial knowledge in young children. Second, and relatedly as will become clear, we need a virtue component for an explanation for such justified testimonial belief and for testimonial belief generally. An agent won't be more justified in their testimonial belief simply in virtue of that agent being more experienced. Our judge is old but she is also wise. Rather, what is required is that the testimonial recipient be virtuous in

⁷² Interestingly the approach based on dual process theory outlined offers a way of seeing the counterfactual sensitivity credulist position and the Humean position as being much closer than the pair are usually represented as being.

her beliefs, though I assume that experience potentially contributes to the overall virtuousness of the testimonial recipient.

§3.1.5 Traits and Abilities

As we have seen in the previous section both the counterfactual sensitivity credulist view and the Humean view look to be badly positioned to account for testimonial knowledge in young children; the contribution, according to those accounts, required of an agent in order to have a justified testimonial belief seems to be such as to rule out young children gaining testimonial knowledge. Now recall the similarly demanding ability requirements of both Greco's robust virtue epistemology and Pritchard's anti-luck virtue epistemology; that the exercise of ability explains the truth of an agent's belief or that the exercising of ability on the agent's part produce that agent's safe true belief ("such that her safe cognitive success is to a significant degree creditable to her cognitive agency"), respectively. (Pritchard, 2012). Both requirements appear to rule out otherwise plausible cases of testimonial knowledge in young children. On the one hand it looks like we should say that there are cases of testimonial knowledge in young children, on the other hand it looks like a young child believing truly in such cases has little or nothing to do with the exercise of ability.

One motivation for saying that there are cases of testimonial knowledge in young children is based on response intuitions to cases such as the following:

Sarah, a four year-old child, asks her mother where, Lucky, the family dog, is. Her mother answers truthfully that her older brother Paul has taken Lucky for a walk.

Assuming this is a normal case, her mother is not for example out to deceive Sarah in this matter, then intuitively Sarah knows that Lucky has been taken for a walk.⁷³

That Sarah could have easily been deceived by her mother, that being a four year-old,

⁷³ It is necessary for Sarah to have understanding of the concepts involved in the proposition in order to be in a position to know the proposition. That such an understanding is necessary for knowledge is uncontroversial.

if her mother had said anything within a very wide range of possibilities, then Sarah would have still believed her, doesn't seem to undermine the intuition that she knows that Lucky has been taken for a walk. At the same time, however, such a consideration does undermine both the claims that Sarah believes truly because of the exercise of her ability and that her safe true belief is the product of ability such that the true belief is significantly creditable to her cognitive agency. That Sarah believes truly that Lucky has been taken for a walk is surely down to her mother's testimony that this is the case, and if someone is creditable to a significant degree for Sarah's true belief that Lucky has been taken for a walk then surely it's her mother.

Such cases of testimonial knowledge in young children might be taken as a challenge to virtue epistemic approaches generally. Such a challenge would involve moving from the plausible claim that young children in Sarah type cases don't exercise a reliable cognitive ability in either of the ways described by Greco and Pritchard and that there is therefore no such requirement for knowledge, to the alternative and less plausible conclusion that no contribution from the cognitive agent to her/his believing truly is required. Before showing that these two conclusions can and should be teased apart, it is worth pausing to consider a salient feature of Sarah type cases, a feature of that type of case which is such that we're inclined to attribute knowledge.

Part of the intuition that Sarah knows seems to be related to the presumed relationship between speaker and hearer; to see that this is so, notice that our intuition that Sarah knows may weaken or disappear altogether if we imagine Sarah asking a neighbour where his dog is. Even if the neighbour answers accurately and sincerely, whether Sarah knows now seems to depend on her relationship with the neighbour; whether this neighbour is a friend of the family, someone who talks to Sarah quite regularly and is playing some, even if small, part in Sarah's upbringing, or whether this is a neighbour with whom Sarah's family have barely any contact, someone with whom Sarah seldom speaks or with whom up until now Sarah has never spoken, someone who is a stranger to Sarah. I come back to how the relationship between testifier and testimony recipient has a bearing on whether there

is knowledge, or more precisely justified belief, later in this section. For now it suffices to flag the consideration.

Having briefly considered the significance of the relationship between Sarah and the testifier, let's return to trying to tease apart the aforementioned conclusions. While the Sarah case indicates that there is no need for the exercising of a reliable cognitive ability in either of the ways that Greco and Pritchard claim in order for an agent to gain knowledge, there may still be a need for the exercise of epistemic virtue. On the one hand, while the Sarah case plausibly is a case in which the protagonist gains knowledge; on the other hand, testimonial cases may be constructed in which the mere reliability of the testifier as well as the good epistemic practices of the broader community that explain the agent believing truly won't be sufficient to get the agent knowledge. For example, suppose the agent possesses apparently strong defeaters against the truth of the relevant belief, or, say, against the reliability of the testifier and yet continues to believe. If this were the case, then, even if the belief is in fact true and the testifier is in fact reliable, we wouldn't claim that an agent in such a case knows. An intuitive explanation for saying so is that an epistemically vicious belief that *p* cannot be a candidate for knowledge that *p*. This is based on the plausible assumption that an agent continuing to believe that *p* in the face of apparently stronger defeaters against the truth that *p*, or against the reliability of the testifier, renders her/his belief an epistemically vicious belief.

While the epistemic virtuousness and epistemic viciousness of beliefs undoubtedly come in degrees and while there are surely difficult boundary cases in which it may be difficult to determine whether the belief is epistemically virtuous or epistemically vicious, the claim that a particular belief is either epistemically virtuous or epistemically vicious seems right. Perhaps that this seems so is related to us thinking that in cases in which a proposition cannot be believed epistemically virtuously an agent should suspend belief or not believe either way. The foregoing commits one to denying that beliefs may fall into a zone of epistemic permissibility in which a belief is neither virtuous nor vicious. This implies, and this is why this is of interest for my purposes, that claiming that vicious belief cannot be a candidate for knowledge

implies that if a belief is a candidate for knowledge, then that belief can only be a virtuous belief.

Now the question which arises is whether there can be a virtuous belief that doesn't involve the exercising of a reliable cognitive ability and whether that kind of virtuous belief can be a candidate for knowledge. Let me deal with the first part first. In certain everyday cases that an agent may find herself in, it seems correct to say that she is not in a position to exercise a reliable cognitive ability such that it will explain her believing truly. Sarah type cases are such cases.⁷⁴ Yet in such cases it seems perfectly appropriate epistemically for such listening agents to believe; in other words belief in such cases seems epistemically virtuous rather than vicious. Given that, in such cases, as has been previously argued, the agent believing truly is not because of the exercise of cognitive ability and the cognitive agent is not primarily creditable for the truth of her belief, or significantly creditable for her true belief, on the basis of the ability exercised, the virtuousness of the belief remains to be explained. An obvious move would be to defend a position according to which less of a contribution from the ability of an agent is required. Aside from such a position being vulnerable to Gettier type counterexamples, it simply seems ill-suited to explaining what is virtuous about the relevant beliefs in Sarah type cases. To put the point differently, it doesn't look like the epistemic virtuousness of the young child's belief is owing to the exercise of her cognitive ability.

A plausible alternative explanation of the virtuousness of the belief is that the virtuousness of the belief is attributable to the exercising of a relevant epistemically virtuous trait. The thought is that if we grant that it's implausible that cognitive ability is doing the relevant work such that the agent epistemically appropriately believes, and assuming that the right account of the nature of knowledge is a virtue epistemological account, an account upon which the cognitive agent must make some contribution towards her believing truly in order to have knowledge, then if an agent believes in such a case as it is epistemically appropriate for her to believe, then

⁷⁴ Typical cases of non-experts forming beliefs on the basis of expert testimony plausibly also meet this criterion. I discuss the non-expert/expert relationship in depth in Chapter V.

it looks like her appropriate belief might be explained by the exercising of an epistemically virtuous trait.

Thinking of the kind of belief required for knowledge, an epistemically virtuous belief, either as being from the exercising of an epistemically virtuous trait or from the exercise of a cognitive ability puts us in a much better position to make sense of intuitive cases of testimonial knowledge in a way that fits into a virtue epistemological framework. I have been focusing on Sarah type cases because they seem to be the hardest type of case for both Greco and Pritchard to account for, but what I say with regard to dealing with such cases also plausibly applies to non-Sarah type cases of testimonial knowledge, though not necessarily all cases of testimonial knowledge. Plausibly, given good epistemic practices of an epistemic community, being trusting is, *prima facie*, an epistemically virtuous trait, although the conditions under which it is virtuous for a normal adult to trust differ from the conditions under which it is virtuous for a child to trust.⁷⁵

A motivation for distinguishing the conditions under which it is virtuous for a child such as Sarah to trust and the conditions under which it is virtuous for an adult such as Jenny to trust is that an account on which a normal adult only had to meet the conditions a child is required to meet in order to have a virtuous belief would yield the wrong results. By yield the wrong results I mean that it would commit us to saying that cases that intuitively aren't cases of knowledge are cases of knowledge. The distinction of the conditions under which it is epistemically virtuous to trust allows us to distinguish between adults who believe testimony from an epistemically virtuous trait and adults who believe from an epistemically vicious state, such as gullibility, while allowing us to hold onto the claim that Sarah's belief is from an epistemically virtuous trait. One reason why categorising the testimonial belief of young children as knowledge is so tricky is that it looks like children, generally being credulist in the way that they are, should mean that we put their belief into the same category as gullible beliefs and a plausible candidate for true testimonial belief being precluded from qualifying as knowledge is testimonial belief formed because of

⁷⁵ By "normal adult" I just mean an adult that is typical in the relevant respects of adults in the actual world.

gullibility, formed because the hearer would believe almost anything said. Yet it seems epistemically appropriate for Sarah to believe that Lucky has been taken for a walk and there remains the intuition that she knows that Lucky has been taken for a walk.

It would be odd to criticise a young child as gullible. And yet it's not as if adults don't criticise or blame children. A parent might say to another that a child is very demanding, or that a child is very aggressive; similarly a parent might say to a child that she shouldn't throw her food on the floor or that she shouldn't have given her baby sister's favourite soft toy to the family dog to chew on. And yet believing almost anything they're told doesn't seem an appropriate criticism of a young child, while in contrast it would be an appropriate criticism of a normal adult. Children are sometimes admonished not to believe everything they're told, but one supposes that this is at a later stage of development than the stage children such as Sarah are at and it is such children whom we have in mind. It seems likely that the children that are admonished are admonished because they are expected to be able to exercise some degree of discrimination. Young children are in a unique situation epistemically, they neither have nor have had the opportunity to build up the experience required in order to be discriminating in the testimony they receive. Even if this wasn't the case, it seems that they don't have the requisite abilities developed to a sufficient degree to be able to make use of any relevant potential experience. For young children, at their early stage of development, given the difficulty they would have in developing further both generally and epistemically in the absence of trusting their caregivers, it is epistemically appropriate, and as such epistemically virtuous, of them as cognitive agents to trust. Admittedly this way of explaining the epistemic virtuousness of the trust of young children is one on which practical factors such as underdeveloped ability and lack of experience have a bearing on whether the basis on which an agent believes is deemed epistemically virtuous. Given some consideration, the manner of the explanation shouldn't be seen as worrisome.

What we suppose counts as virtuous for a normal human adult differs from what we would suppose counts as virtuous for, say, a god. Part of what we suppose

determines what counts as virtuous for a normal adult human being is based on what adult humans are capable of and what their limitations are. For example, a normal adult human agent may virtuously believe that there is a sheep in the field in the Roddy case, albeit believing because of her misidentification of a sheep shaped object in that field, although believing from the exercise of cognitive ability; but a god, and perhaps a sheep expert, plausibly would not virtuously believe on the same basis.⁷⁶ If one holds that the conditions under which a belief's virtuousness can vary like this, then one is well-positioned to defend the claim that Sarah believes virtuously.

Intuitively, it is similarly epistemically virtuous for a person to trust an expert claim, given suitable background conditions. Whether an agent's belief is epistemically virtuous, however, varies according to the role that agent is playing. For example, it wouldn't be epistemically virtuous for a peer tasked with reviewing the expert's claim simply to believe the expert's claim on the basis of her testimony.⁷⁷ It's not as if we'd just criticise him for not doing the job expected of him but grant that his simple belief of the expert testimony, the same way the non-expert believes the testimony, is epistemically virtuous.⁷⁸

While thinking of knowledge as requiring virtuous belief, with the possibility that this includes belief because of the exercise of a virtuous trait rather than only belief that is because of ability, helps us provide a virtue epistemological account of cases of testimonial knowledge, the associated virtue epistemic claim that the cognitive

⁷⁶ Depending on the conception of a god in play, it might be thought either a god doesn't form beliefs or that a god wouldn't be lucky in gaining a true belief in the way described. For my part I'm using "god" as a surrogate term for a being that's massively more advanced in terms of its cognitive abilities than the normal human adult.

⁷⁷ Greco described a case like this in a keynote address at the 2012 *Edinburgh Graduate Epistemology Conference*. On the basis of this case and others like it, Greco ultimately defended the view that the standard for what counts as testimonial knowledge can vary according to the role that the testimonial recipient is playing in the epistemic community, whether he is acting as an epistemic gatekeeper for the epistemic community or not. The address came when I was close to finishing my thesis and I'm not aware of there being a print version of it, so I haven't explored the ideas presented in the address in my work. Nonetheless, on the face of it I think the criticisms articulated of his account in Chapter II stand.

⁷⁸ Consideration of such a type of case also, and relatedly, lends support to the claim that knowledge is an honorific term in the sense of attributing knowledge to someone being a praise of that person in a certain respect.

agent is creditable for their believing truly needs to be explored. The cognitive agency of a young child does make a contribution to her gaining a belief that is true in Sarah type cases. She trusts rather than not trusts. That Sarah is trusting is epistemically virtuous given her stage of development. If requirements external to Sarah's cognitive agency are met, then she will have knowledge. Nonetheless the trait of trusting shouldn't be confused with an ability. That Sarah believes truly isn't plausibly down to a faculty of her cognitive agency, unlike say visual perception, that can be exercised to get true beliefs. A consequence of defending an account on which Sarah has a true belief but not because of the exercise of ability is that one must give up the notion that a knower is always primarily creditable for the truth of her belief or that her safe cognitive success is to a significant degree creditable. What seems right to say about Sarah type cases, and is befitting of cases being social epistemological cases, is that credit is shared. While credit for the truth of the belief is down to the first testifier who believes truly what he testifies to because of the exercise of ability and to the wider community that sustains good testimonial practices, the testimonial belief is down to the testimonial recipient and that belief is intuitively creditable if the recipient believes virtuously.

The account I offer leaves open the possibility that some true testimonial beliefs may be gained because of the exercising of ability. The judge in the Jana case described earlier would plausibly be one such case. The account also leaves open the possibility that in some cases of testimonial knowledge an agent's virtuous belief may be formed from a combination of the exercising of a cognitive ability and the exercising of an epistemically virtuous trait; plausibly it is such a combination that would make Jenny's belief in the Jenny* case epistemically virtuous and put her in the running for knowledge. A task remaining is to explain how testimonial doxastic justification may differ in degrees on my account, an account on which the justification condition of knowledge is satisfied by a belief being virtuous either because of the exercise of an virtuous trait or because of the exercise of a cognitive ability.

§3.1.6 Preliminary Conclusion

In this chapter I argue that neither the traditional approaches to providing accounts of testimonial knowledge, counterfactual sensitivity credulism and Humeanism, can help overcome the challenge encountered by both Greco and Pritchard's treatments of testimonial knowledge. I make the case that testimonial knowledge in young children is best explained by belief in the relevant cases being down to an epistemically virtuous trait, trusting, rather than from the exercising of ability. This explanation allows for a plausible defence of a virtue epistemological account of testimonial knowledge in young children in typical cases, albeit one that results in denying the claim that the truth of a belief in cases of knowledge is always primarily or even significantly creditable to the cognitive agency of the believer.

§3.2.0 Why Epistemic Grace?

In this second part of my chapter I attempt to deal with the challenge raised by the Barney case. Having in the previous section argued for expanding what we think makes for an epistemically virtuous belief to include beliefs held because of the exercise of an epistemically virtuous trait, in this section I make an argument for what knowledge requires in addition to an epistemically virtuous belief. On the one hand, as Greco's account of the nature of knowledge stands, the Barney case looks like a counterexample to that account, on the other hand, Pritchard's attempt to provide an account that rules out such cases as being cases of knowledge looks under-motivated. In this part of my chapter I argue for an alternative to Pritchard's safety condition. I do so initially on the basis of an alternative reading of standard Gettier-type cases, a reading which is equally applicable to Barney-type cases. I'll argue that such a reading is intuitive and ultimately helps us to provide plausible explanations of the relevant cases.

My starting point in attempting to overcome the challenges outlined is a non-standard reading of the almost 50 year-old Gettier problem.

Consider again the Roddy case:

Using his reliable perceptual faculties, Roddy non-inferentially forms a belief that there is a sheep in the field before him. His belief is also true. Unbeknownst to Roddy, however, the truth of his belief is completely unconnected to the manner in which he acquired this belief since the object he is looking at in the field is not a sheep at all, but rather a sheep-shaped object which is obscuring from view the real sheep hidden behind. (Pritchard, 2012: 251).

The standard analysis of Gettier-type cases has been that although the protagonists in the cases have true beliefs that are justified, the justification for their beliefs is not appropriately connected to what makes their beliefs true; it's a matter of luck that their beliefs are indeed true.⁷⁹ Following this reading of Gettier-type cases, the task for epistemologists has been taken to be to provide an account of the nature of knowledge that either directly, via an anti-luck condition, or indirectly, via a strengthened justificatory condition, rules out such luck; these approaches have been pursued by Pritchard (2010) and Greco (2009) respectively.

It's not obvious, however, that we should see the cases this way. Linda Zagzebski (1994: 66) identifies Gettier cases as involving two elements of luck; bad luck that one's justification, which is sufficiently robust to satisfy the justification condition for knowledge, doesn't get one a true belief, and good luck that one happens to get a true belief. This identification of two elements of luck in Gettier cases doesn't by itself, however, challenge the view that a condition that rules out the relevant epistemic good luck is required. Zagzebski's identification of both good luck and bad luck does, however, provide resources for a challenge to the standard view.⁸⁰

A non-standard way of seeing Gettier-type cases is one according to which the agent in a Gettier-type case should not be seen as lucky because he has a true belief; rather he should be seen as unlucky because he doesn't have knowledge. In Gettier-type cases, agents have justified beliefs and ordinarily having a justified belief means that knowledge has been gained.⁸¹ In such cases, although knowledge has not been

⁷⁹ For an early attempt to make good on this reading of the Gettier problem, see Goldman (1967).

⁸⁰ Although Zagzebski (1994) considers various strategies, to rule out Gettier-type cases, her own view is that they are inescapable.

⁸¹ Later in the chapter I return to the claim that ordinarily having a justified belief means that knowledge has been gained. At this point I want to sketch out the alternative way of seeing the Gettier

gained, if knowledge had have been gained then the agent would have had a true belief as well. A natural reading of Gettier-type cases is that the agents therein are better described as unlucky not to have knowledge, rather than lucky to have a true belief.

By way of analogy, consider again the case of Jana, who loses her fortune on a particular day because of bad luck, but, because of good luck, she acquires a lesser fortune that same day. Other things being equal, it would be more appropriate to say she has been unlucky rather than lucky or at least more unlucky than lucky on this day. To keep things simple let's change the case slightly and say that the fortune she loses and gains is in the form of cash in both cases. Interestingly, the appropriateness of the analogy with Gettier-type cases might be contested on the basis of claiming that of knowledge and that which falls short of knowledge, it is only true belief that is of significant value, and therefore that in Gettier-type cases nothing of significance is lost when one loses knowledge but still ends up with a true belief. Such a position, if advocated, however, would come at the cost of committing the advocator to a revisionary solution to the *Meno* problem, at least when considered synchronically.⁸² Wanting to stay neutral on such issues might be thought of as a dialectical reason for not going further than saying that there is good luck and bad luck in Gettier-type cases, but that neutrality won't motivate the inclusion of an **anti-luck** condition in an account of the nature of knowledge.

Alternatively, one may wish to defend the view that the kind of luck present in Gettier-type cases is value neutral, in other words it is neither good nor bad luck. This view however, is problematic. As Ballantyne (2012) writes, significance is widely held to be a requirement for luck, marking the difference between luck and something that is merely unlikely.⁸³ Ballantyne describes the consensus in the debate

problem. I take it that the claim is sufficiently plausible such that its inclusion does not render the sketch unsatisfying.

⁸² One might have a Socratic account according to which knowledge is of more value than a mere true belief, because having knowledge entails having a justified belief, and having a justified belief better enables one to hold on to knowledge over the course of time. Such an account could motivate the view that when we freeze frame matters, and so take them synchronically, then knowledge is no more valuable than true belief. Neither Greco nor Pritchard have such an account.

⁸³ Thanks to Lee Whittington for this point and directing me to Ballantyne's paper.

about luck thus; “[t]he consensus just is, at bottom, that if an event is lucky for an individual, then it’s somehow *good for* or *bad for* her.” (The emphasis is Ballantyne’s own). Such a consensus looks right when we consider a case offered by Pritchard (2005: 132), mentioned by Ballantyne, in support of the requirement. An unlikely landslide that occurs that affects nobody is not lucky because it is significant to nobody. We can further stipulate that the unlikely landslide occurs on a distant, uninhabited planet to make the intuition that significance is a requirement for luck clearer.

While this seems right, it is worth pointing out that luck in everyday language is often used in a comparative way where the event itself is not plausibly thought of as good for the individual but is thought better than some considered alternative that is relatively nearby. Consider the usage of “luck” in the “Four Yorkshiremen” sketch by Monty Python. The sketch begins with four well-dressed Yorkshiremen commenting on the pleasing taste of an expensive wine before they move on to reminisce about their childhoods:

“TG (Terry Gilliam): But you know, we were happy in those days, though we were poor.

MP (Michael Palin): Aye. BECAUSE we were poor. My old Dad used to say to me, "Money doesn't buy you happiness."

EI (Eric Idle): 'E was right. I was happier then and I had NOTHIN'. We used to live in this tiiny old house, with greaaaaat big holes in the roof.

GC (Graham Chapman): House? You were lucky to have a HOUSE! We used to live in one room, all hundred and twenty-six of us, no furniture. Half the floor was missing; we were all huddled together in one corner for fear of FALLING!

TG: You were lucky to have a ROOM! *We* used to have to live in a corridor!

MP: Ohhhh we used to DREAM of livin' in a corridor! Woulda' been a palace to us. We used to live in an old water tank on a rubbish tip. We got woken up every morning by having a load of rotting fish dumped all over us! House!?
Hmph.

EI: Well when I say "house" it was only a hole in the ground covered by a piece of tarpolin, but it was a house to US.

GC: We were evicted from *our* hole in the ground; we had to go and live in a lake!

TG: You were lucky to have a LAKE! There were a hundred and sixty of us living in a small shoebox in the middle of the road.

MP: Cardboard box?

TG: Aye.

MP: You were lucky. We lived for three months in a brown paper bag in a septic tank. We used to have to get up at six o'clock in the morning, clean the bag, eat a crust of stale bread, go to work down mill for fourteen hours a day week in-week out. When we got home, our Dad would thrash us to sleep with his belt!

GC: Luxury. We used to have to get out of the lake at three o'clock in the morning, clean the lake, eat a handful of hot gravel, go to work at the mill every day for tuppence a month, come home, and Dad would beat us around the head and neck with a broken bottle, if we were LUCKY!"

("Four Yorkshiremen Sketch". The only alteration to the quoted text is the bracketed additions of the names for which the initials stand.)

The upshot of the non-standard reading of Gettier-type cases, thinking of the protagonists in Gettier-type cases as unlucky not to know, is that an alternative to the anti-luck direction taken in epistemology emerges into view. The task given this way of thinking of Gettier-type cases isn't to include a condition in one's analysis of knowledge that rules out certain types of good epistemic luck; rather the task is first to identify how bad epistemic luck may be knowledge undermining and, second, based on that identification, offer a suitable account of the nature of knowledge.

§3.2.1 Epistemic Grace: Gettier Cases

The bad epistemic luck I've described in Gettier-type cases, and this can be said of both standard Gettier-type cases and Barney-type cases, can plausibly be conceived

as being an absence of co-operation from the world, or, as McGlynn (2012) might put it, “uncooperative external circumstances”.⁸⁴ Despite Roddy and Barney having justified beliefs, beliefs that are the product of the exercise of cognitive ability, and their beliefs being true, neither has knowledge. Their justified beliefs put them in the running for knowledge, that they don’t have knowledge is due to factors beyond the contribution that is required of them.⁸⁵ If this is accepted, and independent of particular accounts of the nature of knowledge it seems quite a natural way of thinking of the cases, then logically we also get the point that knowledge depends, **at least sometimes**, upon something which is beyond the scope of what is realised by the exercising of an agent’s cognitive ability, and the subsequent belief being true. A further, stronger point seems plausible; that, for **any** belief, despite a human agent forming a virtuous belief, forming a belief because of a cognitive ability or of an epistemically virtuous trait, and that belief being true, the agent may still not gain knowledge.⁸⁶ An agent not having knowledge in such cases, other things being equal, is unlucky.

This way of conceiving of bad epistemic luck and its theoretical relationship to knowledge, presupposes a certain way of thinking of justified belief. The accounts of Alexander Bird (2007), Martin Smith (2010) and Jonathan Ichikawa (draft) offer elaboration on the thought that justified belief should be thought of as would-be knowledge or potential knowledge.⁸⁷ Smith (2010: 12) does so as follows, “[m]y belief is justified just in case I have *done my epistemic bit*—the rest, as it were, is up to fate.” Directly after, and in a way that resonates with thinking of the knowledge

⁸⁴ I say “might” because although McGlynn does use the phrase “uncooperative external circumstances” as an alternative to talk of a lack of cooperation from the world, as far as I’m aware he doesn’t use it as a diagnosis for an alternative reading of the Gettier problem.

⁸⁵ Greco (2010: 167) would claim that an agent also needs to meet a weak internalist condition in order to be in the running for knowledge; the agent must be motivated towards the truth such that if she had any reason not to believe the reliability her way of forming beliefs, then she wouldn’t believe.

⁸⁶ What is here written is not intended as challenging the possibility of infallibilist justification, that an agent’s justification, say that p, may entail the truth of p, rather the claim is that possession of a justified belief, whether the justification entails the truth of the belief or not, is not sufficient for knowledge.

⁸⁷ It should be noted that the debate in which Bird, Smith and Ichikawa are involved is inspired by the knowledge first approach, which eschews analysing knowledge into component parts and according to which justification should be understood in light of an understanding of knowledge. Nonetheless, the account which Smith provides of justification, is not ultimately, however, described in terms of knowledge.

undermining epistemic luck in Gettier cases as being bad rather than good, Smith writes, “[m]y belief will qualify as knowledge provided that the world obliges or cooperates—but *I* am not required to do anything further”. (The emphasis is Smith’s own in both quotes). Although the two ways resonate with one another, they can also be pulled apart, which raises the question of which provides the right answer to what separates justified belief from knowledge.

If one accepts that knowledge must be safe, regardless of whether it requires an independent condition or not, and that, in any case, as discussed in the previous chapter, an ability condition that requires the securing of safety from an ability would be too strong, then one accepts that the world, in the broadest possible sense, needs to be a certain way in order for there to be knowledge. Given this, rather than saying an **absence** is required, a positive construal of what is required is appropriate; therefore an absence of bad luck is not what is required.

The idea of knowledge requiring the cooperation of the world fits better with such a consideration. This idea, however, does not neatly capture the theoretical relationship between the agent’s contribution and the world’s contribution that is necessary and jointly sufficient for knowledge. More precisely, while it seems correct to say that knowledge requires cooperation from the world, simply saying that does not capture the ordering of cooperation in a case of knowledge, the prior work required of the agent before what the world can contribute has a bearing on whether there is knowledge. The differences here are subtle and perhaps minor, and if one wanted to offer a modified and an enlarged account in the language of cooperation from the world, then one undoubtedly could. My theoretical preference, however, is to talk in terms of what I call “epistemic grace” which seems to me to provide an immediate and intuitive grasp on what is required for knowledge.

While conceiving of knowledge as requiring either an absence of bad luck or cooperation from the world is intuitive and puts us on the right track, ultimately conceiving of knowledge as requiring “epistemic grace” is more explanatorily fruitful than a modified bad epistemic luck or epistemic cooperation from the world

requirement. First, however, what is meant by “epistemic grace” should be explained. Epistemic grace can be understood by way of analogy with a religious conception of grace. According to a religious, synergetic account of grace, whether someone receives grace and so is saved is not just down to God, but also depends on the person acting in the right way. In the epistemic case, the thought is that it is through good epistemic work, having a virtuous belief, that one puts oneself in a position to gain knowledge, but whether one does indeed gain knowledge depends on whether one’s virtuous belief enjoys epistemic grace. Claiming knowledge to require epistemic grace reflects, as highlighted earlier, the dependence of any epistemically virtuous belief on the world in order for it to be knowledge. Tying what I’ve argued knowledge requires together, we get the analysis that:

S knows that p iff S has an epistemically virtuous belief that p and that virtuous belief enjoys epistemic grace.

Although this analysis as presented here isn’t a reductive analysis, I hope that what I’ve written elsewhere in this chapter shows that such an analysis, though it would be more complex than the analysis given, is possible.

§3.2.2 Dealing with Gettier Cases

Key to the presence of religious grace is the action of God. The presence of epistemic grace depends on whether the environment in the particular case in which a virtuous belief is formed is such that the way the belief has been formed is in fact appropriate given the environment. There are two dimensions of normativity immediately relevant to knowledge. One dimension, related to virtuous belief, is anchored in what the agent in a particular case should believe, taking into account factors such as what he knows, the evidence that is available to him in the case, his epistemic duties – say as a police witness or as a peer reviewer. The other dimension is anchored in an objective perspective, how the world is, and whether, given how the world is, the way in which the agent formed her belief in the particular case is suited to getting the truth in the environment in which the belief was formed.

Epistemic grace is present when the virtuous belief is virtuously formed in such a way in the particular case that it is suited to getting truth in the environment in which the belief is formed. Saying that the way in which the agent formed her belief is **suited** to getting the truth means that the way is such that it yields a true belief in the environment; environment here refers to both the local and regional environment.

This way of thinking of knowledge helps us deal with the cases that have previously been shown to be of concern. In the Roddy case, Roddy has a virtuous belief but his belief doesn't enjoy epistemic grace, it turns out that given how the environment is, the way he formed his belief is such as to be ill-fitted to getting the truth in his environment. In the Barney case, like Roddy, Barney forms a virtuous belief, like Roddy, it turns out that the world Barney is in is such that the way he formed his belief is not a good way of getting the truth in his environment. In the Sarah case, Sarah forms a virtuous belief. It turns out that the world Sarah is in is such that the way she formed her belief is a good way of getting true beliefs in her environment. The approach argued for gives us the right answer to the cases and does so in a well-motivated way. This way of thinking of knowledge also yields the intuitive answer that there is justified belief in brain in a vat type cases. While the agent with the experiences that she has forms the virtuous belief that she has hands, her virtuous belief does not enjoy epistemic grace – it turns out that the way her belief is formed is such that her belief is not appropriate given her environment.

A belief that enjoys epistemic grace is not just true in most nearby worlds but it is also true in the actual world. Intuitively it wouldn't be appropriate to say that a virtuous belief enjoys epistemic grace if it weren't actually true. Therefore, a belief enjoying epistemic grace implies the truth of the belief. Similarly, if a virtuous belief is formed in such a way as to be true in the actual world but not true in most nearby worlds, then the virtuous belief doesn't meet the second normative criterion, i.e. that the virtuous belief be formed in such a way as to be well fitted to getting truth in the agent's environment. On my account, whether a belief enjoys epistemic grace depends on the belief being an epistemically virtuous belief. It is by a belief being epistemically virtuous that, so to speak, it is in the running for epistemic grace. As

such, the epistemic grace requirement is not independent from the requirement that a belief be epistemically virtuous. This is appropriate given that epistemic grace has been characterised as potentially being enjoyed by beliefs that are candidates for knowledge, by epistemically virtuous beliefs.

§3.3.0 A Possible Worry

I've argued that an epistemically virtuous belief is necessary in order to know. But isn't there an important sense in which an epistemically virtuous belief itself is owing to epistemic grace? The same kind of question may be asked of a synergetic account of a religious conception of grace, that is, whether someone knows how to or is in a position to do the right things to put herself in a position to receive grace also seems to require grace. Rather than this being worrisome for my account, it, on the contrary, seems appropriate if we keep the ordering clear. That S may be born into an epistemic community that has good epistemic practices and so she herself is in a good position to gain good epistemic practices and so is in a position to form epistemically virtuous beliefs may itself be thought of as a matter of epistemic grace. But S, as a cognitive agent, has to be doing what it takes to have a virtuous belief if her belief is to be in the running to be knowledge.

§3.4.0 Conclusion

Motivated by a non-standard reading of the Gettier problem, I have articulated a conception of epistemic grace and have argued that it is a requirement for knowledge. In so arguing, coupled with the epistemically virtuous belief condition argued for in my section on testimonial justification, I have shown how such a requirement allows for satisfactory handling of both standard Gettier-type cases and Barney-type cases. Furthermore, by requiring epistemic grace we get the result that Barney doesn't know. While such a result is secured by a safety requirement, I have made the case that considering there to be an epistemic grace requirement is well motivated, whereas, as discussed in the previous chapter, the safety requirement does not seem so motivated. My chapter addresses the challenge of dealing with both Barney-type

cases and Sarah type cases, yielding the conclusion that an agent knows a proposition if his belief is epistemically virtuous and that virtuous belief enjoys epistemic grace. I believe that this account of the nature of knowledge overcomes the challenges outlined at the beginning of the chapter and has, as we shall see, interesting implications for the value of knowledge.

Chapter IV: Valuing Virtue and Grace

§4.0.0 Introduction

Knowledge on Greco's account is distinctively valuable in virtue of knowledge being an achievement. In Chapter I I argue that his account of the value of knowledge is in one respect too narrow and in another too broad. I argue that his account is unduly narrow on the basis that virtuously held belief also satisfies the criterion for achievement. I argue that his account is unduly broad in so far as his account suggests each instance of knowledge has pro tanto value. In support of my claim that Greco's account is too broad in this respect I make the case that some instances of knowledge do not have pro tanto value.

Having already argued that the account of the nature of knowledge I provide overcomes the challenges faced by Greco's and Pritchard's accounts of the nature of knowledge, in this chapter I argue that this account of the nature of knowledge also provides a foundation for a plausible account of the value of knowledge and epistemic goods generally. I return to the charges developed in Chapter I that Greco's account of the value of knowledge is too narrow in a certain respect and too broad in another respect, considering those charges in relation to my own account of knowledge. With regard to the former charge, I argue that virtuous belief being an achievement on the part of the cognitive agent fits neatly into the account of the nature of knowledge which I offer. I then argue that the value of true belief helps us account for a value difference of knowledge vis-à-vis justified belief. I argue that while the precise value of true belief has a bearing on whether knowledge is distinctively valuable, the value of true belief doesn't help us account for a value difference of knowledge vis-à-vis justified true belief. I make the case that the account of the nature of knowledge which I offer, having as it does the epistemic grace requirement, allows us to account for the value difference of knowledge vis-à-

vis justified true belief. Next, following on from my criticism that Greco's account of the value of knowledge is too broad, and independent but consistent with the account of the nature of knowledge which I offer, I argue that the possession of an instance of an epistemic good may enjoy final value, in virtue of being constitutive of the good life, if possession of that particular instance of the good is virtuous simpliciter. Having discussed the value of knowledge and that which falls short of knowledge, I briefly consider the resources that my account has for explaining the value of understanding and wisdom, two other prized epistemic goods. In the last section prior to concluding, having previously set out the value of various epistemic goods, given the impact of the epistemic environment on our potential attainment of these goods I argue that we have reason to preserve or adapt our epistemic environment accordingly.

§4.1.0 The Value of Knowledge and that which Falls Short of Knowledge

To recap, borrowing from Pritchard (2010: 31), Greco's argument runs as follows:

- (P1) Achievements are successes that are because of ability (Achievement thesis);
- (P2) Knowledge is a cognitive success that is because of cognitive ability (Robust Virtue Epistemology);
- (C1) So, knowledge is a cognitive achievement (Knowledge as Achievement thesis);
- (P3) Achievements are finally valuable (Value of Achievements thesis);
- (C2) So, knowledge has final value.

As can be seen, Greco's account of the value of knowledge neatly falls out of his account of the nature of knowledge. What's more his overall account of knowledge, criticisms already made notwithstanding, appears to offer us ready solutions to the value problems discussed in Chapter I.

Over the course of the previous three chapters I have argued for moving away from Greco's account of knowledge at a number of junctures. In what follows I'll present the account of the value of knowledge that emerges from having made these moves.

In Chapter I I argue that there is a good case for thinking that a justified belief or virtuous belief is an achievement. This case itself depends on the claim, argued for previously, that a virtuous belief is a cognitive success, one from the exercise of ability, and the further claim, accepted by Greco, that an achievement is a success from ability. In Chapter III I further argue that the exercise of virtue, rather than the exercise of a cognitive ability, is a requirement of knowledge. I argue that the exercise of virtue as a requirement for knowledge helps explain testimonial knowledge in young children while a narrower ability requirement does not. Rather than claiming that a belief may only be virtuous if it's from the exercise of ability, I argue that a belief may also be virtuous if it's from the exercise of a virtuous trait.

Now a question that arises is whether all cases of virtuous belief, as I have explained virtuous belief, and by extension knowledge, are cases in which there is an achievement? A second and related question is whether achievement should be understood as broader than has previously been described. Affirming the former requires affirming the latter. Denying the latter requires denying the former. Either move will require another departure from Greco's account. Either not all knowledge will be an achievement or all knowledge is an achievement but achievement is not as previously defined.

It seems right to say that if a cognitive agent has a virtuous belief, then regardless of whether it's from an ability or trait, the success of having that belief is down to the exercising of virtue on the part of the agent. The agent could not have a virtuous belief of course without the agent making a virtuous contribution. On this basis it looks like a virtuous belief should be regarded as an achievement and one that is down to the virtue of the relevant agent. This is supported by our intuitions in other cases. Being a virtuous parent is plausibly an achievement. It is further plausible that a parent may be so because she exercises certain traits rather than abilities, for

example she may be a virtuous parent because she is loving, moral, and understanding. Again, such a consideration lends support to the idea that the account of achievement previously offered should be broadened to success that is because of virtue.

What has been said so far about the value of virtuous belief is relevant to knowledge in so far as knowledge entails virtuous belief. A central task in this chapter is to account for the value of knowledge and not just the value of virtuous belief. Knowledge does seem to be more valuable than virtuous belief. Knowledge entails not only a justified belief but also a true belief. If we think of a true belief as enjoying *prima facie* or *pro tanto* value then we get an immediate ground for saying that knowledge is more valuable than virtuous belief. One argument made in the literature for true belief enjoying such value is that true beliefs are instrumentally valuable for achieving practical ends. (Pritchard and Turri: 2012.)

Aside from this value, there is some ground for believing that true belief may be valuable on another ground, a ground with implications for moves that are made in the dialectic. It's implicit in both Greco (2011) and Pritchard's (2010) discussion of the former's position that mere true belief doesn't enjoy final value as an achievement or enjoy final value on some other basis. If this were otherwise, then Greco's account of the value of knowledge would not appear to be a candidate for solving the tertiary value problem, explaining why knowledge is distinctively valuable *vis-à-vis* that which falls short of knowledge. Neither Greco (2011) nor Pritchard (2010), however, offers detail on what value they regard true belief as having. The question as to the value of true belief remains.⁸⁸

Intuitively true belief is valuable, although cashing out that intuition presents a challenge. True belief is, perhaps, also finally valuable and valuable in virtue of true belief, *prima facie* or *pro tanto*, being constitutive of the good life. Consider the

⁸⁸ Pritchard (*forthcoming*: 2-3) claims truth to be finally epistemically valuable, though he plausible argues that this does not commit one to claiming truth to be *pro tanto* finally valuable simpliciter. One possibility he mentions is that from the non-epistemic point of view the only value of truth is practical value, and that value he suggests is a form of instrumental value.

following case from Shelly Kagan (1994: 311), based on a case first developed by Thomas Nagel (1979):

Imagine a man who dies contented, thinking he has achieved everything he wanted in life: his wife and family love him, he is a respected member of the community, and he has founded a successful business. Or so he thinks. In reality, however, he has been completely deceived: his wife cheated on him, his daughter and son were only nice to him so that they would be able to borrow the car, the other members of the community only pretended to respect him for the sake of the charitable contributions he sometimes made, and his business partner has been embezzling funds from the company which will soon go bankrupt.

Fred Feldman (2002: 299), juxtaposes this case with that of a cousin businessman, who thinks all the above about himself with the difference that in the cousin's case his beliefs are true.⁸⁹ If we were to compare the two lives, based on what we know of them, then it's intuitive to think that the cousin's life has gone better.

There is, however, difficulty in cashing out the intuition in play. The intuition that the second life has gone better may be a response to the fact of the cousin businessman's wife being faithful, and everything else described, rather than the cousin businessman's belief being true. If, however, we compare the first businessman case with a second in which the cousin businessman believes truly that his wife has cheated on him, and everything else described, then the intuition nonetheless arises, though perhaps without being very strong, that it is better to be the cousin businessman; this supports the thought that having a true belief is somehow valuable. To have this intuition may require bearing in mind that we can stay neutral on how true belief weighs up in an all things considered judgement against the likely accompanying pain of such true beliefs. As to how to cash out the value of true belief, Pritchard's (2006: 154-155) claim that the value of truth lies in the value of authenticity looks plausible. To support this claim Pritchard appeals to our intuitions as to whether it matters whether one is a brain in a vat or not. With this in mind he writes that even adding greater pleasure to the envatted, deceived self

⁸⁹ It is perhaps natural to think that in the second case the cousin would have knowledge rather than mere true belief, but regardless of that the value that the case highlights is plausibly one owing to true belief rather than knowledge.

would not make us think that that life is better. Such considerations lend themselves to the thought that like achievement, and requiring similar qualification, true belief is also constitutive of the good life. Admittedly, however, the considerations outlined are a somewhat sketchy attempt to cash out the intuition that true belief is valuable. What is important for my purposes here is just that there are grounds to think that true belief is valuable and to flag a worry, in addition to the problem raised by virtuous belief, with the way in which knowledge is defended as being distinctively valuable.⁹⁰

Returning to our original concern, if we take both virtuous belief and knowledge to be achievements, and taking true belief also to be valuable, then knowledge, entailing true belief as it does, is more valuable than virtuous belief. But this leaves unanswered the challenge of explaining the difference in value between knowledge and justified true belief; both whether knowledge is more valuable than that which falls short of it, and whether knowledge is distinctively valuable vis-à-vis that which falls short of knowledge.

Drawing on the thought that the value of knowledge should drop out of the nature of knowledge, the obvious place to look for an answer as to the difference in the value of knowledge vis-à-vis the value of justified true belief is to a difference in the nature of knowledge vis-à-vis the nature of justified true belief. As I argue in Chapter III, knowledge entails epistemic grace, justified true belief does not. Both involve beliefs that are justified, meaning that the believing agent has formed her belief in a way appropriate given her situation, and as such is in a position to gain knowledge. The difference lies in the knowing agent's virtuous belief enjoying epistemic grace; her virtuous belief being formed such that it is objectively appropriate given her environment, while this is untrue of the agent who has a justified true belief.

⁹⁰ A response to what I've outlined here might be to say that this doesn't show that true belief has prima facie or pro tanto value; rather what it shows is that some true beliefs, presumably those pertaining to important aspects of one's life, are intuitively valuable and are perhaps finally valuable, assuming authenticity is finally valuable.

Now it is from this difference in the nature of knowledge vis-à-vis justified true belief that a difference in the value of these goods emerges. Knowledge is the happy state of affairs of the cognitive agent forming her belief in such a way as to put her in a position to gain knowledge and her in fact gaining knowledge. The value of knowledge lies in the value of a good that requires the right kind of contribution from an agent but which cannot be secured by that contribution alone. In this knowledge is in a way similar in structure to eudaimonia or happiness and is analogous to some of the goods that are constitutive of the good life.⁹¹

On the Aristotelian conception of the good life virtue is a necessary but not a sufficient condition for eudaimonia. (Aristotle, 2009: 672-673). In other words, in order for an agent to enjoy eudaimonia, an appropriate contribution from the agent alone is not enough. The most valuable human good, in terms of content, eudaimonia, has the kind of structure that is most valuable. What can be achieved by virtue for Aristotle, and hence what is praiseworthy, isn't as valuable as that which requires virtue but for which virtue isn't sufficient. For Aristotle the latter is "blessed".

We praise the just and the brave person, for instance, and in general the good person and virtue, because of their actions and achievements.... If praise is for these sorts of things, then clearly for the best things there is no praise, but something greater and better. And indeed this is how it appears. For the gods and the most godlike of men are [not praised, but] congratulated for their blessedness and happiness. The same is true of goods; for we never praise happiness, as we praise justice, but we count it blessed, as something better and more godlike [than anything that is praised]. (Aristotle, 2009: 673). (Bracketed content is from the translated passage quoted).

While eudaimonia like knowledge requires virtue, like knowledge, it also requires more than this. But so far so like justified true belief; justified true belief too is more than just virtue. It is in goods of structure analogous to knowledge, that too are constitutive of the good, that we can appreciate the "blessedness" of knowledge and of happiness itself.

⁹¹ Eudaimonia is sometimes translated as happiness but the reader should be aware that Aristotle doesn't conceive of eudaimonia as some kind of subjective feeling. Eudaimonia is that which the good life consists in and is alternatively translated as flourishing. (Hursthouse: 2012).

Consider the case of friendship. An agent, let's call her Naira, may possess virtues that make her a potentially good friend; say she is funny, generous and understanding of those close to her.⁹² It may be the case that despite Naira's possession of virtues that make her a potentially good friend, she fails to enjoy the feeling of closeness that comes with being a good friend, supposing that she isn't friends with anyone. It may be, for whatever reason, that the agents in her environment are unresponsive to those virtues. Just as there are hostile epistemic environments in which, despite an agent believing virtuously, that agent may still fail to gain knowledge, so too may something analogous be the case when it comes to friendship.

In the case of friendship, it may be that though other agents are unresponsive to the agent's possession of virtues that make her a potentially good friend, she happens to enjoy with someone the feeling of closeness that is always a feature of friendship but which may also exist independent of friendship. Perhaps the person with whom she enjoys this feeling of closeness is a man who simply latches on to others. Even if Naira possesses virtues that make her a potentially good friend and enjoys the feeling of closeness that is a feature of friendship, intuitively, as in the case of an agent who has a justified true belief but not knowledge, even if what she has is valuable to some extent, she is missing out on something more and differently valuable.

Let's consider another case to tease out what this is. Consider the case of Tim. Tim gets a particularly desirable job. Tim has ability in the job relevant domain, though he doesn't get the job based on his ability but rather because of a factor orthogonal to the requirements of the job; say the employer was always going to give him the job no matter what because he owed his parents a favour. What is of note here is that Tim's getting the job doesn't seem as valuable in this case as it would have been if he had got the job based on his ability. In both cases the agents are relevantly virtuous and have something good, but their possession of these goods doesn't

⁹² Here I'm trying to avoid the arduous and diverting task of providing an account of friendship but if you hold that virtues like the ones described are necessary to be a potential friend rather than a potential good friend, then you may take what I write to be that – without it detracting from the import of what I write in these paragraphs.

require a contribution from them and in this way their possession of these goods is less valuable than could otherwise be the case.

The difference between knowledge and justified true belief is like this with a corresponding difference in value. On my account virtue is a precondition for epistemic grace distinguishing it from justified true belief. Knowledge, like goods constitutive of the good life, doesn't just involve the coinciding of virtue with another good/other goods, it involves having the good in part because of virtue. The value of knowledge, like the value of the good life, lies in the world rewarding the cognitive agent for her virtuousness rather than her simply getting goods independent of her efforts. With this in mind, note that getting a desirable job or developing a friendship aren't things that are just down to the contribution of an agent. An agent could have all the virtues of a potentially good friend without meeting someone who is sensitive to those virtues. By being virtuous in various respects an agent puts herself in a position to gain certain corresponding goods; it is by grace if she gains them.

What has been argued here, however, does not show that knowledge is distinctively valuable in terms of the kind of value it enjoys; that is, what has been argued here about the way that knowledge is valuable doesn't show that knowledge and no standing that falls short of knowledge is finally valuable.

§4.1.1 The Final Value of Knowledge?

In support of my criticism of the claim that knowledge has pro tanto final value I discuss a hypothetical case of knowledge, the Pierre case, and argue that the case is not one in which it is intuitive that knowledge is valuable, even on a pro tanto basis. The point made here, although originally directed at Greco's account, equally applies to my own. The point, if accepted, challenges the notion that an answer to the value of knowledge question will simply fall out of an answer to the nature of knowledge question, as some further explanation will be required as to when it has value.

The Pierre case is as follows:

Pierre sits in a café by the window looking out onto a relatively busy Parisian side-street. He decides to gain knowledge by counting the number of people who pass by his table on the street outside between every two sips of his coffee. He comes to know that five people passed between his first sip and his second sip, seven people passed between his second sip and his third sip, etc. Let us further add that there is no opportunity cost worth considering, he couldn't have been doing anything else that is valuable.

Greco's approach to knowledge is a virtue epistemic approach; in every case in which there is knowledge, there is an exercise of cognitive ability which falls into the broader category of epistemic virtue.⁹³ What seems mistaken about saying knowledge is valuable in the Pierre case is precisely the pointlessness of the knowledge. What seems to account for this lack of value is an appropriate relation between the knowledge gained and that which is virtuous simpliciter. So if we investigate that which it is all things considered vicious to investigate, then the resultant knowledge will not be valuable. An epistemic agent, like Pierre, is acting viciously, all things considered, when he is seeking knowledge of pointless truths.⁹⁴

If my diagnosis of the Pierre case is correct, then the difference between an instance of knowledge being valuable and not being valuable mirrors the potential difference between an agent exercising epistemic virtue, and on that basis enjoying epistemic grace, and an agent exercising all things considered virtue. When an agent exercises epistemic virtue and enjoys epistemic grace, but is not exercising all things considered virtue, as in the Pierre case, then it's not a case in which knowledge will be finally valuable.

This challenges the view that knowledge always has pro tanto final value. The original motivation for claiming knowledge to have such value was the claim that knowledge is an achievement and the good life is one that is rich in achievement,

⁹³ I've argued previously that we shouldn't think of a belief being virtuous if and only if it is formed from cognitive ability, rather a virtuous belief may also be formed from the exercise of a virtuous trait.

⁹⁴ The viciousness is in the manner in which he gains knowledge. If he were to gain knowledge of a pointless truth simply by having happened to look in a certain direction, then the viciousness in the manner of knowledge acquisition would be absent. In such a case, although the particular knowledge is pointless, value is still present owing to its form.

though the value of an achievement might in some cases be outweighed by the disvalue, say, of the cost of the achievement. The Pierre case, as a counterexample to the value of achievement thesis, pushes us to produce a more refined story of when an achievement is finally valuable. The more refined story offered is that an achievement may not have pro tanto final value in virtue of being constitutive of the good life if the kind of virtue required in the specific instance, epistemic virtue for example, is in that specific instance contrary to what is virtuous simpliciter. The same consideration likewise applies to epistemic grace which necessarily also requires the exercise of virtue. If the exercise of epistemic virtue is all things considered vicious, then as with achievement that is all things considered vicious, it will be in no way finally valuable.

Judging by the Pierre case, a case of pointless knowledge, and another type of case in the literature, immoral knowledge, objecting to the pro tanto move looks plausible. This looks natural in the case of achievements as they were offered as finally valuable in virtue of being constitutive of the good life and yet if they conflict with what is virtuous simpliciter then they conflict with a requirement for the good life. If attaining such a good ultimately is vicious, then, *ceteris paribus*, there is no reason to think that the particular instance of the good held by the agent may in any way be constitutive of the good life.

A good that is sometimes valuable in virtue of being constitutive of the good life and so finally valuable, in cases in which it's not constitutive of the good life, may not be finally valuable. In cases in which that good is constitutive of the good life, then that good is always finally valuable. As such, the intuition that knowledge is valuable is better cashed out as it being in the nature of knowledge to be valuable.^{[95] [96]}

⁹⁵ Saying that it is in the nature of knowledge to be valuable, given the criterion set for an achievement having final value, assumes that there are many ways in which an agent may be virtuous, which in any case is independently plausible when concerning how life is lived generally; this doesn't commit one to the less plausible claim, however, that there are many different ways an agent may be ethically virtuous.

⁹⁶ Given my first criticism it should be clear that what I've said here is not only relevant to knowledge but to justified belief or epistemically virtuous belief and any other epistemic standing that involves an epistemically virtuous belief.

We have an intuition that knowledge is valuable. We have that intuition, perhaps, because knowledge is the kind of thing that the virtuous agent has. But ultimately knowledge isn't valuable alone; it is knowledge that is virtuously held (simpliciter) that is finally valuable. That's why in cases in which knowledge and virtue simpliciter come apart, such as in the Pierre case, our intuition as to what is valuable doesn't track knowledge.⁹⁷

§4.1.2 The Distinctive Value of Knowledge?

On the account of the value of knowledge that I've provided, knowledge is not distinctively valuable vis-à-vis that which falls short of knowledge in the sense that knowledge enjoys final value while that which falls short of knowledge merely enjoys, say, instrumental value. Knowledge is distinctively valuable vis-à-vis that which falls short of knowledge, however, on the basis of the way in which it is finally valuable. I don't dispute that achievements are also finally valuable and that they are finally valuable in being constitutive of the good life. The difference, however, lies in their respective bases for being constitutive of the good life. Knowledge is constitutive of the good life in that epistemic grace is structured such that it is of a kind of good that is most valuable, whereas, say, a virtuous belief, although as an achievement is also constitutive of the good life, is not the kind of good that is most valuable.

The account provided is a neat one, both with regard to the nature of knowledge and the value of knowledge and the relationship between the two. It is one that equips us well to respond to Kvanvig's (1998) challenge to the propriety of affording the study of knowledge a central place in philosophy.⁹⁸ Knowledge, on the account I've

⁹⁷ An alternative way of motivating the move that is made on the back of the Pierre case is appealing to a revised conception of the unity of the virtues, not to determine what counts as virtuous in a given domain but to determine what has final value in virtue of being constitutive of the good life – a valuation that is naturally understood as neither being restricted to a particular domain, nor, with the revised conception of the unity of virtues in play, necessarily pro tanto in nature.

⁹⁸ In this paper Kvanvig couches the challenge as a challenge to epistemology, but it seems from this and subsequent papers, for example Kvanvig (2005), that the challenge is motivated by worries about the importance of the study of knowledge rather than epistemology as a whole. It should however be noted though that in this paper Kvanvig (1998) explicitly takes the “pre-eminence of knowledge” in philosophical inquiry to be the same thing as the “pre-eminence of epistemology”.

articulated, is different in nature and value than that which falls short of knowledge; knowledge really is special.

While the account of knowledge I have provided makes sense of the eminent place of the study of knowledge in philosophy, I don't take that account to make sense of the relative neglect in post-Gettier epistemology of certain other epistemic goods from study that pre-theoretical intuitions suggest are deserving of study from philosophers. The epistemic goods I have in mind here are understanding, though this is beginning to receive more attention, and wisdom.

The virtue turn in epistemology, however, has the potential to provide a framework that can act as a basis for locating prized epistemic goods such as understanding and wisdom. A virtue approach looks to be in a position to do so if we accept that it is plausible that understanding and wisdom require epistemic virtue for their possession. Requiring epistemic virtue not only puts us in a position to begin to identify the nature of these goods but also to identify their value, the exercising of epistemic virtue itself enjoying a certain value though intuitively the value of understanding and wisdom at the least exceed that value. Depending on our particular account of understanding and wisdom, the value that knowledge enjoys may also be present in the value of understanding and wisdom. While it's not obvious that understanding implies knowledge, for example if we think that understanding doesn't require factivity then understanding doesn't imply knowledge, it seems more plausible that wisdom does imply knowledge, albeit a restricted set of knowledge, perhaps for example knowledge of how to live well.⁹⁹

§4.2.0 Epistemic Value and the Epistemic Environment

So far I've argued that achievements and epistemic grace, in virtue of being constitutive of the good life, are finally valuable. I further argued that the basis on which epistemic grace is constitutive of the good life is such that epistemic grace is different and superior in value to that of achievement. This led me to the claim that

⁹⁹ For more on wisdom see Nozick (1989: 267-278) and Ryan (2008).

the value of knowledge, explicable in terms of the value of epistemic grace, is such as to warrant the attention it receives in philosophy. I observed that the value of knowledge is not such, however, to warrant the relative neglect of other prized epistemic goods. I noted that they too seem to require virtuous belief and if so their possession would bring its associated value – the value of achievement and highlight the fact that they may also imply knowledge and therefore its associated value – the value of epistemic grace. An upshot of the discussion is that the possession of various epistemic goods is in one way or another constitutive of the good life and that virtuous belief, if not a common requirement, for example if it turns out that true belief is also constitutive of the good life, is generally a requirement. Given that virtuous belief is generally a requirement of these various epistemic goods, and given that the possession of these goods is constitutive of the good life, we have a reason to want to have epistemically virtuous belief – at least that which does not conflict with what is virtuous simpliciter.

Whether an agent has virtuous belief depends on the way in which that agent forms her belief. But the epistemically virtuous agent may or may not be able to form virtuous beliefs about certain matters depending on her epistemic environment.¹⁰⁰ For example, an epistemically virtuous agent today is likely in a position to form much richer epistemically virtuous beliefs about Mars than an epistemically virtuous agent a century ago could have done. This is not to say that the agent of a century ago couldn't form epistemically virtuous beliefs about Mars, it's rather that even from their perspective, they'd have much fewer good epistemic grounds to form a range of such beliefs in comparison to an agent today potentially has. Similarly, depending on the direction intellectual exploration has taken and its consequent shaping of the epistemic environment, an agent may be in a position to form beliefs that are much richer all things considered virtuous than a counterpart in an alternative epistemic environment. For example, an epistemic environment in which an agent is in a position to form much richer epistemically virtuous beliefs about who she is,

¹⁰⁰ I define epistemic environment later in this section when I discuss the epistemic environment and epistemic environmentalism in detail. For now I take it that the reader is able to grasp its basic meaning. One thing to note, however, is that I'm using "environment" here in a broader sense than when I discuss environment in relation to epistemic grace.

how she ought to live, would, other things being equal, be better than an epistemic environment in which she is in a worse position with regard to forming such beliefs.

Given that the epistemic environment has this effect on our ability to form virtuous beliefs, and hence acquire other epistemic goods, and given the value that the goods in question enjoy, we have a reason to preserve or adapt our epistemic environment accordingly. I call this approach of seeking to preserve or adapt our epistemic environment based on the impact of that environment on our ability to attain valuable epistemic goods epistemic environmentalism. Given a concern with our ability to attain epistemic goods it is appropriate to investigate factors, both social and non-social, that have a bearing on our acquisition of those epistemic goods.¹⁰¹

There is an ethical dimension to epistemic environmentalism; that epistemic goods are valuable has ethical implications for practices or arrangements that prevent or facilitate their attainment. Simply put, if social arrangements are such that an agent is prevented from attaining a good that is constitutive of the good life, then, *prima facie* at least, those social arrangements are morally bad. In this respect the environmentalism component of epistemic environmentalism mirrors the form of non-epistemic, conventional environmentalism. There are facts about the environment and some of those facts may be taken to be relevant to existing ethical concerns or may inspire new ethical considerations. I use “epistemic environment” when referring solely to facts about the epistemic environment and “epistemic environmentalism” when referring to the approach, with its normative dimension, generally.

The environment component of epistemic environmentalism seems apt given the use of “environment” to describe that to which epistemic environmentalism is, in form, conceptually akin. Central features of our epistemic environment that influence knowledge acquisition are the interconnections and interdependencies that exist between different epistemic agents and agents and their physical environment. Like

¹⁰¹ Epistemic environmentalism is inspired by Alvin Goldman’s (1999) *Knowledge in a Social World* in which gaining true beliefs is taken to be desirable. In the book, Goldman considers social changes on the basis of our interest in truth.

the knitting together of the interconnections and interdependencies of various ecosystems with each other and with non-internal abiotic elements, such as sunlight and air, which provides our picture of the environment, the epistemic environment is formed, on the basis of interconnections and interdependences, by a similar kind of collage. “Environment” in its psychological sense, as is apparent from its definition, conveys a social dimension, one which is paralleled in epistemic environmentalism: “[T]he aggregate of social and cultural conditions that influence the life of an individual or community”. (“Environment”, Dictionary.com, from *Merriam-Webster's Medical Dictionary*). Insert “epistemic” prior to “life” into the definition and it expresses in a basic way a large part of what is meant by “epistemic environment”, albeit without making the relational aspects explicit.

While I’ve set out the approach and scope of epistemic environmentalism, I explore that approach specifically in relation to social epistemic factors that influence the attainment and dissemination of epistemic goods.¹⁰²

Good testimonial institutions in a community can have the effect of making an environment such that it is easy for a recipient of testimony to gain knowledge. While our perception is generally well adapted to the environment in which we form perceptual beliefs, the epistemic environment, in which we form testimonial beliefs, may be better or worse depending upon our testimonial institutions. Whether good testimonial institutions are in place in a community is not something that is down to an individual epistemic agent, though those institutions play a significant role in contributing to whether or not such an agent gains testimonial knowledge.

The practices of a community to which the individual is immersed may be such as to influence whether or not the individual is likely to form virtuous beliefs and also contribute to the individual’s prospects of enjoying grace. Using a religious analogy,

¹⁰² Social epistemology has drawn renewed attention from traditional epistemologists in the last two decades. Craig’s (1990) hypothetical genealogical account of the concept of knowledge has surely been influential in effecting this development. On Craig’s account knowledge originates as a product of social interaction and interests. The story he tells is of the concept of knowledge developing, via ancestor concepts, from a need, because of their importance, for a label to identify good informants. (Craig, 1990: 96-97).

if for example one holds that partaking of the sacraments, which the individual in a particular community does as a matter of course, is one way a person may be helped in receiving religious grace, then one begins to see an analogous way that community may be significant in an individual receiving grace in the epistemic case. Religious believers may hold that the absence of religious community may prevent or be a considerable obstacle for an individual to gain grace seems equally understandable when analogised to the epistemic case. It seems relatively obvious that past epistemic environments for example, which include testimonial institutions, played a significant role in the epistemic goods that individuals forming beliefs in those environments gained. It's plausible to think that whether an individual lives in a community that has structures for gaining, preserving, and disseminating knowledge and, if present, the quality of those structures will have a significant effect on the epistemic goods that they enjoy.

§4.3.0 Conclusion

In this chapter I have argued that it is in the nature of knowledge to be finally valuable in virtue of being constitutive of the good life. I have argued that the basis on which knowledge is finally valuable differs from the bases on which epistemic standings that fall short of knowledge are valuable. More specifically I have argued that knowledge has the value of epistemic grace while that which falls short of knowledge does not. I have further made the case that the value of epistemic goods has normative implications for epistemic arrangements that have a bearing on our attainment of epistemic goods. With this in mind I have outlined what I have called epistemic environmentalism. In the next chapter I examine the epistemic environmentalist approach in relation to trust, a key factor in the attainment and dissemination of valuable epistemic goods.

Chapter V: Trust

§5.0.0 Introduction

In Chapter III and Chapter IV I articulate and argue for an account of the nature and value of knowledge. The account of the nature of knowledge for which I argue is one that can explain cases of testimonial knowledge, including cases of testimonial knowledge in young children. I do so by departing from the claim common to both the accounts of Greco and Pritchard, that knowledge requires the exercising of ability, and instead defend the claim that knowledge requires epistemically virtuous belief, whether that be from the exercise of an epistemically virtuous trait or the exercise of cognitive ability. In the case of testimonial knowledge in young children I argue that the epistemically virtuous trait, the exercising of which allows belief to be a candidate for knowledge, is trust. In my chapter on the value of knowledge, I outlined how my account of the value of knowledge is based on my account of the nature of knowledge. I drew attention to the fact that factors in the epistemic environment, including factors that have a bearing on the epistemic appropriateness of trust relations, have an effect on the attainment and/or holding of knowledge and hence on the attainment and/or holding of the value of knowledge. Accepting that knowledge and other epistemic goods are valuable, we have a reason, though not necessarily a decisive one, to preserve or alter our epistemic environment depending on whether such preservation or alteration will facilitate the holding or attainment of knowledge.

In this chapter my thesis begins the move from establishing the value of epistemic goods to spelling out the implications that their value has for our social arrangements. Rather than trying to spell out such implications in this chapter, I confine myself here to an examination of an aspect of the epistemic environment. More specifically I examine certain types of relationships from which, if trust is present and the trust is

appropriate, much knowledge may be held or attained. I begin by providing an analysis of trust because of its significance to the relationships being examined in particular and the epistemic environment in general. Following my goodwill analysis of trust, I examine the type of relationships described. They are the non-expert – expert relationship, and the relationship between the occupant of a position of epistemic authority and the non-occupant. I argue that similar to intuitive cases of testimonial knowledge in young children, it is the exercising of a virtuous trait, trust, that helps explain how testimonial knowledge may be gained from the testimony of experts and those occupying positions of epistemic authority. Finally, I examine factors, that may either be directly present in such a relationship or that are part of the broader epistemic environment, that bear on the epistemic appropriateness of such relationships with consequent effects on the attaining and/or holding of knowledge and the enjoying of the final value of knowledge. In the next chapter I consider how intervention in the epistemic environment, in order to hold or attain the value of epistemic goods, might be justified or unjustified.

§5.1.0 An Analysis of Trust: Introduction

The topic of trust involves two central aspects: trusting and trustworthiness. I concentrate my attention on examining trusting and ultimately explain trustworthiness by way of my account of trusting.¹⁰³ Before proceeding with an analysis of trust, there are a number of useful distinctions that I believe it is helpful to lay out.

As a basis for my analysis of trust I begin by examining the requirements that McLeod (2011) lists and discusses for trusting. Building on those requirements and motivated by Potter's sexist employer case, I argue for a goodwill analysis of trust. Having so argued, I highlight an attractive feature of my overall account of trust; that

¹⁰³ I take trusting to primarily be a psychological phenomenon in that whether trusting is taking place or not depends on what is going on in the mind of an agent. As such, a philosophical account of trusting, aside from with regard to its logical coherence, can legitimately be assessed according to how it captures our experience of and intuitions with regard to trusting; although such an account should be consistent with psychological and neurological findings.

based on my particular analysis of trust, sense can be made of the appropriateness of particular reactive attitudes that may follow a perceived breach of trust.

§5.1.1 The Nature of Trust: Distinctions

I begin by setting out a series of distinctions that I believe serve to clarify the target for analysis and are generally informative of the nature of trust. Many of the distinctions are ones to be found in the literature on trust. The work of Carolyn McLeod (2011) is a particularly rich source in this regard.

The act of trusting versus a trusting character

The type of trust that is of interest for my purposes is the trust involved in the act of an agent **trusting** another agent with respect to some matter.¹⁰⁴ It should be noted, however, that there is an everyday sense of trusting according to which it does not follow that the trusting agent has a relationship of trusting to a particular agent or a range of agents. The sense in question, rather, may be used to describe an agent's character. This sense of trusting conveys the idea that an agent's character is such that she is disposed to trust. In what follows, and as is in keeping with the literature, see McLeod (2011), I seek to provide an account of the nature of the act of trusting, rather than the characteristic of being a trusting person.¹⁰⁵

Trust versus reliance

Roughly speaking, vulnerability to betrayal demarks trust from nearby notions like reliance. (McLeod, 2011). A perception on the part of the trusting agent that his trust has been breached won't necessarily lead to a feeling of betrayal, but it's the possibility that helps us distinguish trust from nearby notions. Later in my account I argue that whether there is a feeling of betrayal depends on the way trust is perceived to have been breached. I further claim later in this chapter that whether a perceived breach in trust does lead to betrayal will depend on which of the necessary conditions

¹⁰⁴ A trusting relationship may also, however, be a reflexive relationship; that is, someone may trust themselves with respect to some matter.

¹⁰⁵ Part of the way through this chapter, when I discuss testimonial knowledge in children, I consider how an account of the nature of the act of trusting can help us understand trusting as a characteristic.

for trust is violated. To see how this vulnerability to betrayal follows from a trust relation and not other nearby alternatives, consider the following case:

Eusebio depends on his alarm clock to wake him up at seven each morning. On a morning on which it is particularly important that Eusebio gets up at seven, he has to catch a morning flight, his alarm clock, because of running out of battery, fails to go off.

It would not seem appropriate for Eusebio to feel betrayed, though he might perhaps feel annoyed or irritated. It further seems right to say that Eusebio doesn't trust his alarm clock but merely relies on it. Although this is only one particular case, the thought is that in cases in which the relation is such that there is no possibility of there being an appropriate feeling of betrayal, then there is no trust relation and vice-versa. For the claim to be shown to be wrong one could attempt to provide a case in which there is intuitively trust but no possibility of betrayal or a case in which there is a possibility of betrayal but no trust.

Why is it that Eusebio relies on rather than trusts his alarm clock? For one thing an alarm clock is an object that Eusebio perceives as such. Intuitively perceived objects cannot be the objects of trust. Again this is a claim that can be tested by consideration of possible cases. In §5.1.2 I consider such a possible case but reject it as a counterexample to this claim.

It's not, however, just relations with perceived objects that rule out appropriate trust relations. A relation between one agent and another may be such that it is clear that there is mere reliance rather than trust. As Baier (1986) describes, Kant's neighbours could famously set their clocks by Kant's daily walks. We can say that they could rely on him in that regard. Nevertheless, were Kant to break with his routine for some reason, we wouldn't think it appropriate for his neighbours to feel betrayed. That there is not the possibility in this case of appropriate feelings of betrayal marks it out as one in which there is not a trust relation involved and highlights the fact that trust requires more than mere reliance. What it is that rules out Kant's neighbours from having appropriate trust relations in the case described will later be explored.

Trust: Conscious and unconscious

An agent may or may not be conscious of trusting another agent. (Baier, 1986). The claim that an agent may be conscious of trusting seems sufficiently obvious such that it doesn't need defending. The claim that an agent may unconsciously trust another agent looks plausible if we cast our mind to a possible case in which an agent becomes aware that she feels let down or betrayed by another agent's behaviour. Intuitively there seems nothing amiss in considering that in such a case the agent may have unconsciously trusted the other agent.

Similarly, an agent may be conscious of being trusted or be unconscious of being trusted. Again, the claim that one may be conscious of being trusted seems obvious enough not to need defending. It also seems right to say that in some cases an agent may not actually be conscious of being trusted. Consider the following case;

John, Peter and Brian are in a cabin in a forest and Brian is dangerously ill; he is slipping in and out of consciousness. Peter and Brian had a serious falling out before Brian became dangerously ill. When Peter announces that he is going to go to get help, John lets Peter know that he doesn't trust him to do so. Both are unaware of the fact that Brian is conscious at that moment, and hearing what Peter has to say, Brian does trust that he will try to get help. Nevertheless Peter is unconscious of being trusted.

Trust: Voluntary and involuntary

As previously discussed an agent may simply find that they trust another; this suggests that trust, at least sometimes, occurs involuntarily. Whether trust in every case is involuntary is more contentious. Let's first consider a case. Suppose a tourist in an alien culture has his wallet and passport stolen. A stranger, seeing what has happened, offers to drive him to a police station. We can imagine the tourist considering the offer and his situation and we might be tempted to say that, thinking he has no good alternative, he might decide to trust the stranger. The defender of the view that trusting can only be involuntary may make the diagnosis that the agent in the case described is simply acting as though he trusts, and thereby continue to reject the possibility that an agent may simply trust at will. Whether we think there can be

such a decision to trust may hinge on whether we think that an agent by acting as though he trusts another agent, can be said to trust that agent. (Faulkner, 2007).

One thing that makes that claim contentious is what simply being able to choose whether to trust seems to imply for agents like us. If voluntary trusting is possible and if we receive testimony that x, then we can simply decide whether to trust the testifier and so we can simply decide whether to believe her testimony. But now it looks like we have a route to believing at will, at least in some cases. It seems plausible that the voluntary/involuntariness of trust should be understood on a similar model to the voluntary/involuntariness of belief; that we can do certain things to try to bring about trust/belief but that normal human agents can't make themselves believe/trust by a mere act of their own will.¹⁰⁶

§5.1.2 The Nature of Trust: Requirements

Having made some basic distinctions, let's turn to the requirements of trust. Carolyn McLeod (2011) sets out the following requirements for trusting as being "relatively uncontroversial":

Trusting requires that we can,

- 1) be vulnerable to others (vulnerable to betrayal in particular);
- 2) think well of others, at least in certain domains; and
- 3) be optimistic that they are, or at least will be, competent in certain respects.

The vulnerability that McLeod has in mind is minimally, or at best, that whatever an agent trusts the trusted person to do might not get done. In other words, while an agent is trusting, rather than being certain or ensuring, that S will do p, that agent is vulnerable to the trusted agent not doing p. That an agent is vulnerable to the trusted agent not doing p, leaves that agent vulnerable to being let down or betrayed. If that

¹⁰⁶ There may be odd cases in which human agents are not normal human agents and can believe almost at will. For example, consider the protagonist in the film *Memento* who suffers from amnesia. The protagonist manipulates his own beliefs by writing notes, the writing of which will be forgotten and the contents of which will be believed when read.

vulnerability were somehow removed, then so would the trust be removed, at least roughly speaking.

The thought that we have to be actually vulnerable to either being let down or betrayed is surely too strong. By “actually vulnerable”, I mean that it has to be the case that S really might not do p or we really could be betrayed by S. Given that whether or not a trusting occurs depends on what is going on in the mind of an agent, whether we’re actually vulnerable or not is irrelevant.

Consider a case that intuitively supports such a judgement: R trusts S to do p. It turns out, unbeknownst to R, that S is actually a highly advanced android that can’t but do p. In trusting S, R is never actually vulnerable, though he may perceive himself to be. Nonetheless, we shouldn’t take from this that an agent has to perceive or be disposed to perceive himself to be vulnerable to having his trust betrayed or let down, and that without such belief a trusting cannot occur. This is obvious if we accept that unconscious trusting is possible.

Skipping over the second requirement for the time being, the third requirement that McLeod lists relates to how the trusting agent regards the trusted person. Putting it more specifically than McLeod does, the requirement that R can’t be said to trust someone, say T, to Φ , if he does not believe, or is not disposed to believe, that T has the competence to Φ seems right.¹⁰⁷ I can’t be said to trust you to drive me to my important meeting if I don’t believe that you can drive. The possibility that T need not be presently competent also seems right. Parents may trust that a child will look after them in their old age without believing or being disposed to believe that the child is currently competent to do so. R’s belief in the current or future competence of T, in addition to R being vulnerable to others, doesn’t imply that R trusts T. It is plausible to think that R must stand in some further relation to T in order to be said to trust T.

¹⁰⁷ A trusting agent’s belief or disposition to believe that T has the competence to Φ may be based on the trusting agent’s belief that T has competence X which encompasses the ability to Φ and so implies an ability to Φ .

The point of the second requirement, that the trustor must “think well of others, at least in certain domains”, is not immediately obvious. McLeod (2011) writes that paradigmatically “trust involves being optimistic” that the trusted person will do something for us. The idea, as it is presented, is that if someone is easily suspicious of others, then that suspicion can forestall trust relations from forming. However, what is necessary for trusting is not just that there be an absence of suspicion in the way mentioned, but that candidates of trust are well regarded in a certain respect. This is in fact what is suggested by how the requirement is formulated. Although this requirement may be one of a set of requirements that are “relatively uncontroversial” in the discourse, more needs to be said about precisely what thinking well of others, or being optimistic regarding their actions, means.

§5.1.3 The Nature of Trust: The Goodwill View

There is disagreement as to how far a belief in the commitment of the trusted agent to do whatever it is that s/he is trusted to do goes in terms of satisfying what is required for trust. (McLeod, 2011). Intuitive responses to the following case provided by Potter (2002: 5, cited in and adapted somewhat by McLeod, 2011) support the view that the way in which an agent is believed to be committed is also relevant:

If a sexist employer could avoid hiring or get away with firing female employees then he would do so. In fact, he treats female employees “well”, but only out of fear of legal action if he does not. We can say that he is committed to treating his female employees well because he is committed to not risking having legal proceedings brought against him.

It looks like we may be able to say that he can be relied on to treat his female employees “well”, but intuitively it seems mistaken to say that, knowing the reasons for his behaviour, he can be appropriately trusted to treat his female employees well. (Or that, knowing the reasons for his behaviour, it would seem inappropriate for a female employee to feel betrayed rather than say aggrieved if his sexism was to assert itself in a particular instance.)

Below I set out a goodwill account of trust, based on Annette Baier's (1986) account, which I believe can best make sense of the way in which the trusting agent believes the trusted to be committed.¹⁰⁸ Let me begin by making clear what is meant by "goodwill", and how goodwill relates to the listed requirements. The "good" in "goodwill" need not mean a moral will, nor a will perceived as moral by the trusting agent; rather the sense of good involved is that the will is regarded as good given some end. For example, a bank robber may trust his accomplice to be waiting outside the bank in a getaway car. In doing so it seems obvious that he need not be supposing that the will of his friend is good in any moral sense.

The "will" in "goodwill" indicates that a trusting agent may trust a perceived agent. If someone said that they trusted their car, rather than relied on it for example, then intuitively we would have grounds to think that they were misusing "trust" or perhaps that they were making strange and false attributions to their car. While it's true that sometimes people might say that they trust their car, the thought is that if they were to speak carefully, then they would not do so any more than they would refer to their car as she or he, or generally make attributions to their car suggestive of the perception of agency. Furthermore, there's no conceivable scenario in which we would say that a trusting agent can be betrayed by his or her car; betraying being only something that can be done by an agent.¹⁰⁹

The goodwill view of trust helps make sense of the vulnerability that accompanies appropriate trusting. Betrayal may result when the trusting agent discovers that the will of the trusted, that the trusting agent believed to be good given some end, is actually ill, or indifferent to that end, with regard to the end to which the trusting agent trusted the other agent; alternatively, the discovery, for example, that the trusted agent's will is not good given some end may result in the trusting agent judging her trust to have been unwarranted or inappropriate.

¹⁰⁸ Baier (1986) first put forward what McLeod (2011) calls a "goodwill" view in "Trust and Antitrust".

¹⁰⁹ All the same it seems appropriate to interpret agent liberally enough to include non-human actors states and corporations.

This thought can be fleshed out with the use of an example. We can imagine that a feeling of betrayal would result in the case of a politician who campaigned on issue X but when in a position to act with regard to issue X demonstrated no interest in doing so or showed himself to actually be opposed to acting with regard to issue X. Alternatively, we can imagine a person who had believed that a politician had campaigned on issue X but who later comes to realise that his position was more nuanced not feeling betrayed when the politician when in a position to act with regard to issue X demonstrates no interest in doing so or shows himself to actually be opposed to doing so.

Just as a change from trusting and therefore meeting the goodwill requirement to not meeting the goodwill requirement and therefore no longer trusting is linked to the reactive attitude of betrayal that accompanies appropriate trust, so too is trusting related to the competence requirement. The third requirement, the competence requirement, in cases in which there is no longer trust because that requirement is no longer satisfied, and the trusting agent doesn't blame herself for her belief in the competence of the trusted agent, can make sense of how trust is linked to the reactive attitude of feeling let down. A change in belief to the effect that the trusted agent is not sufficiently competent in the relevant respect will end the trusting agent's trusting of the agent in that respect but will not result in feelings of betrayal, but in the feeling of being let down. As with the case of the politician with regard to goodwill, if the trustor comes to believe that she was at fault for believing the agent to be competent, then the agent will not feel let down by the trusted agent.

Consider the following case:

John says he will help Vera with an upcoming tutorial assignment and Vera trusts him to help her with the assignment. Vera knows John to have received good grades for previous assignments and John has confidently assured her that he will be able to help with this assignment. The night before the assignment is due, although Vera can see that John is trying to help her with the assignment, the goodwill condition is satisfied, it turns out that John is lacking the necessary competence to help with this particular assignment.

In such a case, Vera, if reacting appropriately will not feel betrayed; instead, John's lack of competence to perform the action which he has been trusted to perform, will result in Vera feeling, if reacting appropriately, let down.¹¹⁰

Given the connections identified between the competence requirement and the goodwill requirement to the first requirement that McLeod identifies, it is plausible to see the competence and goodwill requirement as necessary and jointly sufficient for trusting to take place and to see the vulnerability that accompanies trust as a consequence from those requirements in cases of appropriate trusting rather than being itself a necessary condition for trusting. In so doing the two conditions also make sense of why vulnerability to the reactive attitudes of betrayal and feeling let down accompany trust. Furthermore, there is an easy and plausible application of this account of trust to the two virtues of the good testifier, identified by Bernard Williams (2002), accuracy and sincerity. Accuracy is neatly covered by competence and sincerity by goodwill. The picture we're left with of the trustworthy testifier is one that falls out of the account of the nature of trusting.

As previously stated, the topic of trust is relevant given that on the basis of exercising trust, when trusting is a virtuous trait, knowledge may be gained. The relationship between children and caregivers has been discussed in this regard. The requirement that an agent has certain beliefs or dispositional beliefs about matters of competence and will might be thought to raise the bar too high for children to qualify as engaging in trusting, and so preclude them from knowing on the basis of trust. The key here, however, is that the beliefs or dispositional beliefs need not be highly refined. If the four year-old Sarah is disposed to believe that her care-givers are omnipotent or something close to it, something four year-olds seem to typically be disposed to believe about their caregivers, and that their will is good with regard to

¹¹⁰ If the case were different, if John had presented himself as being able to help with the assignment when in fact he knew that he didn't have the requisite competence to do so, then it would be natural to think that betrayal would be the appropriate reactive attitude for Vera to have if she were to find out. On my account, this would in turn require the assumption that Vera trust John, consciously or otherwise, not to be dishonest about such matters; a natural assumption to make given the description of the case.

teaching them about the world, then the disposition to believe that her mother has the competence and the will to tell her where Lucky is will be encompassed.

Of further interest, given a concern for the epistemic environment, is the relationship between expert and non-expert and those in a public position of epistemic authority and those not. I will now turn to what the relevant features of those relationships are and what bearing those features have on trusting in those relationships and therefore on the knowledge that can be gained from that relationship.

§5.2.0 Trust: Expert, Epistemic Authorities and the Layperson

This section examines trust relations between non-experts and experts and, later, trust relations between non-occupants of positions of epistemic authority and occupants of positions of epistemic authority. The goal here is to describe features typical of such trust relations and to set out the epistemic benefits that such trust relations may have.

First, I discuss the role of the expert and how best the expert should be defined. Next I examine non-expert – expert relations drawing on the work of Hardwig. In particular, I focus on Hardwig’s claim that the expert’s testimony is necessary in order for us to have knowledge in expert domains. I go beyond Hardwig’s own claims, however, by pointing out that not only are we dependent on the testimony of experts for such knowledge but that in many cases we are also dependent on the transmitters of that testimony, such as the media, and media relations centres at universities.

§5.2.1 Identifying Experts

It is tempting, particularly when thinking of the designation of “expert” as a status, just to say that an expert is a person recognised as such. On the other hand, however, it is conceivable that a person could have advanced technical and theoretical knowledge, though go unrecognised as an expert. That this is so challenges understanding “expert” in terms of recognition. Such a person’s knowledge of a

subject may have the depth and breath equivalent to or exceeding that of a publically recognised expert.

An expert who is not recognised as an expert will not be in a position of epistemic power over others because people will not be dependent on her testimony as an expert for their beliefs about the world. Given that her testimony is not recognised as expert, it can be expected that her testimony will not have the same potential to influence the epistemic environment. Obviously the non-recognised expert may have localised influence. Her acquaintances may defer to her on her subject of expertise as their knowledgeable acquaintance on those matters.

§5.2.2 Epistemic Dependence

John Hardwig (1985) argues that the non-expert is dependent on experts for much of what she takes herself to know and is dependent on experts for knowledge that lies in expert domains. He further argues that this dependence is rationally appropriate. He writes that if a non-expert recognises a putative expert as an expert, then that non-expert is generally rationally committed to accepting the views of the expert. The non-expert can imagine alternatives and raise relevant objections but:

...only someone with A's expertise can make an accurate assessment of the value and validity of the objection or alternative. Under cross-examination by the layman, the expert may admit the cogency of a given point, but he (and his fellow experts) must judge whether it is cogent and germane, since they are the only ones who fully understand what is involved in the methods, techniques, premises, and bases of the expert's training and inquiry and how these affect the resultant belief. (Hardwig, 1985: 342).

Hardwig (1985: 342) acknowledges that the non-expert may rationally refuse to accept an expert's views on the basis of worries related to the particular expert testifying, for example with regard to the sincerity of her testimony, and the non-expert may be correct in doing so, but this itself, Hardwig claims, will not yield knowledge in the expert's domain. If we're to gain knowledge in an expert domain, we have to trust the testimony of the expert.

Hardwig's claim here seems a little strong. It doesn't seem to be necessarily true that non-experts can never have knowledge in expert domains without depending on experts, although it does seem correct to say that ordinarily this will be the case. The thought that there may be exceptions is motivated by a possible though outlandish case, hence it doesn't threaten Hardwig's claim when it comes to ordinary cases. Suppose John is an expert in a particular domain. A question of the structure is p true or not true is posed. There are only two possible answers, now suppose that there is an evil demon who always makes it so that John when reporting on questions of this type will get it wrong and we know that this is so. In such a case we are not dependent on believing the testimony of the expert, John, in order to be said to know. In fact we can draw a justified inference and not believe him and get it right. Nevertheless it is worth noting that the layperson's knowledge of the expert domain is gleaned in a way such that the expert still plays a crucial role. However, this dependence does differ from the dependence that Hardwig claims characterises knowledge within expert domains – that is, we can't have knowledge in expert domains without deferring to the testimony of experts.

Goldman (2001) argues that our dependence on experts is not completely blind. He makes his case by arguing that the non-expert can rationally discriminate between believing the testimony of one putative expert and another in a particular domain. One particularly important means of doing this is judging by past **track record**. Non-experts can judge experts' past track records if experts have made claims with a predictive aspect. Making claims with a predictive aspect may allow the non-expert to see if the expert has gotten things right in the past and use that as a basis as to whether to believe the expert.

Although this approach helps non-experts with the assessment of some expert claims, and so whether to trust some experts, it is an approach that requires a significant investment of time on the part of the non-expert per expert claim. It is therefore unlikely to be a practical way for a non-expert to respond to the variety of expert claim that they encounter. Goldman's conclusion reflects the difficulty faced by the non-expert even when she attempts to make an informed assessment of an expert's

claim. “Establishing experts' track-records is not beyond the pale of possibility, or even feasibility... There is no denying, however, that the epistemic situations facing novices are often daunting. There are interesting theoretical questions in the analysis of such situations, and they pose interesting practical challenges for "applied" social epistemology.” However, the non-expert remains dependent on experts for knowledge of a certain type. She can only gain knowledge of that type by trusting experts. Given this epistemic dependence of the non-expert on the expert, and the trust required if the non-expert is to gain knowledge of a particular type, the non-expert who seeks knowledge of this type is in a position of vulnerability vis-à-vis the expert.

Hardwig (1991: 705) claims that non-experts can't simply presume that experts have decisive prudential reasons to provide accurate and sincere testimony; a result of which is that non-experts cannot simply suppose the reliability of their testimony. He argues that detection methods, peer review and replication, for bad (dishonest or sloppy) scientific research are flawed, and that the consequences of being found out are often not very serious. (Hardwig, 1991).

In support of his claims Hardwig (1991: 702-703, 705) cites (a) M Davis (1989), in an unpublished PhD, as finding, according to a survey conducted, that nearly 25% of respondents from various academic fields claimed to personally know someone to have falsified data and (b) J. P. Tangney (1987), for the claim that, according to survey results published in *New Scientist*, scientists do not believe that proven fraud would result in even temporary exclusion from research. Hardwig uses this research to bolster his claim that the alternatives facing the non-expert are often either trusting or remaining ignorant.¹¹¹

Other research on this topic suggests that the problem isn't as bad as Hardwig suggests. In 2005, *Nature* published a survey of scientists in which only 0.3% of

¹¹¹ Hardwig considers the possibility of increasing the severity of punishment for publishing dishonest or generally unreliable results. He worries, however, that such an approach, by trying to minimise the role of trust, risks creating a suspicious research atmosphere which is not conducive to the growth of knowledge. Instead he advocates encouraging the development of ethical researchers, who he believes will be trustworthy.

anonymous respondents admitted that they had personally falsified data within the previous 3 years, though 6% admitted that they did not present data if it contradicted their earlier research. (Martinson, Anderson, de Vries 2005.) A meta-analysis of research into the percentages of scientists involved in the fabrication and falsification of scientific results found that “[a] pooled weighted average of 1.97%... of scientists admitted to have fabricated, falsified or modified data or results at least once”. 12.34% claimed that colleagues had engaged in the “falsification” or “fabrication” of data. (Fanelli 2009.) This research does offer some reason to worry about the quality of expert testimony, though not to the extent that Hardwig seems to think. At the very least though, given concerns with both sincerity and accuracy and how research is generally carried out, it seems right to say that we cannot simply rely on expert testimony for knowledge, rather we have to trust it if we’re to gain the related knowledge.

A more worrying aspect of the epistemic environment and our attainment or holding of epistemic goods of value than the practices of experts and our dependence on them for a particular type of knowledge is the process by which expert testimony is generally communicated to the non-expert. Failure, either to correctly identify expertise, and so include quacks or exclude genuine experts, or to accurately communicate expert testimony, risks polluting this part of the epistemic environment and undermining the trust of the non-expert in the expert, and encouraging scepticism of testimony offered. A study based on an analysis of a random 200 press releases from 20 medical research centres found that 58% of press releases failed to include relevant caveats and cautions regarding methods used and the results of the studies. (Steven Woloshin et al, 2009). This serves to highlight the fact that not only is the agent in a position of epistemic dependence vis-à-vis the expert, she is also dependent on those involved in the process of communicating the testimony of the expert. While in some cases she can perhaps locate and read the research for herself, practically and plausibly in principle in cases in which the language of the research is too complicated for the layperson she is dependent.

There is a dependence analogous to that of non-expert – expert dependence in the relationship between non-experts and those in positions of epistemic authority. A person in a position of epistemic authority need not be such that they have any special training or methodological knowledge that separates them from others in the way the expert and non-expert are separated. Rather their position is such as to be privileged with certain information in a manner that means that if others want to gain knowledge on the basis of that information then they will be dependent on those in positions of epistemic authority in a like way to how they are dependent on experts. Figures in positions of epistemic authority may include politicians in receipt of confidential intelligence, judges with knowledge of inadmissible evidence, school principals with knowledge of the home lives of some of their students. For those in positions of epistemic authority, their authority will be limited to a particular domain and that they have epistemic authority in that domain will be owing to the position they hold. When they make claims about areas of which others don't enjoy privileged information, others will be in a relationship of trust, or mistrust, vis-à-vis occupants of positions of epistemic authority. We won't know that they don't have information that we don't and they will be able to press their privileged position, with the access to information that that involves, as a reason for action. Here the example of elected leaders claiming that they had solid evidence that Iraq had "weapons of mass destruction" springs to mind; the person not occupying that position could do little to contradict such a claim. If the politician is speaking from a position of epistemic authority then we just aren't in a position to evaluate the relevant evidence and come to an independent conclusion. Like in the case of the expert, however, we might have reason to distrust her motives or we might perhaps see an error in her reasoning.

§5.3.0 Conclusion

In this chapter I provided a goodwill analysis of trust. In doing so I further elucidated what is involved in typical cases of testimonial knowledge in young children and an important aspect of the epistemic environment. I also examined the epistemic dependence of the non-expert on the expert. I generally supported Hardwig's

argument that the non-expert is epistemic dependent on the expert, but I also highlighted the non-expert's practical dependence on those who communicate expert testimony. I further claimed there to be a like epistemic dependence in the relationship between the occupant of a position of epistemic authority and those not in that position. Over the course of the second half of this chapter I flagged concerns with the trustworthiness of those on whom we are dependent for knowledge. The point here was to highlight a way in which our epistemic environment may be better or worse. In the next chapter I examine what we should take into account when considering whether to intervene in the epistemic environment on the basis of regarding an epistemic good or epistemic goods to be valuable.

Chapter VI: Epistemic Interventionism

§6.0.0 Introduction

The aim of this chapter is to provide a lay of the land for work in applied social epistemology. I pose a relatively abstract question about whether the state should assist a citizen in attaining an epistemic good. I fill in details as to the kind of value such a good may have and what that epistemic good may be. I discuss significant ways in which state intervention to promote epistemic goods differs from the promotion of material goods. Considering one justification for supposing epistemic intervention to be appropriate, on the basis of a good having final value in virtue of it being constitutive of the good life, I explore how such a justification fits or not with perfectionist, value pluralist, and value subjectivist theories of the good. Finally I examine various forms of paternalism and consider what justificatory support they might offer to an affirmative answer, or negative answer, to the question posed near the beginning of my chapter. I do likewise with legal moralism and liberalism.

§6.0.1 Background

This chapter has been inspired by reflection upon the significance of **social epistemology** and **epistemic value theory** for one another. I take social epistemology roughly to be the study of the influence of other agents, social institutions, and social arrangements, on the epistemic standings of agents.¹¹² This epistemically significant social dimension, is part of the epistemic environment. Epistemic value theory examines the worth of epistemic standings. Findings from

¹¹² Goldman's (1999: 4) term "social paths" approximately conveys the kind of influencing factors I have in mind here, although he describes them as paths to knowledge whereas I leave this open. In the same passage he claims that social epistemology, in contrast to traditional epistemology, may take as its focus groups of believers, such as a team of co-workers. The possibility or otherwise of group agents does not concern me here.

epistemic value theory can be used as a basis for judging the felicity, or otherwise, of epistemic environments and thereby inform what has previously been described as epistemic environmentalism.¹¹³ For example, if knowledge is valuable and if because of the epistemic environment being a certain way less knowledge is possessed than in alternative epistemic environments, then, *ceteris paribus*, that epistemic environment is worse than the alternative epistemic environments. This suggests how viewing the epistemic environment through the prism of epistemic value theory might motivate intervention in that environment, so setting the ball rolling for applied social epistemology.

§6.1.0 The Question

Question (henceforth Q): Should S assist R in attaining epistemic good x?

Let us specify that S is a state, contemporary and democratic, and that R is a citizen of that state. Other possible specifications are also undoubtedly worthy of exploration, but it is the state-citizen relation that is of interest to me here. If there is to be a safe keeper or manager of the epistemic environment, a shared space – the condition of which interests are bound up with, then a natural candidate for that role is the state, and addressing this broader issue of whether this is an appropriate role for the state to play is also one of my concerns.

I write “assist” for want of a word or term that better captures what I mean. What I have in mind with Q is the role the state should play or not in R ending up with or not ending up with a certain good. When I say “assist” in Q, I want to leave open whether that means S making for friendly conditions with the result that R has a better chance of attaining x, S helping R directly to attain x, or S attempting to ensure that R ends up attaining x. With regard to the first and third cases, I want “assist” to

¹¹³ Although I’m not concerned with examining this particular relationship, it seems natural to think that the value of an epistemic standing may be influenced by its social significance. For example, knowledge might derive value, beyond that of standings that fall short of knowledge, from the licence for the kind of action that its possession grants. (Miranda Fricker’s (2007) work on epistemic injustice might be categorised as falling into this approach of seeing the value of an epistemic good, Fricker’s focus is knowledge, in part at least by way of its social significance.)

be neutral as to whether R wants x, is indifferent to x or doesn't want x. I also want to leave open the possibility that S could encourage R to want to attain x. My neutrality here anticipates the neutrality I adopt with regard to the interventionist and non-interventionist positions which I consider in the next section.

The epistemic good in question may be any of the following: true belief, true belief that is of interest to us, the exercise of virtue, justified true belief, knowledge, understanding, or, perhaps, wisdom. I won't specify x here as one or several of the epistemic goods mentioned. I want to stay neutral here on both which epistemic goods might be more worthy of promotion than others and what values those goods might actually have.

A claim as to the value of an epistemic good may be a claim that the epistemic good has prima facie or pro tanto value rather than all things considered value, while the kind of value may be, say, or final.¹¹⁴ An all things considered account would say what value the good has after external factors that influence overall value are taken into consideration.

A good may be valuable in more than one way. Knowledge might, for example, be finally valuable but also instrumentally valuable for some end. One epistemic good may be valuable in a way that another epistemic good isn't, or it may have the same kind of value but to a different degree. For example, mere true belief and justified true belief may both be instrumentally valuable but one good may be more instrumentally valuable than the other. The thought in this particular case could be that it's more instrumentally valuable to have a justified true belief for end A, because generally a justified true belief will stay in place better than a mere true belief. So, if you have a justified true belief rather than a mere true belief you might be better placed to satisfy end A.

In this chapter, I leave open whether x is a single instance of an epistemic good or a collection of instances of an epistemic good. It should be noted that the degree or

¹¹⁴ For the sake of simplicity I just discuss final and instrumental kinds of value, though other kinds of value could be discussed.

kind of value may differ from single instance to a collection of instances. For example, an instance of true belief may be instrumentally valuable, while having lots of true beliefs and few false beliefs may be finally valuable.

The value of an epistemic good may have various sources. An epistemic good may have a particular value in virtue of a property also shared by non-epistemic goods. Alternatively an epistemic good may have the particular value it has because of a property that is distinctly epistemic. Finally an epistemic good may have the particular value it has because of a property of the good in question, rather than say because of a property it shares with other epistemic goods. For example, knowledge may be valuable because there is something about knowledge per se that is valuable, or because knowledge has an epistemic property that makes it valuable though that property is not exclusive to knowledge, or knowledge may have the value it has because of a non-epistemic property – say because it's an achievement and achievements have a particular value.

The epistemic goods of the kind under consideration are different from everyday material goods, both in terms of the implications of their possession and their not being limited by scarcity with regard to potential possession for beings like us. Firstly, R possessing n amount of knowledge does not lessen the amount of knowledge that others may potentially possess – this is a not a zero-sum game; secondly, the further attainment of a single instance of an epistemic good is always theoretically possible. An example of this latter point is that, say in the case of our candidate epistemic good being true belief, even if a true belief is acquired it may always be possible in theory to gain a further true belief. Similarly, if our candidate epistemic good is the possession of a collection of instances of true beliefs and few false beliefs, then, even if such a collection is possessed, that collection can always be added to.¹¹⁵

¹¹⁵ If the epistemic good is taken to be some limited collection, then the latter point won't apply. This would be the case if, for instance, the epistemic good to be promoted were true beliefs that are of interest to us and such a collection is finite.

That the candidate epistemic goods are goods which are not limited like material goods in terms of the implications of their possession in one way makes a decision to assist their attainment much easier – in principle there need be no worries about losses incurred by other parties.¹¹⁶ On the other hand, however, the second point about epistemic goods means that agreement to intervene, in the absence of additional principles, leads to an imperfect or a limitless open commitment to assisting R's attainment of x. An initial step in establishing a principled basis for determining limits on state intervention might be achieved by calculating the cost in resources of state assistance in the attainment of x and/or by appealing to a sufficiency account of assisting in the attainment of whatever good is to be promoted.

An epistemic good's kind of value and the degree to which it is valuable may influence an answer to Q. For example, if understanding is recognised as having final value then that might be a reason to assist R in attaining understanding. Similarly, if true beliefs are recognised as particularly instrumentally valuable for one or more finally valuable goods then that might be a reason to assist R in attaining true beliefs. S might have reason not to assist R in attaining x if x has no value. Similarly it might be thought that a good merely having instrumental value is no reason **by itself** to assist someone in attaining it. More strongly still, one might think that without further specification one has no reason to assist R in attaining x even if x has final value and to a degree greater than other finally valuable goods; it may be the case that given further specification, one might see no reason to assist R or less reason to assist R than to not assist R.

If the reason for assistance is based on the value of the good, then, depending on the source of the value of the epistemic good of x, by implication there may be reason to assist in the attainment of further goods. If x is an epistemic good which has the value it has not in virtue of the good *per se* but in virtue of a property shared by other goods, be they epistemic or otherwise, then there will also be reason to assist R in the attainment of those other goods. It may be thought that if assistance is given for a good of one kind of value, that is instrumental or final, then that implies that

¹¹⁶ Of course in practice a state may incur a net loss of resources in assisting R's attainment of x, resources which may have been expended in a different way.

assistance should be offered for all goods of that kind. Such a thought would be a mistake. As already discussed, even if two goods are of the same kind of value, that does not entail that they are of the same value. For example, two goods may have instrumental value but one good may have more instrumental value than the other for a given end. Whether final value is also like this, that two goods may be finally valuable but one may be more finally valuable than the other, might be contested. Intuitively, however, it seems perfectly acceptable to say that of two goods that are assumed to be finally valuable, one good may be more finally valuable. For example, both wisdom and knowledge may be finally valuable but it seems intuitive that wisdom may be more finally valuable than knowledge.

The manner of interventions in the epistemic environment can be expected to differ for practical reasons depending on what the epistemic good to be promoted is. For example, it seems plausible to assume that if one wished to promote understanding then one would not intervene in the same way as one would if one wished to promote true belief. Similarly, the level of resources required for assisting the attainment of x may differ depending on the identity of x. Assisting R to attain an understanding of some matter can be expected to require more time and energy than assisting R to attain a true belief about some matter.

§6.2.0 The Status of Final Value Claims

An approach which seeks to justify intervention in the epistemic environment on the basis of findings in epistemic value theory, assuming that such a justification may be on the basis of a good having final value, should flag a key premise of such an approach.¹¹⁷ The key premise that certain goods can be said to be finally valuable may be contested. There is an already existing debate about making value judgements on the basis of conceptions of the good life. What is said in this debate is relevant to some of the goods claimed to have final value; that is the goods we take to have final value in virtue of being constitutive of the good life. The assignment of

¹¹⁷ What I say here is of relevance if the reader thinks that the candidate epistemic goods are no more than instrumentally valuable but are instrumentally valuable for some non-epistemic good that is constitutive of the good life.

final value rests on a normative claim, and this normative claim may equally be challenged by value pluralism and subjectivism, as will shortly be discussed. Any remainder of finally valuable goods that do not fall into the category of goods that have final value in virtue of being constitutive of the good life, for example the first book off the first printing press, may face analogous challenges to their normative status.

Perfectionist theorists generally hold “an objective account of the human good”, or the good life. (Wall, 2008). Just as in value based epistemic interventionism, perfectionists seek to develop a politics or ethics informed by the account of the human good to which they subscribe. In political discourse, therefore, perfectionism holds that institutional arrangements and public policies can either “promote or impede perfectionist values in various ways and to varying degrees.” (Wall, 2008). Aristotelian political theory, Marxism and a strand of liberalism are each considered to be perfectionist.¹¹⁸ (Wall, 2008).

Isaiah Berlin argues for value pluralism, which contrasts with perfectionism if we hold perfectionism to entail a monism about what counts as a good life. (Wall, 2008). Furthermore, according to Berlin, we have no way to interpersonally justify a particular ranking of these values or goods and that the pursuit of one good is necessarily at the cost of other goods. We’re not in a position to justify to another person that a hedonistic good is less valuable than economic equality or that some epistemic good is even in some cases of more value than autonomy. (Gaus and Courtland, 2011). Berlin (2002: 217) writes plainly that the goods which humans strive for are sometimes incommensurable. What we get on Berlin’s view is a possible rejection of undertaking the project ultimately envisaged in this chapter.¹¹⁹

¹¹⁸ Perfectionist liberalism claims that “developing individuality and cultivating capacities” is valuable. (Gaus and Courtland, 2011). “[I]t is only the cultivation of individuality which produces, or can produce, well-developed human beings... what more can be said of any condition of human affairs, than that it brings human beings themselves nearer to the best thing they can be? or what worse can be said of any obstruction to good, than that it prevents this?” (Mill, 1963, vol. 18: 267 quoted in Gaus and Courtland, 2011).

¹¹⁹ For a discussion on whether Berlin is committed to the objective or subjective status of values and the charge that Berlin’s pluralism is indistinguishable from relativism, see Cherniss and Hardy (2010).

Value pluralism should be distinguished from value subjectivism. This is the view, advanced by both Locke and Hobbes, that our ideas about what is valuable stem from our tastes. (Gaus and Courtland, 2011). Contrary to perfectionist claims, the subjectivist maintains that there is no objective human good. If what is valuable comes down to no more than individual tastes, then it might be thought that the state has no business in promoting such goods. I think this dismissal is a little quick and later I'll consider an argument compatible with value subjectivism for value based epistemic intervention in the epistemic environment. Generally though, I won't try to settle which of the three views is the right way of thinking about the status of final value claims.

§6.3.0 Possible Answers

Either the state can intervene or not intervene to assist a citizen in attaining epistemic good x. What I investigate is the possible theoretical underpinnings of either answer provided. Intervention could be advocated on paternalistic and non-paternalistic grounds. The paternalist options which I will examine are paternalism, soft paternalism and "libertarian paternalism"; the non-paternalist interventionist options are legal moralism, a position derived from legal moralism that I will call "prudential interventionism", liberal regulatorism; the non-interventionist option I will consider is political liberalism. I devote more space to discussing what I suspect for most epistemologists are the less familiar theoretical options of paternalism and legal moralism.

§6.3.1 Paternalism

The paternalist's answer is to say that intervention may be justified on the basis of the state acting to improve the welfare or promote the good of the citizen.¹²⁰ As the name suggests, paternalism is modelled after the father – child relationship. I provide

¹²⁰ Unless I make it clear that I have the general category of paternalism in mind, when I say paternalism I mean what is often referred to as hard paternalism.

two analyses of paternalism; the first I describe as paternalism, Gerald Dworkin's account of paternalism, and the second as paternalism*, my own account.

Dworkin (2010) offers the following analysis:

S acts paternalistically towards R by doing (omitting) Z:

- (C1) Z (or its omission) interferes with the liberty or autonomy of R.
- (C2) S does so without the consent of R.
- (C3) S does so just because Z will improve the welfare of R (where this includes preventing his welfare from diminishing), or in some way promote the interests, values, or good of R.¹²¹

I take it that the basic import of the first two conditions can be grasped without much difficulty; in any case I discuss them in subsequent paragraphs. The third condition, however, warrants immediate attention. The obvious reading of C3, to me at least, is that S's action will improve the welfare of R, amongst other things.¹²² But surely an action being paternalistic does not entail that it **will** improve R's welfare, never mind S's action being explained by the fact that that action will improve R's welfare. Rather than success, what seems to be required is that S's act be appropriately motivated. Perhaps this is indeed what Dworkin has in mind.¹²³ If that is so, then my point may be taken to be merely clarificatory in nature.

I've written that what is required for an act to be paternalistic is that it be appropriately motivated, but what would make an act appropriately motivated? One option is to say that S acts paternalistically when S does Z just because S is **motivated by a belief** that doing Z will improve R's welfare. Consider the following case:

¹²¹ Solely for presentational purposes, I've altered the symbols marking the identity of the agents, the labelling of each of the conditions and the spacing.

¹²² I'm saying "improve the welfare" for shorthand. My intention is not to challenge Dworkin's fuller description.

¹²³ Nevertheless I haven't found clear textual evidence that he does have this in mind. Dworkin writes that he is providing an analysis of paternalist action but he does not write that a paternalist act occurs if and only if the conditions he lists are met. If the conditions he provides are intended only to be of paradigm cases of paternalist acts, then let my comments on his analysis be seen as indicating what I believe a more basic analysis would look like.

A father believes that if his adult son takes vitamin supplement tablets, then his son's health and ultimately welfare will be improved. Knowing that his son hates taking tablets and so is likely to refuse to take the tablets if asked to do so, the father hides the tiny tablets in his son's evening meals and by so doing interferes with the autonomy of his son. Even if it turns out that taking the tablets does not improve the welfare of the son, intuitively the father has acted in a paternalistic way.

However, requiring that S be motivated by a belief, as proposed, is overly demanding. It seems conceivable that someone could act paternalistically and be motivated on the basis of something weaker than belief. Suppose S does Z just because S is **motivated by a positive epistemic standing** that doing Z may improve R's welfare.¹²⁴ Acting on such a basis intuitively doesn't seem to discount the possibility of the act being paternalistic. Epistemic standings that fall short of belief that may also be able to play the same role include acceptance and suspicion. To return to our example, if the father suspects that there is even a small chance that his son taking the vitamin tablet will improve his welfare, and acts on this basis, then the act is paternalistic.

Now let us turn our attention to the first two conditions. A broader rendering of C2 could capture the sense in which paternalism is a threat to liberty or autonomy, arguably making C1 unnecessary. My provisional suggestion is for a condition stating that S does Z irrespective of the consent of R. It seems to me that there is intuitive support for the idea that S could be acting paternalistically even if R did offer consent, so long as S would have acted in the same way regardless of whether R consented or not.

¹²⁴ As well as S having some epistemic standing in favour of Z, it might be thought that S should also not have any epistemic standing of equal or greater strength against Z: so S does Z just because S is **motivated by a positive epistemic standing** that doing Z may improve R's welfare and does not also have an epistemic standing of the same or greater strength that doing Z may or will disimprove R's welfare. (If S just had say a strong belief that Z wouldn't improve R's welfare then that wouldn't be a reason against acting on the basis of a weaker epistemic standing to do Z.) Requiring that the standing for the paternalist act be greater than or equal to the standing against seems to presuppose that the paternalist actor must be exhibiting rationality in a particular way, but this seems unnecessary.

We want our analysis of a paternalist act to allow for the possibility of the happy, let's say Victorian, paternalist, whose husbandry isn't met by resistance from wife and children, (and perhaps even his labour force if he has one and is a social reformer of the type typical of his age), but rather is taken by them, being of their time, as legitimate. Even if a charge offers consent to his action on her life, it seems intuitive to say that such an act could remain paternalist. What makes someone a paternalist intuitively isn't necessarily that he act in the absence of consent, but rather that the presence of consent or otherwise isn't a factor in his doing what he takes to be for the good of the person on whose behalf he is acting. In other words, someone doesn't stop being a paternalist when their decisions are met with consent, a person stops being a paternalist when their decisions are conditional on there being consent.¹²⁵

It may be objected that the revised condition which I have suggested relies on an artificial conception of consent. Saying that R consents to Z, suggests that Z was optional for R: it makes no sense to talk of R consenting to Z if Z has already been decided, just as it makes no sense to talk about a child consenting to a particular holiday destination when it has already been decided by his parents, even if the parents would like it if the child was enthusiastic about their choice. This is not a problem for Dworkin, whose condition is simply that there needs to be an absence of consent.

We might think, however, that that one cannot consent to acts of a certain kind, ones in which in fact there isn't an option, is a quirk of the meaning of "consent" and is not something we were trying to capture in our analysis of paternalism. Really the intuition we were trying to capture when we discussed consent and imagined the happy paternalist, was that an act could coincide with the wishes of someone and

¹²⁵ I shift in this paragraph from discussing the paternalist act to the paternalist for stylistic reasons. For the sake of clarity, it should be born in mind that the intended object of my analysis is the paternalist act. The need for clarity here arises from the fact that someone who defends the theory of paternalism, even if they have never acted in a paternalistic way, may be described as a paternalist. This implies that a paternalist need not just be one who acts in a paternalistic way.

remain paternalistic. These considerations yield the following condition: (C1*) S does so irrespective of what S believes the wishes of R may be.¹²⁶

The removal of Dworkin's C1 means that there is no longer any requirement that a paternalist action interferes with liberty or autonomy; what is retained on the alternative analysis however, implied by C2*, is that paternalist action threatens both liberty and autonomy. To say that an act may be one which does not interfere with liberty or autonomy but still may be a paternalist act does not conflict with the idea that the cases of paternalistic action that have traditionally been of interest are the ones that do in fact interfere with liberty or autonomy.

Nor does it seem inconceivable that an action may be paternalistic without interfering in the liberty or autonomy of the object of the paternalist action. Suppose that a mother is worried about the safety of her son in his new apartment. It's his first time living away from home and his mother knows that there is no smoke alarm in his apartment. She thinks that if she were to suggest that he get one, then he would agree, but knowing him as she does, she doesn't believe that he would actually get one. She knows that he is very proud of his new found independence, and she thinks that if she were to ask him whether he would like her to buy him one that he would say no. She decides to buy him one anyway and, by offering it to him already bought, tries to make his acceptance of it a *fait accompli*. Does the mother's action interfere with the son's liberty or autonomy? It certainly doesn't interfere with his liberty. His mother has to resort to such tactics because her powers over him are limited. He can simply refuse her gift. What about his autonomy? While his mother's actions certainly don't respect his right to decide things for himself, it's not obvious that they reduce or even hinder his autonomy.

My alternative analysis, which I will henceforth refer to as paternalism* runs as follows:

S acts paternalistically towards R by doing (omitting) Z (iff):

¹²⁶ I couch this condition in terms of what S believes because we can conceive of a world in which S is actually acting irrespective of the wishes of R but has the false belief that he is acting in accordance with the wishes of R.

- (C1*) S does so irrespective of what S believes the wishes of R may be.
- (C2*) S does so just because S has a positive epistemic standing that Z may or will improve the welfare of R (where this includes preventing his welfare from diminishing), or in some way promote the interests, values, or good of R.

The motivation to do what improves R's welfare irrespective of what S believes R's wishes may be is key to understanding how paternalism differs from other interventionist approaches that may likewise be contrary to the will of R but that may have the effect of making R better off. It's also important to note that paternalism may in some cases require that S not intervene, again irrespective of the wishes of R, so long as the motivation is that to not intervene may be or is better for R. If an approach, like paternalism, supports either intervention or non-intervention depending on further description of a case, then I will list it as an interventionist approach.

Having made clear what paternalism and paternalism* involve, let us return to the question. Both approaches would provide prima facie grounds for the intervention of the state on the grounds that such an intervention would make the citizen better off. An intervention which coincides with what S believes to be the wishes of the citizen but which is insensitive to those wishes can be classed as paternalism*, if C2* is also met. Furthermore, according to paternalism* but not paternalism, an intervention may be classed as paternalist even if it doesn't interfere with R's liberty or autonomy. For example, S adds fact checking resources to the epistemic environment rather than directly trying to intervene in any of the existing practices or institutions of its citizenry. So S may intervene in the epistemic environment on the basis of the motivation mentioned and in the insensitive way mentioned, but the intervention need not interfere with either R's liberty or autonomy.

§6.3.2 Soft Paternalism

According to soft paternalism, intervention is permissible, if it is to ensure that R is acting knowledgeably and voluntarily. (Dworkin, 2010). I focus my attention on the

knowledge component as it seems more likely in the kind of cases with which we are concerned that one might act without knowledge than that one might act involuntarily. Let me make the stipulation that the soft paternalist action is motivated in the way C2* describes.

Mill's example of soft paternalism is of forcibly stopping a man from walking off the end of a broken bridge in order to ensure that he knows what he is doing.¹²⁷ (Dworkin, 2010). The violation of the autonomy of the walker is justified if there is no other means of ensuring that he knows what he is doing; say you shout a warning, but it becomes clear that he doesn't understand English. Say if you find out that he knows what he is doing, that going off the end of the bridge is what he wants to do, then, according to the soft paternalist, it is impermissible on soft paternalist grounds to stop him. The soft paternalist limits herself to violating the other's autonomy in order to ensure that the other knows what he is doing. If the man walking off the end of the bridge knows what he is doing and wants to commit suicide, then the soft paternalist will not stop him from doing so.¹²⁸

Bringing soft paternalism to bear on our question, the state may intervene in order to ensure that the citizen knows the value of the epistemic good at stake motivated by a positive epistemic standing that the good at stake may or will be beneficial to the citizen. If the citizen, however, despite being informed about the value of the epistemic good, does not want to attain that good, then the state may not try to impose the good on the citizen.

But could R really know the value of the epistemic good in question and not seek to attain it? Let us assume for the sake of simplicity that the good is of final value. It might be thought that it follows from knowing the special value of the epistemic good that the agent will have a prima facie desire to attain that good. If R does not

¹²⁷ If it seems too implausible to think that he doesn't know what he is doing, then we can stipulate that it appears that he may be blind.

¹²⁸ This example makes understanding Mill's position somewhat complicated. Mill also writes that selling oneself into slavery should not be permissible. (Mill, 1863: 198-199). In both cases the actions may be knowingly and voluntarily carried out. In both cases the result will be future loss of autonomy for the actor. In one case, however, Mill seems to think the actor's decision should be overturned, while in the other case the decision should not be overturned.

have a prima facie desire to attain the good, then R has not learnt its special value, which would warrant further intervention.¹²⁹ Now soft paternalism, if we accept the above and at least with respect to certain value questions, does not look so soft. Perhaps this just indicates that the soft paternalist is committed to denying that knowledge that a good is valuable implies that an agent will have a prima facie desire for that good.

One alternative rendering of soft paternalism that would make it a good deal softer is not simply to prohibit intervention on the basis of knowing that the agent is acting knowledgeably, but to add the condition that if the agent makes it clear to the paternalist agent that he does not want to act knowledgeably, then the soft paternalist is prohibited from intervening. If the paternalist forcibly stops the walker from going off the end of the bridge and starts to try to ensure that the walker knows what he was doing, and if the walker then interrupts and says that he does not want to know what he was about to do, then according to the alternative rendering the soft paternalist would not be permitted to intervene further on the point. On this view the state could try to let the citizen know the value of x and, if the citizen wanted, assist him in its attainment. But if the citizen said he wasn't interested in knowing the value of x then the state would be required to desist in its attempts to ensure that R acted knowledgeably.

§6.3.3 Libertarian Paternalism

Libertarian Paternalism is the view advanced by Cass Sunstein and Richard Thaler (2003) that paternalist intervention is permitted to help an agent make the choice that she would want to make herself, but fails to do so because of weaknesses of her will or mind.¹³⁰ Weaknesses of will and mind that the libertarian paternalist intervention is motivated to protect the agent from include the bad effects of procrastination,

¹²⁹ It **may** be the case that R does have a prima facie desire to attain x but that the desire is defeated by other considerations.

¹³⁰ The justification for paternalist action on the basis of weaknesses of will or mind is similar to the soft paternalist's already discussed justification for action on the basis of ensuring an action is knowledgeable or voluntary. In fact "libertarian paternalism" is described as "soft paternalism" in literature on libertarian paternalism. For an example, see Douglas Glen Whitman and Mario J. Rizzo (412: 2007).

framing effects, the status quo bias, amongst others. (Sunstein and Thaler, 2003). It's paternalistic because, where there is libertarian paternalist intervention, choices are framed on the basis of findings in the psychology of decision making, in such ways as to influence the choice of the agent, who is the object of the paternalist intervention, such that the agent will be made better off. The intervention is libertarian because, as the authors insist, it does not curtail the freedom to choose, as paternalist intervention is normally taken to do. Even if the choice facing the agent has been framed so as to encourage a good choice being made, the agent is still free to choose how she pleases. A justification for this type of intervention is that in many cases there is no "neutral setting" that will allow agents to make decisions in such a way that non-salient features are neutralised from having any influence on their decision.

To use a favoured example, the libertarian paternalist does not advocate the banning of an opt out option when it comes to making pension contributions, rather given the libertarian paternalist's commitment to freedom of choice and what we know about human psychology and the benefits of making pension contributions, they argue that there should be an opt out option but that it should not be the default option. (Sunstein and Thaler, 2003: 1184.) Pension choices generally should be framed in such a way as to attempt to influence or nudge the object of the paternalist intervention towards what will make them better off.

The authors point out that there is scope for disagreement between libertarian paternalists on the respective weight to afford the libertarian and paternalist components. (Sunstein and Thaler: 2003: 1185-1186). For example, a libertarian paternalist who attaches greater importance to the paternalist component of the approach, while committed to allowing freedom of choice, may try to impose greater costs on an agent whose choice makes him worse off than a libertarian paternalist who attaches greater importance to the libertarian component. The particular libertarian paternalism that Sunstein and Thaler are advancing is one that "... attempts to ensure, as a general rule, that people can easily avoid the paternalist's suggested option." (Sunstein and Thaler, 2003: 1186).

Libertarian paternalism can be seen as lying between soft and hard paternalism. Like the soft paternalist, the libertarian paternalist advocates intervention in actions that are potentially harmful to the good of an agent, like walking off the end of bridges. While the soft paternalist seeks to ensure that the agent knows what he is doing, the libertarian paternalist uses the resources of the psychology of decision theory to influence him to not walk off the bridge. What arguably might be called manipulation, rather than influence, which the libertarian paternalist employs, is very close to the deception, and thus the riding roughshod over autonomy, that the hard paternalist might advocate employing to attempt to make the agent better off; although of course the hard paternalist can employ other tools when attempting to ensure a sought outcome.

The libertarian paternalist would, with regard to Q, *ceteris paribus*, advocate the state framing a decision for R to attain or not to attain x in such a way so as to make it more likely that R will, or let's simply say nudge R to if the first rendition seems uncharitable, decide what the state believes to be in his best interest. This might include attempting to attain x and being happy to accept state help if he can't attain x by himself. The scope of libertarian paternalism is narrower than the various forms Q might take in a real life example. Libertarian paternalism is concerned with decision making, whereas often the attainment or not of epistemic goods does not involve any kind of decision. I can understand something or not, a decision to understand or, more realistically, try to understand need not effect my attainment of the epistemic good.

§6.3.4 Legal Moralism

Legal moralism is the view that the immorality of a kind of conduct is relevant in considerations of the criminalisation of that conduct. Anthony Duff (2008), in his discussion of legal moralism, distinguishes positive legal moralism from negative legal moralism. The former is the view that a kind of conduct being immoral is a good reason for its criminalisation. "Good reason" here means some reason but not

necessarily an all things considered reason; for example, the practical difficulties of policing a law might be a sufficiently strong reason not to make the particular law in the first place. Negative legal moralism holds that the immorality of a kind of conduct is only a necessary condition for the criminalisation of that conduct but that only conduct of an immoral kind should be a candidate for criminalisation. One point of contrast between positive legal moralism and negative legal moralism, a point that arguably suggests the greater plausibility of positive legal moralism, is that positive legal moralism does not require something to be immoral for it to be criminalised; so driving on a particular side of the road may legitimately be criminalised without requiring that driving on that particular side of the road is immoral. In this section I concentrate my attention on developing a positive legal moralist account of epistemic intervention.

Legal moralism should not be confused with paternalism. The legal moralist does not justify the criminalisation of conduct on the basis of a claim that by so doing the welfare of the person whose liberty is threatened or restricted is promoted. The legal moralist simply holds that there is reason to criminalise conduct which is immoral conduct in virtue of it being immoral.

Legal moralism is regarded as a challenge to liberalism. One might find this hard to see if one believes that there is nothing morally wrong beyond that which the harm principle already implies should be regulated. According to the harm principle the restriction of another's liberty is only justifiable on the basis that it prevents another from harm.¹³¹ Of course liberals need not, though undoubtedly some do, think that instances of moral wrongness are exclusive to actions in which one person harms another person; rather, what is at issue here is the basis upon which the liberty of others may be restricted. The legal moralist is in a position to claim that both actions which the liberal thinks should be restricted as well as some actions which they don't, are actions which we have reason to legislate against. Actions that may be regarded

¹³¹ Mill articulates the harm principle as follows: "[T]he sole end for which mankind are warranted, individually or collectively, in interfering with the liberty of action of any of their number is self-protection. That the only purpose for which power can rightfully be exercised over any member of a civilised community against his will is to prevent harm to others. His own good, whether physical or moral, is not a sufficient warrant." (Mill, 1863: 23).

as immoral without at the same time being regarded as involving the non-consensual harming of one person by another include: prostitution, polygamy, incest between consenting adults, the viewing of snuff movies, desecration of the dead. It is to such cases that the legal moralist may point as examples of the kind of actions which she has the theoretical resources to justify the criminalisation of, which the liberal, arguably, has not.

Although having the theoretical resources to say why certain conduct should be criminalised which intuitively we believe should be illegal might seem a strength of legal moralism, the consistent application of these resources burdens the theory with having to defend counterintuitive claims as to what else we have reason to criminalise. Duff (2008) describes a case with this concern in mind. I betray a friend by frivolously revealing a secret with which she has entrusted me. What I've done is deeply immoral but intuitively the immorality of this kind of conduct doesn't seem a reason for its criminalisation. Remember, though, the positive legal moralist can counter that she is not committed to the position that conduct of an immoral kind must be criminalised, it's just that conduct being immoral is a reason to criminalise such conduct. However, for Duff it is counterintuitive to think that there is even a *pro tanto* reason for criminalising such conduct, it's just not an area with which the law should be concerned.¹³² It's worth noting that the liberal has a similar job to do in explaining why some intuitively harmful conduct warrants criminalisation and other such conduct does not.

A defender of legal moralism has one of several dialectical options here. She might simply deny the intuition that there is not even a *pro tanto* reason for criminalisation in the cases discussed. Alternatively she might accept the criticism and adjust her position accordingly. She might argue, as Duff (2008) suggests, that it is the immorality of a certain class of conduct, rather than the immorality of immoral conduct generally, which provides a reason for criminalisation. Another option is to

¹³² This is just an intuition though and it is possibly one that is closely related to the kind of society we live in. It seems likely that there is a widely shared intuition in our kind of society that someone shouldn't be sent to prison for committing adultery, while in another kind of society it seems plausible that people would have an intuition that adultery warrants a prison sentence. Thanks to Liz Ellis for this point.

argue that a community holding that conduct should be criminalised because the conduct is taken to be immoral yields a pro tanto reason for its criminalisation.¹³³

This latter way of responding to the betrayal case provides no a priori guidance as to which conduct might actually be criminalised. In fact, conduct that one community holds to be immoral and therefore worthy of criminalisation might be regarded as permissible or as immoral but unworthy of criminalisation in another. On this view, something that we find morally abhorrent, such as coerced or child female genital mutilation, and worthy of being criminalised, might not be so regarded in another community, and so the other community would lack this pro tanto reason to make the practice illegal.

There are at least two ways in which legal moralism might be drawn on to justify epistemic intervention. The straightforward way would be just to argue for one of the strands of legal moralism outlined and show its relevance for epistemic intervention. That the state could be shown to be obliged to assist R directly is prima facie less plausible than that the state might be shown to be obliged to intervene indirectly in the epistemic environment. The second less straightforward way would be to take the structure of one of the legal moralist strands, but rather than taking moral value as the key component, take what has been called, perhaps somewhat misleadingly, “prudential value” as the key component.¹³⁴ By “prudential value” I mean the kind of value the sum of which is used for judging the good life; to put it another way, that which varies according to our assessment of how well someone’s life has gone. I call intervention on this basis “prudential interventionism”.

Citizens possessing certain epistemic goods might be a demand of collective morality, in which case, according to legal moralism, the state could be required to intervene to promote the attainment of the epistemic goods in question. If moral value is attached

¹³³ Although this position may seem significantly different from legal moralism as it has been discussed so far, it bears a family resemblance to the position articulated by Lord Devlin, which is recognised in the literature as a legal moralist one. Devlin held that the law should reflect common morality because if it didn’t, then the bonds that hold society together would be threatened. (Devlin, 1970: 10).

¹³⁴ See James Griffin (2007: 314) for an example.

to the possession of epistemic goods or moral disvalue is attached to lacking epistemic goods, then the basis on which legal moralism could require intervention in the epistemic environment becomes clear. So long as epistemic goods are morally significant then they may be the concern of collective morality, and depending on exactly what that collective morality is, it may be legitimate for the state to intervene in the epistemic environment. Similarly if an epistemic good is judged to have prudential value by a community then that may provide some justification for legislative arrangements that reflect that judgement.

§6.3.5 Liberalism

Distinguishing political liberalism from comprehensive liberalism is important for the purposes of laying out different liberal positions on value based epistemic intervention. Political liberalism provides a framework within which people holding differing views of the good can pursue their own ends. (Rawls, 1993: xxviii).

Political liberalism is not intended to provide a competing theory of the good, rather it aims to provide a solution to their being competing theories of the good such that if the state adopts the framework, then the state can enjoy reason based endorsement from the adherents of competing theories of the good given the framework's neutrality on the good. On this view, it would be problematic for the state to promote a particular conception of the good which it would be doing if it were promoting epistemic goods on the basis of a judgement about their value.

A softer form of political liberalism might hold that while direct intervention by the state is not desirable, the state does have a role in regulating systems which produce certain common goods which may upon reflection be appreciated as such. Such common goods include functioning markets, clean air and a "clean" epistemic environment. Regulation implies the recognition of the value of certain goods and hence their target for regulation. Regulation can take various forms depending on the nature of the good and the system which produces it. Liberal regulatorism does not advocate imposing particular ends on a citizen population as paternalism and legal

moralism may. The citizen can decide for themselves whether they will seek to attain the good that is the object of regulation.

Liberal regulatorism, in relation to our question, could be interventionist by requiring the shaping of the epistemic environment in a way so as to make it friendly for the attainment of the epistemic goods in question. The liberal regulatorist, however, would need to give an explanation as to why this particular set of goods is worthy of regulation.

§6.4.0 Conclusion

In this chapter I have set out what I take to be important distinctions in epistemic value theory, to be drawn upon for the purposes of applied social epistemology. I have also set out various potential partners to which epistemic value theory might be wedded for the purposes of applied social epistemology. My purpose has been to provide a lay of the land rather than to recommend one particular approach. Nonetheless, I hope that my examination of the lay of the land indicates areas that require more detailed exploration and difficulties to be overcome.

Thesis Conclusion

My thesis began with the question, “why knowledge?” By way of elaboration I followed up by asking why a whole branch of philosophy has developed which has as perhaps its most central tasks the study of knowledge. I noted that there are nearby alternatives to knowledge that have not enjoyed such consideration. I explained that I hoped to go some way towards answering “why knowledge?” by investigating the value of knowledge by way of an examination of its nature.

In my first chapter I set out Pritchard’s articulation of the three value problems of knowledge; (i) why is knowledge more valuable than true belief?; (ii) why is knowledge more valuable than that which falls short of knowledge?; (iii) why is knowledge distinctively valuable vis-à-vis that which falls short of knowledge? I then distinguished different kinds of value, intrinsic from extrinsic, final from instrumental, as well as different ways in which a good may be valuable, pro tanto from prima facie. Having laid out the above, I turned to where discussion of the value of knowledge traditionally begins, with Plato’s *Meno*. Having highlighted shortcomings with the solution offered in the dialogue, following Pritchard I considered Greco’s account as one that is well placed to account for the value of knowledge and solve the three value problems. I began my examination of Greco’s account by showing how his account of the value of knowledge, according to which knowledge is valuable as an achievement, falls out of his account of the nature of knowledge, according to which knowledge is true belief because of ability. I explained that by claiming knowledge to be an achievement and achievements to be constitutive of the good life and so finally valuable, Greco looks well placed to account for the distinctive value of knowledge and the value problems generally. I argued, however, that Greco’s account of the value of knowledge is both in a certain respect too narrow and in another too broad. More specifically, I argued that his account of the value of knowledge is too narrow in that justified belief is plausibly an

achievement. I made the case that justified belief is intuitively a success from ability and thereby satisfies the criterion for achievement. I argued that his account is too broad in that it suggests that knowledge is pro tanto finally valuable. In support of this argument I presented the Pierre case, and claimed there to be the response intuition to the case that the protagonist does not gain anything that is pro tanto finally valuable, even though he gains knowledge and has an achievement.

In Chapter II I turned my attention to Greco's account of the nature of knowledge. With Pritchard I argued that his account is too weak in that it rules in counting Barney-type cases as cases of knowledge and that his account is too strong in that it rules out plausible cases of testimonial knowledge as knowledge. In arguing that Greco's account of the nature of knowledge is both too weak and too strong I followed the moves made in the dialectic and with respect to each charge added a further move.

I argued that Greco's response to Barney-type cases is such that either he has to say the same about standard Gettier cases, thereby denying that they are cases of justified true belief and thus losing the neat explanation of those cases, a feature that makes his account attractive, or he must bite the bullet and count Barney-type cases as cases of knowledge. I did so by arguing that standard Gettier style cases, exemplified by the Roddy case, imply the presence of the kind of environment found in Barney-type cases and I supported this argument with the Tony case.

I argued that Greco's response to the challenge posed by the Jenny* case is such as to undermine his original motivation for moving to an agent reliabilist account of knowledge. I did so by laying out how his treatment of the Jenny* case, expanding on the stated facts of the Jenny case, could equally be applied to the treatment of the Careless Math Student case with the unwanted result that the latter case would come out as a case of knowledge. I claimed that again there is a dilemma; deny that cases such as the Jenny* case are cases of knowledge and thereby allow the preservation of the original verdict in the Careless Math Student case or grant that Jenny has

knowledge in the way described and that, given the sameness of relevant factors present in that description, the Careless Math Student also has knowledge.

Finally I looked to Pritchard's anti-luck virtue epistemology as a possible alternative to Greco's robust virtue epistemology. Considering the merits of the account in relation to the two challenges posed to Greco's account of the nature of knowledge, I argued that Pritchard's account can rule out Barney-type cases as being cases of knowledge via an anti-luck condition but only in a way that seems ad hoc. I argued that while Pritchard's weaker ability requirement is such as to better place his account to deal with accounting for cases of testimonial knowledge, nevertheless the requirement still looks too strong to account for typical cases of testimonial knowledge in young children.

In Chapter III I attempted to overcome the challenges set out in Chapter II within a virtue epistemic framework. I began by laying out what I take to be three central challenges to developing an account of testimonial knowledge. I argued that such an account should, (i) be one that makes sense of our experience of testimonial reception; (ii) be such as to allow us to account for different degrees of justified testimonial belief; (iii) allow us to account for intuitive cases of testimonial knowledge in young children. I argued that an account can do (i) and (ii) by drawing on dual systems theory and in doing so a credulist counterfactual sensitivity position and the Humean position are conceived as being much closer that they are usually conceived as being. I argued, however, that if we are to account for (iii) we need to broaden the virtue requirement to include believing from a virtuous trait. I argued that in the case of young children it is epistemically appropriate for them to trust their caregivers; that believing from trust in such cases is believing from an epistemically virtuous trait. I argued that this doesn't imply that trusting is always epistemically appropriate by making the case that what counts as a virtuous way of forming a belief may vary according to factors such as the kind of agent involved, their social role and epistemic duties.

In attempting to provide a way to account for Barney-type cases I returned to the Gettier problem. I laid out what I take to be the orthodox diagnosis of Gettier-type cases; that the protagonist does not know because she is lucky to have a true belief. I argued that rather than seeing Gettier-type cases as cases in which an agent is lucky to have a true belief, it makes more sense to see them as cases in which an agent is unlucky not to have knowledge given that the agent in such cases, after all, has a justified belief. This diagnosis of Gettier-type cases chimes with the idea that despite an agent having an epistemically virtuous belief, and so doing their part, the world may fail to cooperate such that such an agent does not attain knowledge. Ultimately the position I defended, inspired by the aforementioned considerations, was that knowledge requires epistemic grace, where epistemic grace is understood as analogous to religious, synergetic grace. More specifically I argued that *S knows that p iff S has an epistemically virtuous belief that p and that virtuous belief enjoys epistemic grace*. This account of knowledge can be spelt out as requiring, (i) a virtuous belief which can be understood as a belief formed in an appropriate way given the agent's situation, the kind of agent she is, her epistemic duties, etc.; (ii) that her virtuous belief is formed in such a way as to enjoy epistemic grace in that the virtuous way in which she formed her belief in her environment is actually a right way of forming a belief in that environment, where right way means a way in which to gain a true belief.

In Chapter IV I returned to the topic of the value of knowledge. I argued that the value of knowledge can be understood in the light of the grace account of knowledge. More specifically I argued that while knowledge and virtuous belief have the value of achievement, knowledge and justified true belief have the value of achievement plus whatever value true belief has, knowledge has a value that differs from that which falls short of knowledge in virtue of having the value of epistemic grace. I argued that this value comes from an appropriate contribution being made by an agent, the agent virtuously believing, and that appropriate contribution enjoying grace. In making my point I drew on an analogy with a case involving friendship. I argued that having the virtues of a potentially good friend isn't sufficient to secure friendship or the feeling of closeness that accompanies it. As with justified true belief

it is possible to exercise the relevant virtues and to happen to get the feeling of closeness that accompanies friendship. I argued that knowledge is more and differently valuable than that which falls short of it in that knowledge requires the accompanying good, truth, being secured not independent from the virtue of the agent but due to the virtue of the agent and grace. I argued that while the way in which knowledge is valuable is different from that which falls short of knowledge, I didn't argue that knowledge has a different kind of value than that which falls short of knowledge. In other words I didn't argue that knowledge is finally valuable while that which falls short of knowledge is not. Rather, I argued that although sharing the same kind of value as justified true belief, the way in which knowledge is valuable is different and more valuable than justified true belief. The account of the value of knowledge offered is one on which knowledge is special.

Next, drawing on the argument made with regard to the Pierre case, I argued that the claim about the value of knowledge can be made more precise by saying that it is in the nature of knowledge to be valuable. I argued for this more restricted claim on the basis that the final value of the different epistemic goods being put forward is in virtue of the contribution of those goods to the good life. The thought is that it would be mistaken to think that an instance of an epistemic good is finally valuable, even in a pro tanto way, in virtue of being constitutive of the good life, if that instance lacks virtue simpliciter.

Finally, after noting how my account looks to provide a promising basis for understanding the value of other prized epistemic goods, I argued that a consequence of the value of epistemic goods outlined is such that we have a prima facie reason to want to improve what I describe as our epistemic environment. Depending on whether our epistemic environment is better or worse, we will be in a better or worse position to hold or attain goods which are constitutive of the good life. I described a motivation to alter the epistemic environment on the basis of the value of epistemic goods as epistemic environmentalism.

In Chapter V I turned to an examination of relations in the epistemic environment that are crucial to the attainment of a certain kind of knowledge. These relations are characterised by trust. Given the importance of trust to these relations and its general relevance to testimonial knowledge, I provided a detailed analysis of trust. Drawing on various cases, some from the literature and some of my own devising I argued, roughly, that S trusts R to Φ if and only if S is disposed to believe that R has the competence to Φ and S is disposed to believe that R has a goodwill to Φ . I argued that my analysis of trust makes sense of the reactive attitudes that accompany trust. Assuming the trusting agent doesn't come to think of trusting as inappropriate she will feel let down if it turns out that the trusted agent is not competent to p and she will feel betrayed if the agent turns out not to have a goodwill to p. I next turned to the expert – non-expert relationship. I considered Hardwig's argument that in order to gain knowledge in an expert domain the non-expert has no choice but to trust the expert. While generally supporting the claim I identified research that the problem of outright falsification is not quite as bad as Hardwig suggests. I also identified research on the lack of accuracy of university media centres in communicating expert findings. Finally I described the relationship between the occupant of a position of epistemic authority and the non-occupant. I argued that this relationship is analogous to the expert – non-expert relationship in that the non-occupant is in a similar position of epistemic dependence vis-à-vis the occupant of a position of epistemic authority.

In Chapter VI I set out a framework for the kind of epistemic intervention called for by epistemic environmentalism. Here I was primarily concerned with possible state intervention, though I acknowledged the relevance of intervention by non-state actors too. I laid out the various factors specific to epistemic goods and relevant to possible justifications for intervention in the epistemic environment. I turned to these possible justifications for intervention and considered each in turn. The point here was to provide a framework for epistemic interventionism; I didn't ultimately discuss whether the justifications for intervention (or non-intervention) were good ones. I did, however, attempt to set out each of the bases for justification in as positive a light as possible. In this regard I did significant explanatory work to set out what I take

paternalism to involve and I did the same, albeit to a lesser degree, with legal moralism.

The way I have attempted to answer the question – why knowledge?, has been to offer an account of the unique value of knowledge that falls out of an account of the nature knowledge. Knowledge really is special on the account I put forward. The specialness of knowledge, on my position, does not entail that other epistemic goods aren't valuable in their own ways. I examine some of the normative implications of the value of epistemic goods generally and in so doing articulate epistemic environmentalism.

Bibliography

- Adler, Jonathan, "Epistemological Problems of Testimony", *The Stanford Encyclopedia of Philosophy (Winter 2010 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2010/entries/testimony-episprob/>>.
- Adler, Jonathan (2002). *Belief's own ethics*. Cambridge, MA: Bradford/MIT.
- Aristotle (2009). "The Nature of Virtue", extracts from *Nicomachean Ethics* (1999) translated by Terence Irwin, in *Ethical Theory: An Anthology* edited by Russ Shafer-Landau. Oxford: Blackwell Publishing.
- Baier, Annette (1986). "Trust and antitrust", *Ethics* 96 (2):231-260.
- Ballantyne, Nathan (2012). "Luck and interests", *Synthese* 185 (3):319-334.
- Berlin, Isaiah. "Two Concepts of Liberty." *In Liberty*, . (, n.d.). Retrieved 22 Mar. 2012, from <http://www.oxfordscholarship.com/view/10.1093/019924989X.001.0001/acprof-9780199249893-chapter-4>
- Bird, Alexander (2007). "Justified judging", *Philosophy and Phenomenological Research* 74 (1):81-110.
- Cherniss, Joshua and Hardy, Henry, "Isaiah Berlin", *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2010/entries/berlin/>>.
- Chignell, Andrew, "The Ethics of Belief", *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2010/entries/ethics-belief/>>.
- Chisholm, Roderick. (1977). *Theory of Knowledge*, (2nd ed.). Englewood Cliffs, New Jersey: Prentice-Hall.
- Chisholm, Roderick (1973). *The Problem of the Criterion*. Milwaukee, Wis.: Marquette University Press.

- Coady, C.A.J. (1992). *Testimony: A Philosophical Study*. Oxford: Oxford University Press.
- Cohen, Stewart (2002). "Basic Knowledge and the Problem of Easy Knowledge", *Philosophy and Phenomenological Research*, 65: 309-329.
- Craig, Edward (1990). *Knowledge and the State of Nature: an essay in conceptual synthesis*. Oxford: Clarendon Press.
- David, Marian (2005). "Truth as the Primary Epistemic Goal: A Working Hypothesis", in *Contemporary Debates in Epistemology*, eds. E. Sosa and M. Steup, 296-312, Oxford: Blackwell.
- David, Marian (2001). "Truth as the Epistemic Goal", in Steup (ed.), *Knowledge Truth and Duty*. New York: Oxford University Press.
- Devlin, Baron Patrick (1970). *The Enforcement of Morals*. Oxford: Oxford University Press.
- Doris, John and Stich, Stephen, "Moral Psychology: Empirical Approaches", *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2011/entries/moral-psych-emp/>>.
- Doris, J. (2002). *Lack of Character: Personality and Moral Behavior*. Cambridge: Cambridge University Press.
- Duff, Antony, "Theories of Criminal Law", *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/criminal-law/>>.
- Dworkin, Gerald, "Paternalism", *The Stanford Encyclopedia of Philosophy (Summer 2010 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/sum2010/entries/paternalism/>>.
- Evans, Jonathan St. B. T. (2008). "Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition" in *Annual Review of Psychology*. Vol. 59: 255-278.
- Evans, Jonathan St. B. T. (2003). "In Two Minds: dual-process accounts of reasoning", *Trends in Cognitive Sciences*. Vol. 7, No. 10.
- Fanelli D. (2009). "How Many Scientists Fabricate and Falsify Research? A Systematic Review and Meta-Analysis of Survey Data", *PLoS ONE* 4(5).

- Faulkner, Paul (2007). "A genealogy of trust", *Episteme* 4 (3):305-321.
- Faulkner, Paul (1998). "David Hume's Reductionist Epistemology of Testimony", *Pacific Philosophical Quarterly*, 79: 302-313.
- Feldman, Fred (2002). "The Good Life: A Defence of Attitudinal Hedonism," *Philosophy and Phenomenological Research*, 65(3) 605-27.
- Frankish, Keith (2009). "Systems and Levels: Dual-system theories and the personal-subpersonal distinction", in Evans, Jonathan and Frankish, Keith eds. *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press, pp. 89-107.
- Fricker, Elizabeth (1994). "Against Gullibility", in *Knowing from Words* edited by Matilal and Chakrabarti, 125-161. Boston: Kluwer.
- Fricker, Miranda (2009). "The Value of Knowledge and The Test of Time", *Royal Institute of Philosophy Supplement* 84 (64):121-138.
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Gaus, Gerald and Courtland, Shane D., "Liberalism", *The Stanford Encyclopedia of Philosophy (Spring 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/liberalism/>.
- Gelfert, Axel (2010). "Hume on Testimony Revisited", *Logical Analysis and History of Philosophy* 13: 60-75.
- Goldberg, Sandy (2011). "The Division of Epistemic Labour", *Episteme* 8.1: 112-125.
- Goldman, Alvin, "Reliabilism", *The Stanford Encyclopedia of Philosophy (Spring 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>.
- Goldman, Alvin I. (2001). "Experts: Which ones should you trust?" *Philosophy and Phenomenological Research* 63 (1):85-110.
- Goldman, Alvin I. (1999). *Knowledge in a Social World*. Oxford: Clarendon Press.
- Goldman, Alvin I. (1991). "Epistemic Paternalism: Communication Control in Law and Society", *Journal of Philosophy*, Vol. 88, No. 3, pp. 113-131.
- Goldman, Alvin I. (1976). "Discrimination and Perceptual Knowledge", *Journal of Philosophy*, 73: 771-791.

- Goldman, Alvin I. (1967). "A causal theory of knowing", *Journal of Philosophy*, 64 (12):357-372.
- Greco, John and Turri, John, "Virtue Epistemology", *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), Edward N. Zalta (ed.), forthcoming URL = <<http://plato.stanford.edu/archives/spr2011/entries/epistemology-virtue/>>.
- Greco, John (2011). "The Value Problem", in *The Routledge Companion to Epistemology*, edited by Sven Bernecker and Duncan Pritchard. New York: Routledge, pp. 219-231.
- Greco, John (2010). *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Greco, John (2009). "The Value Problem", in *Epistemic Value*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press, pp. 313-321.
- Greco, John (2009b). "Knowledge and success from ability", *Philosophical Studies* 142 (1):17-26.
- Greco, John (2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: University Press.
- Griffen, James (2007). "The Informed Desire Account", in *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau, 306-314. Oxford: Blackwell Publishing. Originally published in James Griffin, *Well-Being* (Oxford: Oxford University Press, 1986).
- Habermas, Jürgen (2003). *Truth and Justification*, translated by Barbara Fultner. Cambridge: MIT Press.
- Hardwig, John (1994). "Toward an Ethics of Expertise" in *Professional Ethics and Social Responsibility*, edited by Daniel E. Wueste. Lanham: Rowman and Littlefield Publishers, pp. 83-101.
- Hardwig, John (1985). "Epistemic Dependence", *The Journal of Philosophy*, 82: 335-349.
- Hume, David (1999). *An Enquiry concerning Human Understanding*, edited by Tom L. Beauchamp. Oxford and New York: Oxford University Press.
- Hume, David. (1978). *A Treatise of Human Nature*, L.A. Selby-Bigge, ed, second edition, revised ed. P.H. Nidditch, Oxford: Oxford University Press.

- Hursthouse, Rosalind, "Virtue Ethics", *The Stanford Encyclopedia of Philosophy* (Summer 2012 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2012/entries/ethics-virtue/>.
- Ichikawa, Jonathan (*draft*). "Justification is Potential Knowledge".
- Kagan, Shelly (1994). "Me and My Life", *Proceedings of the Aristotelian Society* 94: 309-324.
- Kallestrup, Jesper and Pritchard, Duncan (2012). "Robust virtue epistemology and epistemic anti-individualism", *Pacific Philosophical Quarterly* 93 (1):84-103.
- Kallestrup, Jesper and Pritchard, Duncan (2011). "Virtue Epistemology and Epistemic Twin Earth", *European Journal of Philosophy* 20 (2).
- Kenny, Anthony (2007). *Philosophy in the Modern World: A New History of Western Philosophy, vol. 4*. Oxford: Oxford University Press.
- Kvanvig, Jonathan (2005). "Truth is not the Primary Epistemic Goal", in *Contemporary Debates in Epistemology*, eds. E. Sosa and M. Steup, 285-96, Oxford: Blackwell.
- Kvanvig, Jonathan L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Kvanvig, Jonathan L. (1998). "Why Should Inquiring Minds Want to Know?", *The Monist* 81 (3):426-451.
- Lackey, Jennifer (2008). *Learning from Words - Testimony as a Source of Knowledge*. Oxford: Oxford University Press. *Oxford Scholarship Online*. Oxford University Press. 17 March 2011
<http://dx.doi.org/10.1093/acprof:oso/9780199219162.001.0001>
- Lackey, Jennifer (2007). "Why we don't deserve credit for everything we know", *Synthese*, 158: 345-361.
- Lehrer, Keith (1990). *Theory of Knowledge*, Boulder: Westview Press.
- Locke, John (1975). *An Essay Concerning Human Understanding*, edited by Peter H. Nidditch. Oxford: Oxford University Press.
- Lycan, William G. (2001). "Moore Against The New Skeptics", *Philosophical Studies*. Volume 103, Number 1, 35-53.
- Lynch, Michael (2004). *True to Life: Why Truth Matters*. Cambridge: MIT Press.

- Maier, John, "Abilities", *The Stanford Encyclopedia of Philosophy (Fall 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2011/entries/abilities/>.
- Martinson, Brian C., Melissa S. Anderson and Raymond de Vries (2005). "Scientists behaving badly", *Nature* 435, 737-738.
- McGlynn, Aidan (2012). "Justification as "Would-Be" Knowledge", *Episteme* 9 (4):361-376.
- McLeod, Carolyn, "Trust", *The Stanford Encyclopedia of Philosophy (Spring 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/trust/>.
- Mill, John Stuart (1963). *Collected Works of John Stuart Mill*, J. M. Robson (ed.), Toronto: University of Toronto Press.
- Mill, John Stuart (1863). *On Liberty*, second edition. Boston: Ticknor and Fields.
- Nagel, Thomas (1979). *Mortal Questions*. Cambridge University Press.
- Newman, John Henry (2007). *An Essay in Aid of a Grammar of Assent*. London: Elibron Classics.
- Newman, Lex, "Descartes' Epistemology", *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2010/entries/descartes-epistemology/>.
- Owens, David J. (2003). "Does Belief have an Aim?", *Philosophical Studies* 115 (3):283-305.
- Plato (2008). *Meno*, translated by Benjamin Jowett. Forgotten Books. http://books.google.ie/books?id=p3oQoPm-b_cC&printsec=frontcover&dq=Meno,&hl=en&ei=Y35aTeqxH8OShAePidCqDQ&sa=X&oi=book_result&ct=result&resnum=1&ved=0CCcQ6AEwAA#v=onepage&q&f=false
- Plato (1980). *Republic*, translated by Paul Shorey, in *Plato: the collected dialogues* edited by Edith Hamilton and Huntington Cairns. Princeton: Princeton University Press, pp. 575-848.
- Plato (1963). *Meno* in *Plato: The Collected Dialogues*, translated by W. K. C. Guthrie, 353-384, Princeton: Princeton University Press.

- Potter, Nancy (2002). *How Can I be Trusted? A Virtue Theory of Trustworthiness*, Lanham, Maryland: Rowman & Littlefield.
- Pritchard, Duncan (*forthcoming*). "Truth as the Fundamental Epistemic Good", *The Ethics of Belief: Individual and Social*, (eds.) J. Matheson & R. Vitz. Oxford: Oxford: University Press.
- Pritchard, Duncan (2012). "Anti-Luck Virtue Epistemology", *Journal of Philosophy* 109, 247-79.
- Pritchard, Duncan and Turri, John, "The Value of Knowledge", *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), forthcoming URL = [<http://plato.stanford.edu/archives/win2012/entries/knowledge-value/>](http://plato.stanford.edu/archives/win2012/entries/knowledge-value/).
- Pritchard, Duncan (2010). "Knowledge and Understanding" in Haddock, A., A. Millar and D. Pritchard, *The Nature and Value of Knowledge: Three Investigations*, 3-88, Oxford: Oxford University Press.
- Pritchard, Duncan (2010b). "Achievements, Luck and Value", *Think* 25, 1-12.
- Pritchard, Duncan, "The Value of Knowledge", *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = [<http://plato.stanford.edu/archives/fall2008/entries/knowledge-value/>](http://plato.stanford.edu/archives/fall2008/entries/knowledge-value/).
- Pritchard, Duncan (2007). "Recent Work on Epistemic Value", *American Philosophical Quarterly* 44, 85-110.
- Pritchard, Duncan (2006). *What is this thing called knowledge?* London and New York: Routledge.
- Pritchard, Duncan (2005). *Epistemic Luck*, Oxford: Oxford University Press.
- Pritchard, Duncan (2004). "The Epistemology of Testimony", *Philosophical Issues* 14 (1):326-348.
- Rawls, John (1993). *Political Liberalism*, New York: Columbia University Press.
- Reid, T. (1983). *Inquiry and Essays*, (eds.) Beanblossom, R. and Lehrer, K., Indianapolis: Hackett.
- Ridge, Michael (2011). "Getting Lost on the Road to Larissa", *Noûs* 47 (1):181-201.
- Rolston, Holmes (1988). *Environmental Ethics: duties to and values in the natural world*. Philadelphia: Temple University Press.

- Schroeder, Mark, "Value Theory", *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/value-theory/>.
- Scott, Dominic (2006). *Plato's Meno*. Cambridge: Cambridge University Press.
- Smith, Martin (2010). "What Else Justification Could Be", *Nous* 44 (1):10-31.
- Sosa, Ernest (1991) [1985]. "Knowledge and Intellectual Virtue", *The Monist* 68; reprinted in *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press, 225-244.
- Sunstein, Cass R. and Thaler, Richard H. (2003). "Libertarian Paternalism Is Not an Oxymoron", *University of Chicago Law Review* 70(4): pp. 1159.
- Tangney, J. P. (1987). "Fraud Will Out-Or Will It?", *New Scientist* CXV, pp. 62-3.
- Unger, Peter. (1975). *Ignorance*. Oxford: Oxford University Press.
- Unger, Peter. (1968). "An Analysis of Factual Knowledge", *Journal of Philosophy* 65:157-70.
- Wall, Steven, "Perfectionism in Moral and Political Philosophy", *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/perfectionism-moral/>.
- Waterfield, Robin (2010). *Why Socrates Died: Dispelling the Myths*. New York: Norton.
- Whitman, Glen and Rizzo, Mario J. (2007). "Paternalist Slopes", *NYU Journal of Law and Liberty* 2.3: 411-443.
- Williams, Bernard Arthur Owen (2002). *Truth & Truthfulness : an essay in genealogy*. Princeton and Oxford: Princeton University Press.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Woloshin, Steven, and Lisa M. Schwartz, Samuel L. Casella, Abigail T. Kennedy, Robin J. Larson (2009). "Press Releases by Academic Medical Centers: Not So Academic?", *Annals of Internal Medicine* May; 150 (9):613-618.
- Zagzebski, Linda. (2004). "Epistemic value and the primacy of what we care about", *Philosophical Papers* 33, 353-377.
- Zagzebski, Linda (1994). "The inescapability of Gettier problems", *Philosophical Quarterly* 44 (174):65-73.

Zimmerman, Michael J., "Intrinsic vs. Extrinsic Value", *The Stanford Encyclopedia of Philosophy (Winter 2010 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2010/entries/value-intrinsic-extrinsic/>.

Internet Sources

“Four Yorkshiremen Sketch”, 2012. Available at:

<http://www.davidpbrown.co.uk/jokes/monty-python-four-yorkshiremen.html>

[Accessed on October 8th, 2012].

Dictionary.com. 2010 "environment," in *Merriam-Webster's Medical Dictionary*.

Source location: Merriam-Webster, Inc. Available at:

<http://dictionary.reference.com/browse/environment> [Accessed: September

10, 2010].