

Developing an Understanding-Oriented Epistemology for the Environment

Alice Rose Wheatley

100176714

PhD in Philosophy

University of East Anglia

School of Politics, Philosophy and Area Studies

July 2025

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with the author and that use of any information derived therefrom must be in accordance with current UK Copyright Law. In addition, any quotation or extract must include full attribution.

Abstract

An understanding of the natural environment and present environmental challenges involves a vast amount of complex information. This produces significant epistemic difficulties that epistemology should have something to say about. This thesis makes three important points about an adequate epistemology for the environment. First, a focus on propositional knowledge is not sufficient to make progress on the epistemological questions that arise from our relationship to the environment, so we should shift our attention to the epistemic state of understanding. Second, we should take interest in what character traits make a person a virtuous environmental thinker. I argue for understanding-oriented accounts of some environmental-epistemic virtues: ‘big picture thinking’ virtues, like outward-lookingness, and the virtue of honesty. Third, I contend that epistemologists should refer to empirical research and engage in interdisciplinary approaches to explore how epistemic difficulties are experienced in relation to actual environmental interventions. I carry out a case study on the Nearshore Trawling Byelaw (conservation management in Sussex that bans harmful fishing practice) to illustrate the point. I review relevant documents and use semi-structured interviews to understand how epistemic successes and difficulties were experienced in practice.

Access Condition and Agreement

Each deposit in UEA Digital Repository is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the Data Collections is not permitted, except that material may be duplicated by you for your research use or for educational purposes in electronic or print form. You must obtain permission from the copyright holder, usually the author, for any other use. Exceptions only apply where a deposit may be explicitly provided under a stated licence, such as a Creative Commons licence or Open Government licence.

Electronic or print copies may not be offered, whether for sale or otherwise to anyone, unless explicitly stated under a Creative Commons or Open Government license. Unauthorised reproduction, editing or reformatting for resale purposes is explicitly prohibited (except where approved by the copyright holder themselves) and UEA reserves the right to take immediate 'take down' action on behalf of the copyright and/or rights holder if this Access condition of the UEA Digital Repository is breached. Any material in this database has been supplied on the understanding that it is copyright material and that no quotation from the material may be published without proper acknowledgement.

Table of Contents

Acknowledgements	viii
Introduction	
0.1 The Epistemic Problem.....	1
0.2 The Plan	3
Chapter 1: Epistemic Aims and Concerns	
1.0 Introduction.....	5
1.1 The Epistemic Point of View	6
1.2 Truth is not the Sole Fundamental Epistemic Aim or Concern	9
1.2.1 Is Grasping Epistemic?	11
1.2.2 Knowledge and True Belief?	12
1.2.3 Explaining Grasping as an Epistemic Aim or Concern in Terms of Truth?.....	14
1.3 Why is Grasping Important?	15
1.4 The Problem of Trivial Truths	16
1.4.1 What is the Trivial Truths Problem?	17
1.4.2 Truth and Reasons.....	20
1.5 Objections.....	22
1.5.1 Wrong Kinds of Reasons	23
1.5.2 Distinctiveness of the Epistemic.....	24
1.5.3 Epistemic and Other Sources of Normativity	25
1.6 Conclusion.....	25
Chapter 2: Objectivism and Externalism	
2.0 Introduction.....	27
2.1 Objective and Subjective Reasons.....	28
2.2 What are Facts that Objectively Count in Favour?.....	31
2.3 Objective Reasons and Environmental Ethics.....	35
2.3.1 Non-Anthropocentrism in Environmental Ethics	36
2.3.2 Why Does Environmental Ethics Need Objective Reasons?.....	38
2.4 Internalism and Externalism	43
2.5 Conclusion	48
Chapter 3: Grasping	
3.0 Introduction.....	49
3.1 Being in an Epistemic Position to Respond to Normative Reasons	49

3.1.1	Condition (i): Mental Representation	51
3.1.2	Condition (ii): Ability	53
3.1.3	Condition (iii): Non-Accidental Connection	55
3.2	Grasp Improves with Reason-Responsiveness.....	57
3.2.1	Why Normative Reasons?	58
3.3	Other Accounts of Grasping	61
3.3.1	Is Understanding Always Normative?	62
3.3.2	Varieties of Understanding	62
3.3.3	Accessibility and Reflection	65
3.3.4	Epistemic Luck	66
3.3.5	Stevens – Recognition	70
3.3.6	Hills – Cognitive Control	71
3.3.7	Grimm – Modal Grasping.....	72
3.3.8	Bourget – Phenomenal Account	73
3.3.9	Bengson – Noetic Account	74
3.4	Advantages	75
3.4.1	Distinctiveness of Understanding	76
3.4.2	The Value of Understanding	79
3.5	Conclusion.....	83

Chapter 4: Understanding and Truth

4.0	Introduction	85
4.1	How Exemplification Enables Reason-Responsiveness.....	86
4.1.1	Elgin and Understanding	86
4.1.2	What is Exemplification?	87
4.1.3	What are the Success Conditions?	89
4.2	Truths in the Content of Understanding	93
4.2.1	Approximately, p	94
4.2.2	Counterfactual Truths	96
4.3	Understanding as a Means to Truth	98
4.4	Truth-Oriented Values	99
4.4.1	Approximate Truth	100
4.4.2	True in a Model	102
4.4.3	‘Something True’	103
4.5	Conclusion.....	105

Chapter 5: Being a Good Environmental Thinker

5.0	Introduction.....	107
5.1	Environmental Virtue Epistemology.....	107
5.1.1	Guidance	108
5.1.2	Moral Responsibility and Related Attitudes	110
5.2	Epistemic Virtues	114
5.3	Environmental Virtues.....	117
5.4	Being Epistemically Virtuous as an Epistemic Aim or Concern.....	119
5.5	Plan for Chapter 6 and 7	121
5.6	Conclusion.....	122

Chapter 6: Outward-Lookingness and Seeing the Bigger Picture

6.0	Introduction.....	123
6.1	Outward-Lookingness	123
6.2	Other Virtues	126
6.2.1	Curiosity.....	127
6.2.2	Open-Mindedness	129
6.2.3	Far-Sightedness.....	130
6.2.4	Holistic Thinking	132
6.3	The Intelligibility Problem	134
6.3.1	The Intelligibility Problem.....	134
6.3.2	Disposition to Make Intelligible	136
6.3.3	Order of Priority.....	138
6.3.4	Awareness of Unawareness.....	139
6.4	Conclusion.....	140

Chapter 7: Honesty

7.0	Introduction.....	141
7.1	Is Honesty an Environmental Virtue?.....	141
7.2	What Does it Mean to Be Honest?	144
7.2.1	The Unification Challenge and the Scope of Honesty.....	144
7.2.2	Avoid or Explain Counterintuitive Results	145
7.2.3	Corresponding Account of Vices	147
7.2.4	Explanation of Value.....	149
7.3	Honesty, Deception and Truth	150
7.4	Honesty and Truth	153

7.4.1	Not All Deception Matters to Honesty	154
7.4.2	Truth and Deception	157
7.5	My Account	160
7.5.1	Explanation of Value and the Scope of Honesty	161
7.5.2	Account of Corresponding Vices.....	162
7.6	Applications.....	163
7.6.1	Testimony	163
7.6.2	Paternalism	164
7.6.3	Self-Deception	165
7.7	Conclusion.....	167

Chapter 8: Why Use Case Study Research in Developing an Epistemology for the Environment?

8.0	Introduction	169
8.1	Defending the Use of Empirical Research in Applied Epistemology and Clarifying the Research Question.....	170
8.1.1	Applied Epistemology and Empirical Research.....	170
8.1.2	Interdisciplinary Approach	172
8.1.3	Research Question	173
8.2	Philosophy and Qualitative/ Quantitative Research Methods.....	174
8.2.1	Ontological Position	175
8.2.2	Epistemological Position.....	177
8.2.3	Theory and Research	179
8.2.4	Qualitative vs Quantitative	179
8.3	Research Design	180
8.4	The Nearshore Trawling Byelaw	182
8.4.1	What is the Nearshore Trawling Byelaw?	182
8.4.2	What is the Help Our Kelp Campaign and the Sussex Kelp Recovery Project?.....	183
8.4.3	Selection of Case Study.....	184
8.5	Interviews.....	184
8.5.1	Research Methods.....	184
8.5.2	Planning and Organising Information Gathering	186
8.5.3	Interview Questions	187
8.6	Approach to Analysis.....	190
8.7	Ethics	192
8.7.1	Professional Integrity.....	192
8.7.2	Informed Consent	193

8.7.3	Anonymity	193
8.7.4	Safety of Participants	194
8.7.5	Data Storage.....	194
8.8	Conclusion.....	194

Chapter 9: Themes and Reflections

9.0	Introduction.....	197
9.1	Epistemic Themes	197
9.1.1	Evidence.....	198
	Evidence in the Consultation Process	198
	Anecdotal Evidence	200
9.1.2	Awareness	203
	Out of Sight, Out of Mind.....	203
	The Film and Visual Imagery.....	204
9.1.3	Simplicity.....	206
9.1.4	Hope.....	209
9.1.5	Understanding Value.....	211
	Arriving at an Understanding of Value	211
	Sharing Understanding of Value	212
	Misunderstanding Value.....	213
9.1.6	Trust	213
9.1.7	Memory.....	215
9.2	Other Important Themes	217
9.2.1	Objections	217
9.2.2	Collaboration	218
9.3	Further Conclusions	219
9.3.1	What is in the Data and the Theory?.....	221
9.3.2	What is in the Data but not the Theory?	221
9.3.3	What is not in the Data but is in the Theory?	222
9.4	Reflections.....	223
9.5	Conclusion.....	224

Conclusion	225
-------------------------	------------

Appendices

Appendix 1	229
Appendix 2	231
Appendix 3	235

Works Cited	237
--------------------------	------------

Acknowledgements

I would like to thank my supervisors James Andow and Susan Hodgett. Thank you both for believing in my ideas, for all the helpful discussions and for giving me so much good advice and support throughout this process, for which I am incredibly grateful.

This research was funded by a Faculty of Arts and Humanities Postgraduate Studentship.

Thank you to the participants in this research for taking the time to share your perspectives.

I also massively appreciate all the support I've had from friends and family. Thank you to my friends in Norwich. Thank you to my Grandparents. And thank you to my sisters

Michaela, Grace and Hannah.

This would not be possible without my amazing parents. Thank you Mum and Dad for your encouragement throughout this project and for always being there for me. It's very much appreciated.

Introduction

Many of the philosophical questions that arise from our relationship to the natural environment and present environmental changes concern what we ought to believe, what we are in a position to know or understand, or other aspects of our epistemic position. Along with others, I think epistemology should have something to say about this (Coady, 2019; Levy and Brownstein, 2021). This thesis contributes to the task of developing an epistemology for the environment; by that, I mean an epistemology that focuses on questions that matter to those living an environmentally conscious life. It's not possible to develop a comprehensive epistemology for the environment within this thesis, so the following section introduces a more specific focus and motivates the three questions that the thesis will address.¹

0.1 The Epistemic Problem

In discussion of obstacles to meaningful action to mitigate climate change, Dale Jamieson says that climate change poses threats that are 'probabilistic, multiple, indirect, often invisible, and unbounded in space and time. Fully grasping these threats requires scientific understanding and technical skills often in short supply' (2014: 61). This means that:

Climate change must be thought rather than sensed, and we are not very good at thinking. Even if we succeed in thinking that something is a threat, we are less reactive than if we sense that it is a threat. Consider the difference between touching a hot stove and being told that the stove is hot. Scientists are telling us that the world is warming, but we do not sense it and so we do not act. This is the hardest problem to overcome. Any approach to coming to terms with climate change must respect these facts about ourselves (Jamieson, 2014: 103).

¹ In order to demonstrate a range of relevant epistemological questions, here are some examples of existing debates in epistemology that engage with environmental issues. Given that it is demanding to be in the know about environmental issues, some address the question of culpable ignorance about the environment (Robichaud, 2017). Some epistemologists seek to answer questions surrounding expert testimony on environmental issues (Anderson, 2011; Almassi, 2012). Moreover, an issue that I address directly later in this thesis is the role of epistemic virtues in Environmental Virtue Ethics (Stafford, 2010; Kawall, 2018).

We are ill-equipped to grasp the complex reality of a changing climate, because it is unfamiliar and cannot be sensed and observed in the way that problems that we are more accustomed to can be.²

The complexity of environmental issues presents an everyday challenge for those who want to live a moral life. Jason Kawall writes:

To be a good environmental citizen will require information concerning the environmental impacts (both positive and negative) of various activities in which we might participate...in particular, to be a good environmental citizen will likely be especially demanding epistemically, given the wide range and complexity of issues at stake (2009: 110).

Kawall (2021) also points out that the epistemic demands on the environmentally virtuous person are significant, because of the complexity of the issues, but also because of the sheer amount of relevant information that bears on their everyday choices.

We cannot do the right thing by the environment unless we can think about the environment in the right way and thinking about the environment is difficult. It's difficult to think about environmental issues because this will involve processing vast amounts of complex information. We might worry that those who want to live an environmentally conscious life appear to be persistently faced with the task of comprehending the incomprehensible. However, we can often make enough sense of very complex things, without having to fully comprehend them, so the epistemic problem seems to be that environmental issues require us to be able to consistently find a way to process many complex things. I will focus on this epistemic problem because it arises from structural features of environmental issues, and is therefore pervasive in this context, and because it clearly presents a serious barrier to those who want to do right by the environment.

An epistemology for the environment must recognise the features of ourselves and our natural environment from which the epistemic problem arises, features that are often obscured by highly ideal theories.³ To see why, consider an argument for when a non-ideal

² For more detail on difficulty understanding complex systems, see Sterman (2011).

³ Many question the assumptions of ideal theories in epistemology. We can see this in the distinction between traditional and social epistemology. For example, feminist epistemologies take into account the social context

approach to epistemology is preferable: certain idealisations of ourselves and our environments recommend the wrong epistemic norms and ideals, not because they are unattainable in light of our limitations or unfavourable environments, but because attempting to live by these norms and ideals will make us worse off epistemically (McKenna, 2023: 24). There is good reason to think that failing to recognise the epistemic problem may lead us astray in this way; if we are working from an epistemic ideal that ignores the epistemic problem, and the challenges it produces in thinking about the environment, then attempting to live by this ideal may not help us overcome it. In the following section, I set out the plan for this thesis by clarifying three questions that arise from the epistemic problem and my approach to answering them.

0.2 The Plan

I aim to make progress on three questions that arise from the epistemic problem I have outlined in Section 0.1. This thesis primarily addresses the epistemic problem by endorsing an understanding-oriented approach to an epistemology for the environment, according to which understanding takes priority, over knowledge and truth, because it is understanding that renders the complex world intelligible by bringing what is important to the forefront of our minds.

The first question asks which epistemic states matter to an epistemology for the environment that seeks an answer to the epistemic problem. I argue that an epistemology for the environment should shift its attention to understanding, in particular its grasping element, because this is the epistemic state that renders complex matters intelligible. Chapter 1 motivates the idea that grasp is of particular interest to an epistemology for the environment, due to its capacity to enable agents to respond to the world in the right way, where knowledge may fail. I argue that this presents a challenge to veritists who claim that epistemic normativity is fundamentally about truth. This veritist view cannot explain the importance of grasp. To defend these claims, we need to say more about the nature of grasp. The following chapters argue that grasp is a relation to facts that objectively count in

of knowers (Anderson, 1995: 54). By recognising the social context of the knower, other important concepts are revealed, such as epistemic injustice (Fricker, 2007). In *Non-Ideal Epistemology*, Robin McKenna (2023: 3) complicates this distinction by recognising that we can take both an ideal or a non-ideal approach to social epistemology.

favour of responses from certain agents (sometimes called objective normative reasons). Chapter 2 defends some of the minimal commitments that follow from conceiving of grasp this way. Chapter 3 further explicates the nature of grasp and its relationship to understanding. I argue that the degree to which one grasps *X* corresponds to the degree to which one is in an epistemic position to respond to objective normative reasons presented by *X*. Chapter 4 argues that representations of the world contribute to understanding by selectively emphasising our most important objective normative reasons and this can require deviating significantly from the truth. Building on the account of grasp and its relationship to understanding developed in Chapters 1-4, I argue that an epistemology for the environment should be understanding-oriented because understanding, in virtue of its grasping element, selectively emphasises important features of the world and thereby renders the complexity of the world intelligible. This is essential for thinking about the environment in a way that enables us to do right by the environment.

The second question asks what matters from the perspective of the epistemic agent who wants to think about the environment in the right way. In Chapter 5, I argue that an adequate epistemology for the environment should give guiding epistemic advice and grounds by which to evaluate others, and this requires answering the question of what makes a good environmental thinker. I turn to responsibilist epistemic virtues to answer this question. I then provide two demonstrations of this, adopting the understanding-oriented approach. Chapter 6 concerns epistemic virtues that help us to ‘see the bigger picture’. Chapter 7 concerns the epistemic virtue of honesty.

The third question concerns whether my focus on the epistemic problem of processing vast amounts of complex information related to environmental issues is the right focus. I have chosen this focus because this epistemic problem is pervasive in environmental contexts and clearly presents an obstacle to action, but these are not decisive reasons to think that this is the only relevant epistemic problem or even that it is the most important. The nature and prevalence of other relevant epistemic problems is an empirical question about peoples’ experiences. In Chapter 8, I argue that an epistemology for the environment should engage in empirical research and I defend the methodology for a case study I conducted to demonstrate this. Chapter 9 presents the results of this case study research.

Chapter 1: Epistemic Aims and Concerns

1.0 Introduction

Due to the vast amount of complex information we must navigate in our attempts to understand environmental issues, it is my contention that an adequate epistemology for the environment will not take truth to be fundamental to epistemic normativity. In other words, it will not be veritist. As will become clear, many arguments against veritism are motivated by considerations about the epistemic lives of real agents such as their limitations and the complexity of their environment. For this reason, it is unsurprising that anti-veritist arguments are relevant in this context. It is in our attempts as limited beings to apprehend the complexities of our changing environment that these issues are most salient.

The central aim of this chapter is to challenge the claim that epistemic normativity is fundamentally about truth. Section 1.1 clarifies what is meant by the veritist claim that truth is the sole fundamental epistemic aim or concern. Section 1.2 argues that truth is not the sole fundamental epistemic aim or concern because, if this were the case, we could not account for the variety of epistemic standings that we aim for and value. In particular, I argue that grasping, an epistemic standing often required to move us to do the right thing, is an epistemic aim or concern that cannot be explained in terms of truth. Section 1.3 proposes that, because grasp has this guiding quality, and knowing or believing the truth can sometimes fall short, grasp seems to be the epistemic relation that facilitates our response to reasons. This makes grasping an important epistemic aim or concern. Another veritist claim might be that, although truth is not the *sole* fundamental epistemic aim or concern, it is particularly important or somehow central to our epistemic evaluations. Section 1.4 builds on a related complaint against veritism, known as the trivial truths problem, to argue that, even if truth is among our fundamental epistemic aims or concerns, it is less important than grasping. Section 1.5 demonstrates how a picture of epistemic normativity that centres grasping, rather than truth, can avoid the following unpalatable outcomes: (i) a wrong kinds of reasons problem, (ii) an inability to account for distinctiveness of epistemic normativity and (iii) the result that epistemic normativity is entirely instrumental to other sources of normativity.

1.1 The Epistemic Point of View

We can shed light on what is meant by ‘epistemic normativity’ by comparison to the normative realm of morality. While morality concerns evaluations that are made from a moral point of view, epistemic normativity concerns evaluations that are made from a distinctively *epistemic* point of view. This is more than simply an evaluation that is made about a belief. To see this point, consider that a belief can be right or good from the moral or practical perspective and right or good from an epistemic perspective. It is also possible for these evaluations to conflict. For example, it might be epistemically good for us to suspend judgement on whether someone is a thief (if we have no evidence either way), but morally good for us to believe they are not a thief (if morality says we are to treat others as innocent until proven guilty). It might also be epistemically good to tell our friend that, truthfully, we hate their new haircut or say nothing at all, but it might be practically and morally good for us to tell them that we love their new haircut, since sharing our opinion or revealing it via silence is not worth the upset.⁴ So, epistemic normativity is not simply normativity that governs belief, but evaluations that are relative to a distinctively *epistemic* set of aims or concerns.

The nature of distinctively epistemic aims and concerns is the subject of much debate in epistemology. However, many influential thinkers maintain that epistemic normativity is in some sense fundamentally about the truth, because we, as epistemic agents, seek to acquire truth and avoid error:

[T]he primary function of cognition in human life is to acquire true rather than false beliefs about matters that are of interest or importance to us... I don't know how to prove that *the acquisition, retention, and use of true beliefs about matters that are of interest and/or importance* is the most basic goals of cognition...but I suggest that anyone can see its obviousness by reflecting on what would happen to human life if we were either without beliefs at all or if our beliefs were all or mainly false (Alston, 2005: 29).

[W]hat makes us cognitive beings at all is our capacity for belief, and the goal of our distinctively cognitive endeavours is *truth*: we want our beliefs to correctly and accurately depict the world (BonJour, 1985: 7).

⁴ Some think that that practical and moral considerations enter in to justification or knowledge attributions (Fantl and McGrath, 2010; Basu, 2021).

Information seeking is a pervasive activity of human life. We scan the horizon to see if rain is imminent, we watch the news to learn who has been elected, and we listen to the traffic report to anticipate delays in our drive to work... What we seek in all such cases is true or accurate information, not misinformation. No newspaper reader wants the sports page to misreport outcomes or the financial section to falsify market transactions. We commonly seek truth, or a close approximation to the truth (Goldman, 1999: 3).

Any view that says epistemic normativity is fundamentally about truth could be described as veritism, but this is a vague claim that incorporates many views and so I will focus on a precise claim.⁵

One precise, and intuitive, claim that supports the idea that epistemic normativity is fundamentally about truth is that truth is a *fundamental* epistemic aim or concern. An epistemic aim or concern is fundamental if it is not explained by any other aim or concern – a fundamental epistemic aim or concern might also explain derivative aims or concerns.⁶ For example, if truth is the fundamental epistemic aim or concern, then truth is not explained by any other aim or concern and might explain why evidence-gathering is a derivative epistemic aim (because by gathering evidence we arrive at truth) or perhaps why knowledge is an derivative epistemic aim (because by knowing, we have a true belief that is in some way robust).⁷

To get even more clear on how truth is fundamental to epistemic normativity, we should attend to the distinction between the following claims:

⁵ There are different ways to make sense of truth as an epistemic aim or concern. It is commonly thought to be about the two aims of believing truth and avoiding errors, which William James describes as ‘first and great commandments as would-be knowers’ ([1987] 2014: 17).

⁶ I deliberately opt for this terminology of ‘aims and concerns’ to remain neutral on other questions. For example, I am not committed to the thesis that epistemic normativity depends, at least in part, on our goals or that we can make sense of epistemic normativity in terms of aims at all. Another option would be to talk of fundamental ‘values’, but to use this terminology to describe a reductive view of epistemic normativity would imply that value is prior to other normative concepts.

⁷ Note that the question of whether truth is the fundamental epistemic aim or concern is a separate question from what the *source* of epistemic normativity is. Some think that epistemologists are required to produce an unmysterious account of the source of epistemic normativity, e.g. why are we obliged to believe the truth (Kornblith, 1993: 158-9). The question of the content of epistemic aims and concerns will speak to an explanation of the source of their normativity, but these are separable questions. We might claim truth is the fundamental epistemic aim or goal, but we can still ask a further question of in virtue of what truth has this status.

- (i) Truth is a fundamental epistemic aim or concern.
- (ii) Truth is the *sole* fundamental epistemic aim or concern.⁸

It is important to distinguish the above claims. Note that claim (ii) is a strong veritist view that says all epistemic normativity bottoms out in truth. On this view, truth is important to epistemic normativity because anything that counts as an epistemic aim or concern is more fundamentally explained by reference to truth. We cannot say that anything is an epistemic aim or concern, unless it is related to truth in some way. Claim (i) is weaker because it just says that truth is among the fundamental epistemic aims or concerns, which entails it is not merely derivative of some other aim or concern. This does not in principle suggest anything about how important truth is to our epistemic evaluations. There could be other fundamental epistemic aims or concerns that are much more important than truth, such that truth is almost always overridden by them. As a result, there is no guarantee that such a claim preserves the ‘spirit’ of veritism, which plausibly requires that truth is of central importance. From here, I can see two clear ways to say epistemic normativity is fundamentally about truth. The first is that truth is the sole fundamental epistemic aim or concern, such that all other epistemic aims or concerns reduce to truth. The second way includes truth among the fundamental epistemic aim or concerns, but maintains that truth is an *important* epistemic aim or concern relative to any others.⁹

Although veritism has been met with a lot of recent objections, veritism does have an intuitive pull and a variety of benefits.¹⁰ Most notably, there is a certain obviousness to this way of thinking. Consider William Alston’s comments on this:

[T]he idea that it is important for human flourishing to be guided by correct rather than incorrect suppositions about how things are, where this is of interest or importance to us, is so obvious that it would seem to be unnecessary to belabor the point (2005: 30).

⁸ Ralph Wedgwood defends a claim like this, proposing that the fundamental epistemic norm of correct belief is that ‘a belief is correct iff the proposition believed is true’ (2002: 272). Pritchard (2014) defends claim (ii) against prominent counter arguments.

⁹ For an argument for pluralism about epistemic desiderata, see Alston (1993).

¹⁰ Duncan Pritchard (2014) goes as far as to say, that the ‘new orthodoxy’ is that truth is not fundamental.

In other words, we want our beliefs to match reality because if they do not match reality, then we will be led astray in our efforts to live a good life. A point in favour of this view is that it illuminates something distinctive about epistemic normativity that contrasts with other normative sources in the right sort of way. The epistemic is thought to be about matching the content of our beliefs to the reality of the world, i.e. having beliefs that are true, which would explain the conflicts with practical and moral evaluations. From the epistemic point of view, we want our beliefs to be true regardless of practical or moral impact. Any challenge to veritism will have to answer to these intuitions.

In what follows, I reject both ways of saying that epistemic normativity is fundamentally about truth. I present a challenge to veritism by making the case for another fundamental epistemic aim or concern and therefore denying claim (ii). I do not deny claim (i), that truth could be *among* our fundamental epistemic aims or concerns. I will, however, go on to argue that, where believing or knowing the truth presents itself as an *important* aspect of epistemic normativity, it inherits this importance from another epistemic aim or concern. These arguments in conjunction decentre truth in epistemic normativity.

1.2 Truth is not the Sole Fundamental Epistemic Aim or Concern

Some claim that veritist views are too reductive, because if all epistemic normativity is explained in terms of truth, then we cannot explain why we should care about or value other epistemic standings, such as knowledge, justified belief, understanding and wisdom.¹¹ This section advances a similar argument. I argue against veritism, by arguing that there is an important and distinctive epistemic success, best described as ‘grasping’, that is among our epistemic aims and concerns.

¹¹ One example of this is the *swamping problem*. Some argue that veritism cannot explain why knowledge is more valuable than true belief because, according to a view where all that matters is truth, once we have true belief, knowledge cannot add to its value (Zagzebski, 2003).

Let's start by considering the following cases:

Animal Suffering

S knew that they were morally required to avoid contributing to unnecessary animal suffering, but this did not 'sink in' such that it affected their life. However, some of their recent experiences led them to see things differently. S saw a video of a cow attempt to run after their calf, as the calf was being carried away. After this, S began to feel guilt for consuming products that lead to unnecessary animal suffering and chooses to avoid them.

Anthropocene

S gains new knowledge about how our planet is changing. S also gains new vocabulary that helps them view the world in a new light. For example, S uses the term 'endling' to refer to the last living member of a species. This idea evokes a sense of loneliness and great loss. S also adopts the concept 'anthropocene' as a way of thinking about their own life within a series of geological epochs.

Plastic Pollution

Although S knows that plastic pollutes the oceans and they know certain facts about this problem, e.g. that the Great Pacific Garbage Patch is a 1.6 million square kilometres floating mass of artificial waste in the middle of the Pacific Ocean, S is not able to fully appreciate the scale of the problem.

These cases demonstrate significant variety. *Animal Suffering* is a moral realisation case, where the realisation is not fully described in terms of knowing or believing that the morally significant conditions obtain or that they are morally significant. *Anthropocene* features a hermeneutical improvement that is not captured in terms of knowledge or belief about the concepts and meanings. *Plastic Pollution* is a difficulty apprehending scale. Initially, the epistemic success in question appears to be to do with not just affirming that something is the case but making sense of what it means and being able to respond to it. It is natural to describe these cases in terms of 'grasping'.¹² In *Animal Suffering* and *Anthropocene*, S moves from just knowledge to grasping. In *Plastic Pollution*, S has relevant knowledge but fails to grasp.

¹² David Bourget (2017) uses cases like this to demonstrate instances where grasping makes a difference; the first two are cases of *transition*, where S transitions from non-grasping to grasping, and the last case is non-grasping, where it is not possible for S to arrive at grasping. I will go on to propose a different account of these cases than Bourget, which will be defended at length in Chapter 3.

Recently, there has been much interest in the notion of grasping, in particular its role as a distinctive feature of the epistemic state of understanding. It is crucial to point out though that there is much disagreement about the nature of grasping and, as observed by Michael Strevens, ‘grasp is a word with many meanings, related to understanding in various ways’ (2024: 742). In light of this, it is not clear that the sense of ‘grasp’ at issue in the above cases is the same sense of grasp discussed by epistemologists in the understanding literature. I should remain open, at this stage, to the possibility that a different sense of ‘grasp’ is at issue in these cases, but I will continue to use this terminology to describe them.

In light of the above cases, there are at least three ways to maintain a reductive, veritist epistemology. The first is to exclude these cases from the remit of epistemology by denying that they pertain to an *epistemic* success or failure. The second is to claim that grasping is epistemic, but that it can be reduced to knowledge or true belief. The third is to say that, while the nature of grasping might not reduce to knowledge or true belief, the status of grasping as an epistemic aim or concern can be otherwise explained in terms of truth. Although I will need to develop an account of grasping to fully support my claims, I will now address each of these possible objections with the aim to make it clear how these cases are supposed to present a problem for veritism.

1.2.1 Is Grasping Epistemic?

As established in Section 1.1, epistemic evaluations are made from the perspective of distinctively epistemic aims and concerns. It might be argued that what is going on in cases like *Animal Suffering*, *Anthropocene*, and *Plastic Pollution* is not epistemic and is therefore properly understood in terms of considerations that fall outside our epistemic aims and concerns, such as an improvement in S’s practical rationality or moral development.

One could attempt to explain *Animal Suffering*, for example, as a case where S becomes (or fails to become) practically rational. The argument seems to be that this case does not demonstrate a change in epistemic position but is instead a case where S goes through something that reminds them to adjust their preferences, desires or motivations. However, this line of argument overlooks a crucial point. The alignment of S’s beliefs and intentions, such that S becomes practically rational, does not come out of nowhere. It is caused by a

prior improvement in S's cognitive contact with the world. It is because the world affects S differently than it did before, due to some kind of input beyond what they knew to be the case already, that S is moved to change their intentions. S is not merely reminded to update preferences or intentions in accordance with their values, S undergoes a realisation that leads them to reform their values.

A similar point can be made for the objection that that S improves morally. While S might learn something about morality by watching the video in *Animal Suffering*, the transition that takes place in doing so is not merely that S becomes a better person who responds to knowledge in a moral way, but that S achieves contact with the world that pushes them to do so. S arrives at a better epistemic position in virtue of watching the video. Similarly, in *Anthropocene*, S's epistemic position improves in virtue of learning the relevant concepts.

1.2.2 Knowledge and True Belief?

Another way to deny that grasping is a candidate for a fundamental epistemic aim or concern is that the nature of grasping reduces to knowledge or true belief. Let's first consider some ideas about true belief and knowledge. Consider the idea that a belief that p is true just in case it is descriptively successful, in that its content describes reality in a way that fits reality.¹³ Many consider that a belief that p aims for truth in some way; from the first-person perspective, the question of whether p and whether you believe p, seem to have the same answer. Nishi Shah and J. David Velleman (2005) say that to conceive something as a belief is to apply a standard to an attitude of 'regarding-p-as-true', such that it is correct if p is true and incorrect if p is not true. On this view, true belief is correct belief. From this, I gather that there is some agreement that true belief is an attitude that affirms that p is the case and succeeds in the limited sense that p is in fact the case.

Moving on to knowledge and its relationship with truth, knowing appears to be factive, in the sense that if we know p, we can infer that p is true and if p is not true, then we do not know p. Some evidence for this claim is the distinction between factive and non-factive expressions; factive expressions such as knows that, sees that and notices that can only

¹³ This interpretation is from Goldman's (1999: 59) discussion of a version of the correspondence theory of truth.

apply in relation to truths and non-factive expressions such as believes that, thinks that and hopes that, can apply in relation to either truths or falsehoods (Nagel, 2017: 530; Williamson, 2000: 34). Knowing that p also seems to involve an affirming attitude towards p and, in this sense, it is a success partly because p is true. Although grasp in the above cases may correspond with a success or failure in knowing or truly believing that, grasping is not explained on these terms.

In *Animal Suffering*, S already knows and truly believes that unnecessary suffering is wrong. But, in virtue of their direct experience with suffering in nonhuman animals, S gains some further appreciation for its moral significance via grasping. In order to show that grasping in these cases can in fact be reduced to knowledge, it might be argued that there is some other piece of acquired knowledge that explains this improvement. One candidate might be that believing or knowing that “this animal here is suffering”. However, this is not the epistemic improvement in question. S does not improve their behaviour because they see any particular animal suffering, S goes on to improve their behaviour because they appreciate something about the moral significance of animal suffering in general. I am interested in what it is about S’s epistemic position that moves them to feel differently about animal suffering and to act on it.

Consider the other two cases. In *Anthropocene* S improves by acquiring certain concepts. I include this case because, it seems to me, S improves by gaining these concepts, not because S knows or truly believes something new about the world, but because they frame or organise the knowledge in some way. S might gain knowledge such as ‘the word ‘endling’ makes me feel like ‘x, y, z’ or ‘endling means last individual in a species’ or ‘humans have a significant impact on the planet’, but none of these beliefs seem to capture the work these concepts do to improve S’s epistemic position. In *Plastic Pollution*, S knows facts about the plastic problem, but it is difficult for them to fully appreciate its scale. What is epistemically missing can seem quite mysterious. For example, despite the fact S might be able to make many calculations with the number 1.6 million and identify its place on a number line, something is missing in their capacity to appreciate the size of this space and what it means for the plastic problem. This does not seem to be a matter of failing to know or believe that something is the case. In fact, the more we know and the more accurate our beliefs are about larger scales, e.g. as they are represented in statistics, the more unintelligible they may be to us.

At least on the face of it, grasping has less to do with affirming what is the case and more to do with is to do with making sense of the world, such that we can figure out what it means for us and how we should respond, and this goes beyond believing and knowing the truth. I have not fully established that grasping is distinct from knowledge, which is something I aim to do in Chapters 2, 3 and 4, but I have attempted to at least highlight the features of these cases that suggest that it is.

1.2.3 Explaining Grasping as an Epistemic Aim or Concern in Terms of Truth?

Finally, even if grasping is epistemic and does not reduce to true belief or knowledge, we could still argue that its status as an epistemic aim or concern is more fundamentally explained by reference to truth and any additional value that grasping has can be attributed elsewhere. This is Pritchard's (2014) response to the 'swamping problem' objection to veritism. The swamping problem says that if knowledge is instrumentally valuable to true belief, then we cannot maintain that knowledge is more valuable than true belief; once we have already arrived at the truth, no further value is added by the fact the belief is known. Pritchard argues that this outcome is not so bad once we consider the difference between epistemic value and the value of the epistemic. Even if true belief and knowledge have the same *epistemic* value, we may value knowledge more than true belief because knowledge has additional non-epistemic value, such as practical utility and contribution to a good life. Perhaps a similar point applies to grasping.

One attempt to produce this sort of counter argument could be that believing the truth carries epistemic value in virtue of its truth, but believing the truth in a certain way (e.g. via certain visual imagery in *Animal Suffering* or certain concepts in *Anthropocene*) carries practical or moral value over and above this, by moving us to do the right thing more effectively, for example. However, the issue that grasping raises is disanalogous. It is not plausible that truly believing and grasping have the same epistemic value in the same way that one could make the case that knowing and truly believing do.

It is only plausible that knowing and truly believing have the same epistemic value if they are the same way of getting things right, i.e. representing the world accurately, and whatever knowledge adds is just instrumental to this way of getting things right – that's

how the swamping problem arises. That's also why it makes sense to claim that knowledge and true belief have the same 'epistemic value' and that we might find the value of knowledge exceeds true belief along other axes of value. In contrast, grasping and truly believing (on my view) are not the same way of getting things right and the former is not instrumental to the latter. Grasping is supposed to go beyond affirming what is the case, meaning it cannot be construed as merely instrumental to the acquisition of true belief. I, therefore, cannot see the same sort of basis for claiming that the epistemic value of grasping can wholly be explained in terms of truth and any additional value must be non-epistemic. In the next section, I say more about how we should think of grasping.

1.3 Why is Grasping Important?

In order to raise a related objection against veritism, we need a better idea of what is going on in the above cases and in what sense the use of the term 'grasp' here goes beyond knowing or believing what is the case, such that we are better able to use information or respond to information. This section proposes that grasping is to do with reason-responsiveness, in that it provides the sort of cognitive contact with facts that facilitates our ability to respond to them as reasons.¹⁴ If this is right, then grasping is an important epistemic aim or concern, without which our beliefs could not fulfil their function of guiding us in our reasoning.¹⁵

It seems to be that by grasping the world, we become able to do the right thing. This is what is notable about the above cases. In *Animal Suffering* and *Anthropocene*, S achieves 'grasp' and becomes better able to do the right thing. In *Plastic Pollution*, such a shift is not possible for S. For our epistemic states to guide us in this way, such that we not only endorse the truth but become able to respond to it appropriately, then there must be a sense in which these epistemic states facilitate this responsiveness to reasons. Reasons are considerations that count for or against certain responses and when we are in a position to respond to them, we are in an epistemic position to respond in the right way. To my mind,

¹⁴ Linda Zagzebski (1996: 45) uses the phrase 'cognitive contact with reality' as a placeholder to include different epistemic states, such as knowledge, understanding and wisdom. I am interested in a certain kind of cognitive contact with a certain part of reality: our reasons. I have chosen this way of talking because it allows me to remain neutral at this point about the nature of this epistemic relation.

¹⁵ It is an important part of many arguments in epistemology that our epistemic states have a role in guiding or explaining actions (Littlejohn, 2012; Railton, 1994; Whiting, 2010: 217; Williamson, 2002: 101).

this is just what is meant by being able to ‘use’ information or see what it means for us. What precisely is required for one to be in this sort of cognitive contact with one’s reasons is up for debate and I will contribute to this debate in the following chapters.¹⁶ For now, I want to highlight how this proposal about grasping bears on veritism.

What is particularly problematic for veritism is that grasping appears to provide cognitive contact with reasons that knowledge and true belief can be lacking. Note that the fact that true beliefs figure in whether we can respond to the world appropriately, and false beliefs will lead us astray, is part of the veritist argument for why true beliefs are of value and false beliefs are not (Alston, 2005: 30). Even veritist arguments recognise that what is true or known may guide, enable or move us in our responses to the world. This supposed to be part of the appeal of true belief being fundamental to epistemic normativity. If these sentiments are on the right track, then what I have argued in the previous section suggests that beliefs may be unable to fulfil this crucial function, unless grasping is involved. In the next section, I explore (and ultimately reject) the idea that, due to its relationship with reasons, grasping is more fundamental than truth because the status of truth as an epistemic aim or concern is more fundamentally explained by grasping.¹⁷

1.4 The Problem of Trivial Truths

I have suggested in the previous section that grasping has an important relationship to reason-responsiveness. This idea lends support to another argument against veritism: *the problem of trivial truths*. To see this point, consider the following passage from Ernest Sosa:

At the beach on a lazy summer afternoon, we might scoop up a handful of sand and carefully count the grains. This would give us an otherwise unremarked truth, something that on the view before us is at least a positive good, other things equal. This view I hardly understand. The number of grains would not interest most of us in

¹⁶ As I will go on to discuss in Chapter 2, there are different views on what it takes to possess a reason. Some examples include that possessing a reason requires knowledge (Littlejohn, 2017; Williamson, 2002), being in position to manifest relevant know how (Lord, 2018: 97-124), and awareness (Silva, 2021).

¹⁷ A similar point, that the role of belief in action might be prior to the demand on belief to be true, is made by Daniel Whiting: ‘One might view the aim of avoiding falsity as reflecting the more fundamental aim of belief, namely that of enabling one to act in such a way as to realise one’s goals, an aim which is typically thwarted by false belief’ (2010: 217).

the slightest. Absent any such antecedent interest, moreover, it is hard to see any sort of *value* in one's having that truth (2003: 156).

Some truths, such as the number of grains of sand in a handful, are so trivial that they don't seem to be worth having at all.¹⁸ This section uncovers a problem for veritism that is raised by trivial truths, which is that often our interest in the truth is more fundamentally explained by the fact this truth presents a reason.

Initially, the idea that believing or knowing the truth is only important when that truth presents a reason appears to support the claim that grasping is more fundamental to epistemic normativity than truth because truth, where it is an epistemic aim or concern, would be derivative of grasping. However, I end up opting for a weaker position that says, while the problem of trivial truths cannot exclude truth from being a fundamental epistemic aim or concern, it does suggest that truth itself is not an *important* epistemic aim or concern relative to grasping. Before I get to this, I first address existing literature on the trivial truths problem in order to make clear why this is the genuine problem at issue rather than other possible versions of the problem.

1.4.1 What is the Trivial Truths Problem?

In our approach to the trivial truths problem, we should be aware that the fact that some truths are more worthy of belief *simpliciter* than others is no challenge to veritist epistemology, because veritism is a thesis that specifically concerns what matters from the *epistemic* point of view. The veritist can say both that true beliefs about trivial matters have the same epistemic value as true beliefs about important matters *and* that beliefs about important matters have more value overall, in virtue of other non-epistemic values. It is, however, important to note that some truths appear more worthy of belief because of what they are about, such as truths about the beginning of the universe.¹⁹ Moreover, these beliefs appear to have more epistemic value than beliefs about trivial matters, such as grains of sand on a summer afternoon, because of what they are about, even when both are morally and practically irrelevant.

¹⁸ Given that both accuracy and triviality come in degrees, the picture is likely to be more complicated.

¹⁹ Many epistemologists accept the idea that some truths are better, from an epistemic perspective, because of what they are about (Goldman, 1999; Sosa, 2003; Grimm, 2008).

One possible veritist explanation for why we might prefer to believe truths about the beginning of the universe, over truths about grains of sand, is that the former involves *more* truth than the latter. Nick Treanor (2014) argues that, despite our tendency to describe a true proposition as one truth, we can see that one proposition can contain more than one truth and thereby one proposition might contain more truths than another. For example:

- (i) Arran has a son.
- (ii) Arran has a son and a daughter.

Proposition (i) and (ii) might be both be true, but (ii) contains more truth than (i), in the sense that if you know (ii) you know *more* than if you know (i). Moreover, not that we cannot easily determine how much truth is in a complex proposition by deducting from grammatical structure, such as the use of ‘and’ above. Consider that, from use of the word ‘daughter’, (ii) contains more than one additional truth, e.g. Arran has another child and Arran has a daughter. As Treanor puts it, ‘every word brings a sea of truth with it’ (2014: 557). It is difficult to count amounts of truth.

Treanor is right that this difficulty counting truth means the fact we prefer to believe some true propositions over others does not establish that truth is not fundamental to epistemic normativity, because some true propositions contain more truth than others. However, this does not explain why some trivial truths appear to have no value at all.²⁰ In other words, there is not one good thing to say, from the epistemic perspective, about a true belief in trivial matters.²¹ In my view, we should deny that trivial truths have a small amount of epistemic value (alongside the claim that truth is the sole epistemic aim or concern) because, if they do, it remains a theoretical possibility that a long string of higher order beliefs about one’s own beliefs could have more epistemic value than a smaller number of beliefs about a significant scientific theory or moral argument (DePaul, 2001: 173). It is conceivable that an introspective process with enough power would beat any set of beliefs about interesting matters from the epistemic perspective, all other things equal. Moreover, a plausible explanation of the contrast between believing truths about grains of

²⁰ Some epistemologists do argue, however, that we have an interest in having true beliefs, simply because they are true, because we have an inherent interest in the truth (Foley, 1987: 11; Kvanvig, 2003: 41; Lynch, 2005: 55).

²¹ This concern isn’t specific to talk about epistemic value. Some endorse ‘clutter avoidance’ norms that say we ought not believe in trivial matters (Harman, 1999; Friedman, 2018).

sand and believing truths about the beginning of the universe should not only refer to the complexity of propositions, but also what they are about. Some truths are valuable from the epistemic perspective because of what they are about, i.e. because they are interesting, and not just because they are more complex propositions and therefore contain more truth.

Despite all this, it's still not clear what problem trivial truths are supposed to raise for veritism. Even if we grant that some truths have no value and others have value because of what they are about, this does not entail that truth is not a fundamental epistemic aim or concern:

[S]omeone who opts for the deep scientific truth about the universe over the trivial truth about, say, the number of spoons in a kitchen cupboard does so because she cares about truth. This person reveals their concern for the truth in displaying this preference. The same goes for someone who pursues the important line of inquiry, albeit knowing that it might be unsuccessful, over the trivial inquiry that she knows in advance will likely lead to success (albeit of a pointless kind). We immediately recognise that this is just what caring about the truth involves (Pritchard, 2023: 16).

It seems right that caring about the truth does not entail that we care about all truths. It also seems right that we care about beliefs about the beginning of the universe because we want to know the truth about how the universe began. If this is right, then believing or knowing some truths over others is among our epistemic aims or concerns in a fundamental sense, since it is worth knowing and believing them *because* they are true and this is not explained by any other aim or concern.

The fact trivial true beliefs have no value does not present a genuine problem for veritism, but we are now in a position to see what would present a genuine problem of trivial truths. A genuine problem of trivial truths could be that the truth about how the universe begins is valuable from the epistemic perspective because of *a feature of this belief that is not truth*. This result would threaten the idea that truth is even among our fundamental epistemic aims and concerns because, where truth apparently matters for its own sake, this turns out to be more fundamentally explained by some other feature. The next sub-section puts forward an argument with this structure.

1.4.2 Truth and Reasons

One candidate for a feature of true belief that always more fundamentally explains its worth from the epistemic point of view is that the truth one believes is a reason. We must explore this possibility in light of what I have already argued in this chapter. It is easy to see why practically or morally relevant truths are not trivial: they present moral and practical reasons. We can also argue that truths about how the universe began, that are interesting because of what they are about, also present reasons because we can make sense of interestingness in terms of reasons.

One way to understand the interesting is to say that what is interesting is whatever we take interest in. However, along with others, I think this gets things the wrong way round. It is not that the interesting is whatever we take interest in, but that we should take interest in whatever is interesting (Brady, 2009; Grimm, 2011). In other words, we need an interest-independent way to make sense of what makes a truth interesting because there could be agents that are interested in how many grains of sand are in a handful and we want to say that they are wrong to have this interest. We can make sense of the interesting in terms of truths that present a reason for us to have a range of responses, such as curiosity, wonder or the feeling that our curiosity is satisfied and the cherishing of truth.²² This is a way to make sense of the claim that trivial truths should not mean anything to us (we have no reason to find them interesting) and important truths do mean something to us (we have good reason to find them interesting). This could suggest that providing cognitive contact with reasons, i.e. grasping, is more fundamental than truth because it is only when a truth presents a reason that it is worth having.

However, I find this argument unconvincing. Firstly, it does not solve the issue at hand because even trivial truths may present reasons, such as reasons to be bored or direct our inquiries elsewhere, so we still cannot say that trivial truths are not worth believing because they do not present any reasons. Secondly, this view could not account for truth norms on belief (Shah and Velleman, 2005). This is problematic, given that there does seem to be something wrong with having a false belief about trivial matters, even if there is nothing particularly valuable about having true belief in trivial matters:

²² Two views could be distinguished here. A view that says that truth itself is valuable and a view that says truth is valuable because of how we respond to it, e.g. because we seek truth and we cherish truth once it is found. Nenad Miscevic (2007: 258) calls the latter view *response-dependence*.

If, despite my utter lack of interest in the question of whether Bertrand Russell was left-handed, I stumble upon strong evidence that he was, then I have strong epistemic reasons to believe that Bertrand Russell was left-handed. Indeed, my epistemic reasons will be no different than they would be if I had acquired the same evidence deliberately, because I *did* have the goal of finding out whether Russell was left-handed. Once I come into possession of evidence which strongly supports that claim that p, then I have epistemic reasons to believe that p, regardless of whether I presently have or previously had the goal of believing the truth about p, or any wider goal which would be better achieved in virtue of my believing the truth about p (Kelly, 2003: 625).

There seems to be something that counts against a false belief from the epistemic point of view. Because of these considerations, it seems that truth is among our epistemic aims and concerns in a way that is not explained by the truth presenting a reason.

Despite this, I do see a lesson to learn from the genuine problem of trivial truths and it is a lesson that still manages to decentre truth in epistemic normativity. For me, the lesson is that, where believing or knowing the truth is of any plausible importance, the source of this importance is to do with the fact this truth presents a reason. The argument goes that knowing or believing the truth is as important as the response that this truth licenses from us, whether this is a practical, moral, or purely epistemic response. Therefore, the importance of believing the truth is derivative of a more fundamental epistemic aim or concern that enables our response to these reasons: grasping.

I acknowledge, however, that this move sides with one party in debate about what epistemology should centrally concern itself with. I have argued against the spirit of veritism, but I have done so by prioritising substantive normativity. Substantive normativity concerns prescriptive claims, e.g. questions about what we ought to believe, what we should inquire about, where to turn our attention etc., as opposed to non-substantive normativity, which may concern what makes an instance of a particular type of thing good or bad with respect to its type, e.g. how we should evaluate instances of knowledge and belief.²³ I have already provided some reasons to think that the important

²³ Jane Friedman (2020) argues that zetetic norms (norms of inquiry) are epistemic norms and are worthy of concern in epistemology. Whereas J. Adam Carter (2024) argues that substantive norms like these may be a

questions in epistemology concern substantive normativity, and I will attempt to defend this position against some anticipated objections, but one point I particularly want to emphasise here is that these are the questions in epistemology that have the greater bearing on our epistemic lives.

We can see that both arguments in Section 1.2 on grasping and Section 1.4 on trivial truths stem from similar concerns I have raised about the epistemic problem of processing vast amounts of complex information about the environment in the Introduction, Section 0.1. The grasping cases, like those I have listed in Section 1.2, often draw attention to how our minds and our epistemic practices work; we are not just interested in listing true information, we want to be able to use it in our reasoning. Additionally, the problem of trivial truths recognises that we are finite beings with limited time, resources and attention. It makes sense that these same worries would arise for our inquiries into the natural environment and environmental challenges, given that our limitations are most salient in our attempts to comprehend complex matters. At the very least, these are the epistemological issues that matter to an epistemology for the environment.

1.5 Objections

One important thing to make clear is that I have argued against reductionist views about epistemic normativity that claim that the truth is the sole fundamental epistemic aim or concern and views that claim truth is central, but I am still open to the possibility of other epistemic aims or concerns that are neither explained by truth or reasons. Nevertheless, we might think that there are some easily anticipated objections to the idea that grasp is an important epistemic aim or concern due to its capacity to guide us, whereas truth is not. This section addresses the following concerns: (i) that we cannot solve the wrong kind of epistemic reasons problem, (ii) we cannot explain the distinctiveness of epistemic normativity and (iii) we must admit that epistemic normativity is always in service to other sources of normativity.

part of intellectual ethics, but this is not within the remit of epistemology proper, which concerns norms governing attempts at knowledge.

1.5.1 Wrong Kinds of Reasons

I take it that it would be a bad outcome if we could not distinguish between the right and wrong kind of reasons for belief, such that the right kind of reasons are the important epistemic reasons relative to our epistemic aims and concerns. A wrong kind of reason for belief might arise when someone offers you a £10 reward to believe that the sky is blue. The £10 reward might be a reason to have the belief that the sky is blue, but it is not the right kind of reason to believe that the sky is blue because it does not have a bearing on whether it is true that the sky is blue. Given that the right kind of reasons for belief refer to whether the belief is true, and it is in some way constitutive of what it means to believe that we tend to be persuaded by them, it would be a bad outcome if these reasons are not epistemic reasons or not important epistemic reasons relative to the picture of epistemic normativity I have proposed.²⁴

Fortunately, my view does not prevent us from excluding the wrong kinds of reasons. The suggestion that grasping is a fundamental epistemic aim or concern does not entail that epistemic normativity is simply about having epistemic states that result in the right responses to the world or that lead us to the best outcomes, such that we could have an inaccurate take on the world, but it would be epistemically good if it results in us doing the right thing. This view of epistemic normativity would admit the practical benefit of being in a belief state within the set of epistemic reasons for belief, but this is not what my view says about epistemic normativity.

We can see this by paying attention to the distinction between epistemic states that are practically best for us and epistemic states that provide cognitive contact with practical reasons. Consider the choice between two toasters, where we carry out a cost-benefit analysis on each to determine which toaster has the highest expected value and therefore which toaster we ought to buy. There is an important difference between beliefs that provide cognitive contact with practical reasons to buy the best toaster and beliefs that are practically best. In Hilary Kornblith's words, assuming that what is practically best is to do with maximising happiness, the former 'will not tell us what the actual consequences are of

²⁴ Pamela Hieronymi (2005) argues that the distinction between the right and wrong kind of reasons for belief is that the wrong kind of reasons are extrinsic, in that they bear on whether a belief is good to have, and the 'right' kind of reasons have to be constitutive, in that they bear on a question that, when settled, amounts to having that belief.

purchasing each toaster; instead it will tell us what we would be happiest to believe the consequences to be' (Kornblith, 1993: 371). In order to be in cognitive contact with practical reasons to buy a particular toaster, such that we are able to make a decision by responding to these reasons, we need to have epistemic states that represent the reality of the toaster and the consequences of our purchase. These are the sorts of epistemic states that we need to mobilise our practical reasons in our practical reasoning in a way that facilitates a decision about the toaster, based on those reasons. This might have little to do with what is practically optimal to believe, e.g. perhaps believing that your toaster is much better than it is brings about better practical outcomes.

This means that epistemic reasons for belief are not directly related to the practical benefits of a belief, e.g. a £10 reward or false beliefs that overestimate the value of the best toaster, but instead bear on whether one's epistemic states reflect the reality about the world in a way that can guide us to respond in the right way. Normally, these will be considerations that bear on the truth of our beliefs, but I argue in Chapter 4 that this picture is not so simple.

1.5.2 Distinctiveness of the Epistemic

As I mentioned in Section 1.1, we tend to think of epistemic evaluations as distinct from other sorts of evaluations. Firstly, epistemic aims and concerns appear to be distinct, and may even conflict with moral or practical considerations, e.g. it is coherent to say that we can do morally right by carrying out epistemic wrongs (Kornblith, 1993: 368). The second is that what makes an evaluation epistemic is not that its object is belief, because considerations such as moral or aesthetic or practical benefits of having a particular belief do not bear on its epistemic evaluation (Kornblith, 1993: 368). We might worry that similar criticisms can be brought against my view that grasping is a fundamental epistemic aim or concern or that it is more important from the epistemic point of view than truth.

Once we recognise the distinction between what it means to have cognitive contact with practical, moral or epistemic reasons and what it means for a belief to be good from the practical and moral perspective, we are in a position to see how epistemic normativity is distinct from other normative sources in a way that produces conflicts. If we are offered £10 to believe something false, e.g. that the moon is made of cheese, the conflict is

between the practical benefit of being in the belief state that the moon is made of cheese and the epistemic loss that occurs when we are not in cognitive contact with reality and the reasons it might present. At least with respect to this isolated belief, it makes sense to say that, from the perspective of practical normativity, we should believe the moon is made of cheese, but from the perspective of epistemic normativity, we should not.

1.5.3 Epistemic and Other Sources of Normativity

It might be objected that, if we think that grasping is more important to epistemic normativity than truth because of its relationship to reasons, this means epistemic normativity is only important when in service to other sources of normativity such as morality or practical normativity. This seems to be an unwelcome result because it suggests the epistemic is only of importance because of its relationship to morality or practical normativity, such that worthwhile epistemology is just a form of ethics.²⁵

However, recall the case of believing how the universe began discussed in Section 1.4. The question of how the universe began is interesting, regardless of the practical or moral benefits of answering this question. I have suggested that these truths are reasons for responses such as curiosity, finding something interesting, or cherishing truth. These reasons appear to be purely epistemic, in the sense that they are not in service to other sources of normativity. This means that, on my view, epistemic normativity still has important aspects that are not in service of other sources of normativity.

1.6 Conclusion

In this chapter, I have argued that truth is not the sole fundamental epistemic aim or concern and that it is not an important epistemic aim or concern, because this view cannot explain the epistemic standing of grasping or solve the genuine trivial truths problem. The approach to epistemic normativity I propose is more in line with the epistemic aims and concerns of our epistemic lives, as instantiated by our inquiries about the environment. However, the notion of ‘grasping’ remains mysterious. Unless grasping is unpacked

²⁵ Some explicitly take issue with the collapse of epistemic normativity into ‘ethics of belief’ or the ‘moralising’ of epistemology (Miscevic, 2007; Code, 2020 [1987]).

further, the arguments in this chapter are not sufficiently supported. In Chapter 2, I begin to present a detailed account of grasping by arguing that it involves cognitive contact with facts that objectively speak in favour of certain responses.

Chapter 2: Objectivism and Externalism

2.0 Introduction

In the previous chapter, I argued that veritism cannot explain the status of grasping as an epistemic aim or concern. I also argued that, where believing and knowing the truth is important to our epistemic aims or concerns, it inherits its importance from this truth presenting an important reason. However, the notion of grasping I refer to in Chapter 1 is still quite vague. I suggested that grasping is a certain form of cognitive contact with reasons, but two very important questions arise from this. The first is the question of what these reasons are. The second is the question of what precisely this cognitive contact with reasons amounts to; without an answer to this second question, grasping might just be a form of knowledge. This chapter focuses primarily on the first question. Chapter 3 and 4 make progress on the second.

Section 2.1 proposes that grasp is an epistemic relation to facts that objectively speak in favour of certain responses, such facts are sometimes referred to as *objective normative reasons*. I also clarify the sense in which this proposal might take sides in debates between subjectivists and objectivists about reasons and between externalists and internalists in epistemology. However, the overall aim of the following sections is not to settle the contentious debates between objectivists and subjectivists about reasons and externalists and internalists, but to demonstrate that that the commitments that follow from aligning grasping with objective normative reasons turn out to be minimal and not as controversial as they might first appear.

Many are sceptical of objective normative reasons, claiming that they do not qualify as reasons at all or do not inform normative notions, such as rationality or ought claims. Section 2.2 argues that my view does not rest on controversial claims about the relationship between objective normative reasons and these other normative notions. However, there may still be room to claim that objective normative reasons do not matter and therefore it is not valuable to be in cognitive contact with them. Section 2.3 argues that it is implausible to deny that objective normative reasons matter, particularly in the context of this thesis

because environmental ethics needs them. Because objective normative reasons are external to our perspective, my proposal conceives of grasp as an externalist notion. Section 2.4 presents some reasons why we should think of grasping as a valuable externalist success.

2.1 Objective and Subjective Reasons

I start this section with the following idea. There are features of the world that mean something to me, in that they matter to me, and some features of the world that don't mean much to me, in that they don't matter to me. Importantly, when a feature of the world genuinely does mean something to me (and I am not merely mistaken in thinking that it does), it is not just that this feature of the world *does* bring about a response in me (in fact, it may not), it is that there is a sense in which it would be a good thing if this feature of the world did bring about this response in me. The features of the world that matter to me license a particular response in me by producing pressure to respond in a certain way or at least making a particular response appropriate. For example, the presence of a poisonous snake counts in favour of believing there is a poisonous snake, feeling afraid and leaving the area. Moreover, the extent to which a feature of the world matters to me plausibly corresponds to the importance of the range of responses that it licenses. I suggest we make sense of all this in terms of objective normative reasons. There is much disagreement about the nature of reasons and their relationship to rationality, ought claims and what is reasonable. I will introduce some of these distinctions here in this section before clarifying where I stand on these debates.

Firstly, it is usual to distinguish normative reasons from other reasons, e.g. explanatory reasons and motivating reasons.²⁶ A normative reason for S to ϕ is a consideration that counts in favour of S ϕ -ing. My thesis about grasping will concern normative reasons because, as argued in Chapter 1, grasp is an epistemic success that can guide us to do the right thing or respond in the right way to world. While reasons that motivate us may also guide us, having motivating reasons is not indicative of the epistemic success at issue. A motivating reason is what operates in my reasoning, such that it makes my actions intelligible from my own perspective. What motivates me to take a sip from a glass might

²⁶ Maria Alvarez (2010) distinguishes three roles of reasons: justifying, explaining and motivating.

be that I think it contains gin, even if it in fact contains poison.²⁷ Motivating reasons like this do not necessarily guide us to do the right thing or respond in the right way. It is normative reasons that guide us to respond to how things are in the world.

Secondly, there's a divide about whether what we ought to do is in some way relative to our perspective. Some think that what we ought to do is objective, such that it is independent of our perspective (Thomson, 1990; Graham, 2010). Some think what we ought to do is dependent on our perspective (Zimmerman, 2008; Kiesewetter, 2011; Gibbons, 2013). Objective reasons are independent from our perspective, e.g. facts, true propositions, or states of affairs, and subjective reasons depend on some aspect of our perspective, such as our apparent evidence or beliefs.²⁸ These are separate issues, given that we could argue that reasons are objective facts, but only those we know or have evidence for determine what we ought to do. In this chapter, I do not commit to any particular view on ought claims, but I suggest that grasping concerns the objective conception of normative reasons.

Grasping concerns the objective conception of normative reasons, I argue, because grasping is supposed to be a way that we have got things right about the world. Grasping is an epistemic success. We should be immediately suspicious about the task of defining an epistemic success in terms of epistemic limitations because the epistemic success I discuss in Chapter 1 is supposed to be a way of getting the world right and subjective reasons arise when we get the world wrong. Subjective reasons, that are in some way dependent on our perspective, can be a product of false belief or misleading apparent evidence.

While it seems intuitive to say that grasp is an epistemic success and therefore should not be thought of in terms of mistake, this view does carry some commitments that need a defence, especially given that, on some views of subjective reasons, it is coherent to say that they can be either possessed or unpossessed. To see this, consider a view of subjective

²⁷ This example is from Bernard Williams (1979: 102).

²⁸ Some relativise subjective reasons to desires. Williams (1979) argues that p is only a reason if you are able to respond for that reason and, given the Humean view that says you cannot ϕ for p unless you think that ϕ -ing for p will serve one of your desires or motives, p can only be a reason if it serves your desires or motives. For more in favour of desire-relative reasons, see Schroeder (2007). Although I think there are important points to make about this issue, some of which may even follow from the points I make in this chapter, I unfortunately do not have the space to present these arguments and doing so would lead the discussion too far astray from the aims of this chapter.

reasons (the factoring account) on which this would not be coherent. The factoring account says subjective reasons *just are* possessed objective reasons, e.g. reasons we know, are aware of or have evidence for. This account says that there is an epistemic relation we can have to our reasons such that we possess them and subjective reasons just are objective reasons when we achieve this. However, following Mark Schroeder (2008), subjective reasons do not seem to be simply possessed objective reasons because objective reasons are facts about the world, whereas some of our subjective reasons concern the content of our false beliefs, such as there being gin in our glass when there is not. When I falsely believe that there is gin in my glass, I may have a subjective reason to take a sip, but my subjective reason in this case cannot be a possessed objective reason, because there is no gin in my glass.

Moreover, not all subjective reasons are possessed:

Consider an ordinary participant on the famous *Let us Make a Deal* show (also known as Monty Hall show). Let us call him Ben. He stands in front of three doors. There is a luxury car behind one of those doors and goats behind the two others. Ben has been given a chance to choose one of the three doors. Let us say Ben chooses door number 1. Now, the show's host, Monty Hall, is obliged to open one of the three doors (that is the rule). Yet, he is not permitted to open the door that Ben has chosen. Neither can he open the door with the car behind it. Ben knows the rules of the game. Let's say Monty Hall opens door number 2. Of course, there is a goat behind it. Next, Ben is offered the chance to change his initial choice. That is, Ben can change his choice from door 1 to 3 (Logins, 2018: 265).

Given Ben's epistemic circumstances, Ben should opt to change his choice to door 3. Although it may not seem like it, Monty Hall provides this information by opening door 2 and revealing the goat. Ben's first choice had a $1/3$ chance of being the car, meaning there is a $2/3$ chance that Monty Hall had to choose door 2 (because the car was in door 3). This means that, from Ben's perspective, there is a $2/3$ chance the car is behind door 3 and only a $1/3$ chance it is behind door 1. When playing this game, the participant will win $2/3$ of the time if they switch doors.

In this case, participants have a subjective reason to switch doors. This is because, relative to Ben's epistemic circumstances, the fact that Monty Hall opened door 2 is a

reason to switch doors. This is a subjective reason because it is relative to the participants' available evidence, which is an aspect of their perspective. Given that Ben, along with many other participants, is unable to recognise this reason, this subjective reason is unpossessed in the sense that Ben is not in an epistemic position to respond. If Ben is unaware about how Monty Hall's reveal of the goat affects what he should do next, he has an unpossessed subjective reason to switch doors. This means that it would be a coherent task to define what it takes epistemically to possess subjective reasons.²⁹ If it is subjective reasons that determine what we ought to do, for example, then they are a part of normative reality and possession of them would involve a notable epistemic success after all. Contrary to my initial dismissal of this idea, we have to say more about why the epistemic success of grasping is to be construed in terms of objective normative reasons, rather than subjective reasons.

In fact, my proposal that grasping is to do with objective normative reasons is entangled in two debates. The first is the debate between objectivists and subjectivists about reasons. I need to defend the role of objective normative reasons in grasping and show that my view is not the result of mistaken assumptions about what reasons are, what they are for and their relationship to doing the right thing. The second debate is whether grasping is an externalist success that concerns one's relation to objective facts or an internalist success to do with responsiveness to subjective reasons that are relative to our perspective, rather than the facts. This chapter tackles the important task of clarifying the sense in which I take side in these surrounding debates (in particular between the subjectivists and objectivists about reasons and the internalists and externalists in epistemology) and therefore defending the role of objective normative reasons in this thesis.

2.2 What are Facts that Objectively Count in Favour?

Objective normative reasons are quite mysterious and we might be sceptical of their status as reasons, especially because it is unclear how they can behave like reasons, e.g. how they figure in our practical and theoretical reasoning. The purpose of this section, and the following section, is to argue that I must endorse only some of the least controversial claims about objective normative reasons in order for them to play the part they are

²⁹ This idea could possibly fit well with other accounts of grasping that are more internalist-leaning.

supposed to play in an account of grasping: (i) there are facts that count in favour of certain responses for particular agents and (ii) it is a good thing to be in an epistemic position to respond to them.

Claim (i) and (ii) are crucial to my arguments. If grasping is to be about cognitive contact with ‘objective normative reasons’, there must be facts that objectively count in favour in this way and it must be valuable to be in an epistemic position to respond to them. These claims are, in principle, separable from more controversial claims about the relationship between reasons and other normative notions, such as ought claims and rationality. Although Chapter 1 suggests that grasping guides us to do the ‘right thing’, what is meant by this is just that grasp guides us to properly respond to features of the world in a fitting manner, even if this does not amount to doing what we ‘ought’ to do or what is rational. If my view is only committed to uncontroversial claims like (i) and (ii), it will not be subject to many of the criticisms made against objectivism.

To see this, let’s consider why we might think objective normative reasons are not genuine normative notions. Because objective normative reasons are independent from our perspective, they can be hopelessly out of our reach:

Day’s End: Every day, B comes home and flips a light switch. One day, flipping the light switch leads to a series of unpredictable coincidences that results in B’s neighbour, A, being badly injured.³⁰

In this case, B has an objective reason not to flip the light switch, but subjectively, given their apparent evidence and their belief that the light switch is safe, they have a reason to flip the light switch. B, in their epistemic circumstances, has no way of responding to their objective normative reasons. Of course, there is something that counts in favour of B not switching the light on, which is that B can avoid causing A’s injury. However, we might question whether this is sufficient to confer any normative status.

Perhaps the normative status stems from the fact that it would be a good thing if B were to respond to the fact that turning on the light would injure A. However, John Gibbons

³⁰ Paraphrased from Thomson (1990).

denies that we can generate normative status in this way by inviting us to compare cases like *Day's End* with cases like *Avalanche*:

Avalanche: There's an avalanche coming down the mountain, about to hit a ridge. If it goes left, it will destroy the village. If it goes right it won't. Clearly it would be better if the avalanche went to the right (Gibbons, 2010: 346).

Even though it would be a good thing if the avalanche went to the right, this does not mean the avalanche has a reason to turn right.

We could say that this is to do with the fact that B has a mind and the avalanche does not; the idea is that reasons are normative notions and normative notions apply to things with minds. But, it seems *ad hoc* to simply restrict reasons to things with minds (Gibbons, 2010: 347). A theory needs to not only restrict the sort of objective facts that can amount to normative reasons to the facts that can apply to agents with minds, but also explain why normative reasons are restricted in this way. Gibbons argues that objectivism cannot provide this explanation because the presence of B's mind makes no difference in *Day's End*. If B were to walk in and fail to turn on the light, they would not be doing so intentionally for the reason of avoiding A's possible injury.

Another attempt to explain objective normative reasons might appeal to the idea that the fact of A's possible injury would be a normative reason if B were omniscient, i.e. normative reasons are to do with what would be reasonable if one knew all the facts. But this doesn't seem to work either. Firstly, if normative reasons are to do with whatever an omniscient version of oneself would want one to do, then we would never have most reason to learn anything, because rather than advising oneself to learn some important fact p, we would rather just advise ourselves of the truth of p directly (Gibbons, 2010: 349). Moreover, the omniscience approach fails for similar reasons as the good outcome approach, because if B does what the omniscient version of themselves wants them to do, the presence of B's mind is still irrelevant to the status of the fact of A's possible injury as a reason. B is not capable of doing what the omniscient version of themselves wants them to do intentionally.

The underlying issue here is that objective normative reasons do not seem to have a hold on the agent in the actual circumstances from which they arise. Because of this, we might suspect that these objective facts do not deserve to be called reasons and should not be considered normative things.³¹

I do not counter these arguments against objectivism, but instead point out that they do not prevent us from making claim (i), that there are facts that objectively count in favour of certain responses, or claim (ii), that it is a good thing to be in an epistemic position to respond to them. I do not need to compete for the word reason to defend claim (i) or (ii) (although, for convenience I will continue to use this term, even if it is ultimately figurative). Neither do I need to insist that these reasons are normative in the restrictive sense that an agent must be capable of acting for this reason intentionally, in the actual circumstances in which the reason arises. If subjectivists wish to deny that objective facts can be genuinely normative reasons, that's fine with me. I only insist on claims (i) and (ii), which turn out to be very implausible for subjectivists to deny.

First, notice that it is the *fact* of A's possible injury that counts in favour of not flipping the switch and not B's beliefs.³² If B were to list reasons for or against turning the light on, we would expect them to list facts, not that they *believe* certain facts. That B has a belief that the light switch is safe is not itself a fact that speaks in favour of turning on the light. This suggests that, even from the first-person perspective, deliberation concerns facts and not beliefs or the possibly false contents of beliefs. Although we have considered the position that objectivism undermines the role of reasons in reasoning, because facts cannot always move us in cases like *Day's End*, it seems that our practical and theoretical reasoning concerns facts that count in favour of certain responses. Therefore, the subjectivist who denies the very existence of facts that objectively count in favour of certain responses cannot explain the role of reasons in our reasoning either.

³¹ For what it's worth, I am optimistic that these challenges can be met. After all, it is not just that it would be a good thing if B did not flip the switch, but it would be a good thing if B did not do so *for the reason of preventing A's injury*. This would afford the action moral worth, for example, and this may be the relevance of B having a mind and therefore being capable of acting for reasons and not just acting in accordance with them.

³² Many make a similar point about reasons. Jonathan Dancy (2000: 52) claims it is 'blatantly obvious' that many of our moral duties have their grounds in facts, not beliefs; it is the fact that someone is in trouble that makes it such that I ought to help them, not that I think that they are in trouble. One point in favour of this view is that, when asked to explain an action, an agent will cite the *fact* that counts in favour of that action, not their *beliefs* (Scanlon, 2014: 54-55).

In addition, it is implausible to deny that it is a good thing if B is in an epistemic position to respond to the reason that A's possible injury presents. The fact that it is good for B to become better informed is not the main point of contention in subjectivist and objectivist debates. In fact, an important objection raised against the subjectivists begins from this idea that it must be in some way better for B to become aware of A's possible injury and respond accordingly. Therefore, the subjectivist theory must accommodate this fact, by means other than the obvious explanation, which is that B becomes informed of their *reasons* (Graham, 2010: 91; Kiesewetter, 2017: 211-212). So far, it seems that I am committed to two claims that are quite uncontroversial and potentially not even the subject of debate between subjectivists and objectivists about reasons.

Regardless of how we settle the question of whether objective normative reasons are real reasons, whether they contribute to genuine ought claims or how they relate to rationality, I have suggested that we can maintain that there are facts that objectively count in favour of certain responses from particular agents and that it is good to be in an epistemic position to respond to them. For clarity, I will continue to refer to such facts as objective normative reasons.

2.3 Objective Reasons and Environmental Ethics

I have defended what I take to be an uncontroversial position that there are facts that objectively speak in favour of certain responses and that it is a good thing to be able to respond to them. However, it is possible to maintain that there are these things called objective normative reasons, and perhaps they are real reasons or perhaps they are not, but either way these objective normative reasons are not an important part of normative reality. This view also appears to have something going for it, particularly if you deny that objective normative reasons are the sorts of things that determine what we ought to do, what is morally required, or what is reasonable and rational for us to do, because if this is the case, then it is not obvious why such things would matter or why it is a good thing that

we are able to respond to them.³³ This would undermine the idea that grasping is an important epistemic success.

In this section, I argue that objective normative reasons matter to normativity, even if they do not determine what is rational or what we otherwise ought to do, because if we abandon this idea, we arrive at many unpalatable outcomes, some of which are particularly pertinent in environmental ethics. A view that denies the importance of objective normative reasons has a hard time making sense of right and wrong beyond the human perspective. In other words, there may be some scope for denying the importance of objective normative reasons (even if we accept that they exist and that it would be good if we could respond to them), but this will not be a position that appeals to many environmental ethicists and, therefore, this view is even less relevant in the context of this thesis which seeks the epistemic state that enables responsiveness to environmental issues.

2.3.1 Non-Anthropocentrism in Environmental Ethics

The reason that denying the importance of objective normative reasons is unlikely to appeal to environmental ethicists is that many take interest in a sense of right and wrong beyond the human perspective.³⁴ This sub-section clarifies why this is important to Environmental Ethics.

Non-anthropocentrism can be defined in terms of stronger and weaker versions of the idea that nonhuman beings have direct or independent moral significance. Weaker forms of non-anthropocentrism that say non-human beings can have ‘non-instrumental value’: value that is independent of contribution to some human end or benefit (O’Neill, 1992: 119). Other definitions may decentre humanity in morality further than this, by claiming that human concerns are no more central to morality than nonhuman concerns. The general idea behind a non-anthropocentric ethic is that it is not solely human concerns that matter to morality and the mistake that only human concerns matter has led to environmental crises.

³³ In Gibbons’ words, ‘either normative reasons don’t tell you what to do at all, or they tell you the wrong thing’ (2010: 353). If normative reasons don’t tell you what you ought to do, what is rational, or what is reasonable, then it seems that whatever they tell you to do is the wrong thing. If normative reasons tell you do the wrong thing, then perhaps it is unclear why we should care about them.

³⁴ Many have argued that nonhuman beings (or nature as a whole) have direct or independent moral significance to some extent. There are examples of this argument in animal ethics (Regan, 1985; Singer, 1991; Jamieson, 1998). There are also examples in virtue ethics (Hursthouse, 2006; Sandler, 2007).

However, many rightly take interest in the claim that the correct anthropocentric view, that is informed and appropriately far-sighted, will converge with non-anthropocentrism and both will ultimately support the same actions and policies. This is known as the *convergence hypothesis* (Norton, 1991: 237-243). The convergence hypothesis may well be broadly accurate for most policies and actions, but this does not necessarily undermine the importance of morality beyond the human perspective.

An important argument here is Katie McShane's (2007) argument that, while anthropocentrism might license the right actions in line with non-anthropocentrism, anthropocentrism cannot explain which are the right feelings. We might think it is important to have certain feelings for their own sake. For example, it is good to love your friends. Consider that to love one's friend is to have an other-centred emotion, in the sense that we see one's friend as having a value beyond and independent of what one's friend can add to your life. McShane argues that anthropocentrism is incompatible with having the right attitudes and feelings toward nature, such as love, respect and awe. When we love nature, we should also think of it as having a value that does not serve our interests:

insofar as anthropocentrism is true, we are making a kind of mistake when we love the land, respect nature, are in awe of the vastness of the universe, or take other attitudes that are incompatible with thinking that their object's only value is in serving our interest (McShane, 2007: 177).

To say that we should love nature, in a way that does not involve this mistake about its value, requires a non-anthropocentric view.

Non-anthropocentrism might also be relevant to epistemic normativity. In Section 1.4.1, I argued that we should be curious about aspects of the world for their own sake. To be curious about the natural world, for its own sake, and not because it serves our interests in some way, may be an important epistemic attitude that we might have toward the natural environment. However, anthropocentrism must say that this attitude involves a *mistake* about the things we ought to be curious about. It would be a strange result if being curious about aspects of the natural environment for their own sake involves a mistake about its epistemic value (because according to anthropocentrism, we should only be curious about the natural environment if this in some way serves our interest). This remains a

strange result if we can still justify being curious about the natural environment for its own sake in terms of how having this attitude benefits human beings in some way, e.g. being curious about the natural environment for its own sake ends up furthering understanding – if this were the case, the value of having this mistaken attitude could be justified by the convergence hypothesis. Just as it would be strange if loving the environment for its own sake involved a mistake about its value, it would also be strange if being properly curious about the natural environment for its own sake involved a similar mistake about why we should be curious. I conclude from these considerations that, while the convergence hypothesis might be right, in that anthropocentrism and non-anthropocentrism tend to support the same actions, admitting that it is not solely human concerns that matter remains an important part of environmental ethics and an adequate epistemology for the environment.

In light of these considerations, there are bad outcomes we might try to avoid. We should avoid saying that human belief is the sole source of the moral significance of non-human beings, such that what we believe about non-humans produces their moral significance, and we should avoid saying that what we believe about non-human beings can significantly constrain or eliminate their moral significance. If what we believe about nature produces the moral significance of nature or limits it, then we risk an anthropocentric ethic that says that nonhuman beings are only morally significant if we believe they are. This is in tension with a picture of environmental ethics that says there is value in the natural environment for its own sake, irrespective of what we believe about it, and that allows for the possibility of love of nature and curiosity about nature for its own sake. I argue that any view which denies that objective normative reasons matter will result in this bad outcome.

2.3.2 Why Does Environmental Ethics Need Objective Reasons?

The bad outcome follows from some ways of relativising normative reasons to one's epistemic position and, in my view, these are views we should reject. Other ways of thinking about this do not seem to lead to the bad outcome, and perhaps we should not reject them, but these are the views on which objective normative reasons do matter. Again, the key takeaway here is not a controversial claim. My aim is to highlight the

tension between the denial that objective normative reasons matter and the commitments of a plausible environmental ethic.

To see which views seem to lead to the bad outcome, we should pay attention to the distinction between subjective perspectivism and objective perspectivism (Kiesewetter, 2018). The former says that what we ought to do depends on the content of potentially false beliefs (actual or counterfactual) and the latter says that what we ought to do depends on facts that are available to us. Notice how subjective perspectivism means that some of the reasons that determine what we ought to do will not be or correspond to objective facts, but will instead be the contents of false beliefs (such as the proposition that there is no poison in my glass, in the case where there is). In contrast, objective perspectivism entails that all reasons that determine what we ought to do are objective facts, but not all objective facts that count in favour of responses contribute to what we ought to do.

Objective perspectivism does not lend support to the bad outcome. This view does not entail that human believers produce or significantly restrict normative reasons that govern our interactions with nonhuman beings, because normative reasons will always be objective facts. Of course, objective perspectivism also supports the view that objective normative reasons matter; on this view, being an objective normative reason is a condition on being a reason that contributes to what we ought to do. Objective perspectivism avoids the bad outcome and does so in a way that makes objective normative reasons an important part of normative reality.

The extent to which subjective perspectivism results in the bad outcome depends on which version of this view one endorses. One way to claim that reasons depend on our epistemic circumstances is to say that it is possible for us to have reasons to do something, just because we believe we ought to do something. This is a heavily criticised view, as many rightly take issue with the idea that we can simply believe genuine reasons into existence (Kiesewetter, 2017: 84). Any version of subjective perspectivism that permits believing reasons into existence in this way delivers implausible results in general, but we should find it particularly implausible in the context of this thesis because any view that makes it such that we can produce genuine normative pressure by simply believing that we ought to do something will lead to the bad outcome. Not only does this view say that we would have no moral reasons to prevent extinction events or other environmental harm if

we don't believe we are required to, but it also says we would have moral reasons to perpetuate them just because we believe that we are required to. This means all our moral reasons that concern our interactions with non-human beings are limited by what humans believe about their value.

However, there are more complex ways that our reasons might depend on our beliefs. Consider the following case:

While walking in some desert, you have disturbed and angered a poisonous snake. You believe that, to save your life, you must run away. In fact you must stand still, since this snake will attack only moving targets (Parfit, 2011: 34).

Parfit says that you have beliefs in this case that, if true, would present a reason to act (apparent reasons) and you are rational to act on these beliefs and run away from the snake. We can distinguish this factual mistake and a mistake about what our reasons count in favour of. Consider *Future Tuesday Indifference*, in which Parfit (1984: 123-124) imagines a man who is indifferent to agony as long as it occurs on a future Tuesday. This takes the fact that the 15th October is a Tuesday to book his operation on the 15th October. In this case, there is no mistake about the facts, but this man is irrational because there is no reason to be indifferent to agony just because it is a Tuesday. In this case, the man makes a mistake about what reasons count in favour of and, unlike your response to false beliefs about snakes, this appears irrational. It appears that beliefs about the snake produce genuine normative pressure or force that the mistake about the normative relevance of agony does not.

Perhaps this is a matter of the distinction between *de re* and *de dicto* interpretations of apparent reasons.³⁵ On *de dicto* interpretations, an apparent reason is whatever we take to

³⁵ Although normative reasons are often thought to be objective, in the way I have described them, the sorts of reasons that affect the rationality of beliefs and actions appear to come apart from objective normative reasons, even though they are related to them in some way. The objective fact that there is poison and not gin in one's glass is an objective normative reason not to drink from it, but if one does not recognise that there is poison in the glass, it may still be *rational* to drink from it. This is where apparent reasons come in. It is rational to drink from the glass, even though there is no objective normative reason to do so, because there is an 'apparent' reason to drink from the glass. There is disagreement about what constitutes an apparent reason and the relationship between objective normative reasons and apparent reasons, notably the disagreement between *de re* and *de dicto* interpretations discussed here.

be an objective normative reason.³⁶ On *de re* interpretations, p is an apparent reason if (i) p would be an objective normative reason if true and (ii) we believe p.³⁷ Notice that, on the *de dicto* interpretation, you have an apparent reason to run away and the Tuesday indifferent have a reason to book the operation on the 15th, because in both cases you judge that you have a normative reason. However, on the *de re* interpretation, you have a reason to run away, because the snake being poisonous would be an objective normative reason were this the case, but future Tuesday indifference cannot make the fact that the 15th is a Tuesday an apparent reason to book an operation for the 15th, because this does not objectively count in favour of booking on the 15th.

If we endorse a view that says *de re* interpretation of apparent reasons matter, we can avoid the bad outcome. This approach concerns *facts*, facts about what reasons one would have if one's beliefs were true (Way, 2009: 4). The snake presents a reason to run away, given one's false beliefs about the behaviour of snakes, because of something about the objective danger of being poisoned and facts concerning the reasons this danger presents. Conversely, the risk of agony presents a reason not to be indifferent on a Tuesday due to something about the objective unpleasantness of agony, independently from what we believe about what reason agony provides us. Any view that acknowledges the importance of *de re* apparent reasons (rather than for example, only *de dicto* apparent reasons) has the resources to protect us from the bad outcome. On these views our moral reasons with respect to non-human beings cannot simply be produced or significantly limited by what we happen to believe about them, because our normative reasons must relate to objective moral significance. Because of this very fact, we would also have to maintain that objective normative reasons matter to normative reality.

This section uncovers a dynamic that explains why we end up either committed to the bad outcome or validating the importance of objective normative reasons; we need to recognise objective normative force that is more than just believing some consideration is reason-giving, in order to explain why the moral significance of non-human beings is not

³⁶ Kurt Sylvan (2015) rejects the claim that *de re* interpretation of apparent reasons relevant to rationality, as an alternative to *de dicto* approach to rationality, but neither accepts nor rejects the *de dicto* interpretation. Niko Kolodny (2005) takes the apparent reasons that are relevant to rationality to be the reasons that we think we have.

³⁷ Derek Parfit thinks of apparent reasons as beliefs whose truth would give us a reason (Parfit, 2011: 111). Jonathan Way (2009) argues that an important advantage of this view of the apparent reasons relevant to rationality is that it can accommodate the view that rational requirements are normative.

produced or significantly limited by what we believe about them. I do not claim to establish anything about which conception of reasons determines what we ought to do, or which sense of ought is most important. The claim I make is much more modest, which is just that we are missing something about normative reality if we claim that objective normative reasons don't matter. This is something that is not particularly contentious, as Section 2.2 might have already established, but this section has shown that if there is a defensible view that does not accept these claims about the importance of objective normative reasons, it will be deeply incompatible with certain commitments we might expect from a plausible environmental ethic.

Before I conclude this section, I address a remaining objection to the importance of objective normative reasons, an objection to do with their mysterious nature. We might think that a theory of reasons must answer the question of what makes something a reason. Even if facts that objectively speak in favour of certain responses are not strictly reasons, there is a sense in which they demand something from us, by counting in favour of a response. We might seek an explanation for why this is. This would not be an explanation for why we take ourselves to have reasons, e.g. explaining moral reasons in terms of how moral practice came about, but instead an explanation of what justifies or grounds their authority. In other words, whether or not we claim facts are normative or deserving of being called reasons, a version of Christine M. Korsgaard's 'normative question' arises, which would ask what justifies or grounds the authority of fact that objectively counts in favour (1996: 23).³⁸

Objectivism about reasons may have a particularly hard time answering this question. Subjectivism can say that something is a reason because of its relationship to an element of our perspective that grounds its authority, e.g. if reasons arise from what we accept about the world and our desires, then their authority is grounded in an imperative to do what we take to be the right thing. However, objectivism will not appeal to elements of our perspective to explain the authority of reasons or why we should care about them. I cannot respond to this concern directly, because an adequate response probably requires an in-depth discussion about the nature of objective normative reasons and what grounds them

³⁸ Some argue that reasons are fundamental in the sense that they do not derive their authority from any other source, in which case there is no further question about what grounds their authority (Scanlon, 2014: 68).

and I have left much of this unsaid. However, my arguments for the importance of objective normative reasons suggest that the problem should be reframed. In this chapter, I have argued that to deny the existence and importance of objective normative reasons incurs unacceptable consequences. This suggests that the mystery surrounding objective normative reasons prompts inquiry into their nature and what grounds their authority, but it is not by itself a reason to suspect they do not exist or are not important.

It is important to point out that my commitment to the importance of objective normative reasons, much like the commitments I discussed in Section 2.2, is relatively uncontroversial. There seems to be a place in our moral practice for different ways to say that someone acts wrongly, e.g. relative to the facts, their beliefs, or their evidence, and as Parfit points out ‘if we don’t draw these distinctions, or we use only some of these senses, we shall fail to recognise some important truths, and we and others may needlessly disagree’ (2011: 151). I conclude this section by saying that we can relativise reasons to an agent’s epistemic circumstances, but we must admit the importance of objective normative reasons, even if they do not determine what we ought to do. I have also highlighted why we might think the case for the importance of objective normative reasons is strongest in the context of environmental ethics.

2.4 Internalism and Externalism

If grasping is an epistemic relation to objective normative reasons, then whether we achieve grasp extends beyond our perspective. This section establishes that, despite its association with internalism, grasp can plausibly be conceived of as an externalist success. Nevertheless, we do not need to endorse particularly controversial externalist commitments to explain the nature of grasping and its value.

Some think that what matters from the epistemic perspective (and what should be of central concern to epistemology) is internal to our perspective and others think that what matters extends beyond our perspective. We can see this in debates about epistemic justification for belief. There are two prominent internalist theses about epistemic justification. The first says that the factors that determine whether a belief is justified must be accessible (or potentially accessible) to the believing agent upon reflection. This is

called *accessibilism*.³⁹ Accessibilism is motivated by cases like Laurence Bonjour's (1985: 41) *Clairvoyance*, in which Norman forms a belief based on a reliable process but has no accessible reason for thinking he has this power. Bonjour (1985: 42) claims that these beliefs are 'epistemically irrational and irresponsible' and, even if they are reliably true, they are nevertheless unjustified.

The second internalist thesis says that the factors that determine whether a belief is justified supervene on the mental states of the believing agent. This is called *mentalism* (Feldman and Conee, 2001; Feldman, 2005; Egeland, 2019). Mentalism is often motivated by cases in which two agents have identical mental states; agent 1 is in a world where their beliefs are true and agent 2 is tricked into having false beliefs by an evil demon (Lehrer and Cohen, 1983). The verdict about these cases is supposed to be that both agents are justified and therefore justification supervenes on one's mental states. As Zoë Johnson King (2022: 55) suggests, we should look beyond the supervenience claim itself and focus on the deeper intuitions that might drive us towards this verdict, such as the priority that the internalist gives to the agent and their efforts to understand the world and do the right thing, even when they are unfortunate and the external world prevents them from succeeding.

Externalism denies these internalist theses about epistemic justification. Grasp is, therefore, more in line with externalist considerations.⁴⁰ Objective normative reasons are facts and they extend beyond our perspective. Grasp, if it is to facilitate responsiveness to objective normative reasons, is unlikely to be reflectively accessible because it is not necessarily accessible to us whether we are in an epistemic position to respond to objective facts. Moreover, grasp does not supervene on our mental states because objective normative reasons are not part of our mental states and therefore whether we have achieved grasp is not settled by our mental states.

³⁹ For defence of accessibilism, see Bonjour (1985), Egeland (2020), Smithies (2012) and Steup (2003).

⁴⁰ It is worth noting that there are different ways to determine whether the externalist or the internalist is right about epistemic justification, which we might think of as the 'terms of engagement' (Goldman, 2009: 2). For example, the externalist could be right if there is *one* external factor that bears on justification of our beliefs, or we might think that whether the internalist or externalist is right depends on whether the majority of factors that bear on justification are internal or external.

It could be called into question whether grasp concerns what is beyond our perspective in this way, especially given that many align grasp with internalist ideas. Linda Zagzebski (2001: 246) argues that understanding is *transparent*, such that one cannot understand without understanding that one understands, Duncan Pritchard (2009: 39) claims when we understand we have some reflectively accessible grounds for our understanding, and Jonathan Kvanvig (2003: 202) argues that understanding requires an internal grasp of coherence relationships and this affords subjective justification. When we understand, we can “see” or “grasp” how pieces of information hang together and there is some initial plausibility to the idea that this has some affinity with internalist concerns.

Despite this tendency to align understanding with internalism, there are very good reasons to question this idea that grasping affords access to whether we understand. Understanding something in the world does not centre around concerns that are accessible from the first-person perspective. Stephen R. Grimm (2017) provides complaints of this kind, arguing that understanding is subject to accuracy and reliability conditions and whether we meet them is not accessible by reflection alone. We might make this same argument for grasping, if we think that the object of the grasp is something in the world, such as natural phenomena, rather than something internal to our minds. Grimm writes:

[O]ur understanding of natural phenomena seems conspicuously *factive* – what we are trying to grasp is how things actually stand in the world, and there is no reason to think that how things stand in the world is consciously transparent to us (2006: 518).

For a person to grasp the *world*, they need to have an accurate picture of the world, e.g. in *Animal Suffering*, we do not grasp the moral wrongs of unnecessary animal suffering if animal suffering is not wrong. We also cannot grasp the severity of pollution if our beliefs about pollution are just a lucky guess. We cannot tell from the first-person perspective whether our representations of the world are accurate or reliable and, therefore, we cannot tell by reflection whether we grasp.

These arguments suggest that whether we grasp is not accessible to us, but perhaps our grasp provides access to how pieces of information are related and, therefore, to what we understand. This is initially convincing when we associate grasping with a distinctive a

phenomenology of seeing how things fit together – sometimes called the ‘aha moment’.⁴¹ When we experience the aha moment, it is plausible that an idea of how things are related is accessible to us upon reflection. However, it is not obvious to me that once we arrive at an understanding via an aha moment, the relationships involved continue to be accessible by reflection or even potentially accessible. Consider Catherine Elgin’s example of a commuter who understands the New York Subway system; she knows which route is faster dependent on other factors, she knows where to stand to get a seat, and she has a feel for when it will be quicker to switch lines.⁴² It seems to me that the commuter grasps the subway system in virtue of this, whether or not such relationships are (or have ever been) accessible to her by reflection, e.g. she cannot draw a map or articulate the rules that she is following. This does not appear to be simply having less grasp, given that we can imagine a case where the commuter does have access these relationships and this interferes with her grip on the subway and the speed with which she can exercise understanding.

In terms of mentalism, some do describe grasping in a way that implies responsiveness to our own mental states. For example, one suggestion is that grasping is centrally to do with the relationships between beliefs (Kvanvig, 2003; Gardiner, 2012). Whether this is tantamount to saying that grasping supervenes on mental states is unclear, but this, for me, does not seem plausible for the cases introduced in the previous chapter. Consider *Anthropocene*, where the salient feature does seem to be about organising information in one’s mind. Even in this case, the natural thing to say is not that S arrives at a better grasp of their mental states, such as their beliefs, by adopting the concepts ‘anthropocene’ and ‘endling’, but that S arrives at a better grasp of the world *via these concepts*. This suggests that grasping success extends beyond S’s mental states. Although the systematising of information in one’s mind is important, this appears to be instrumental to the improvement of cognitive contact with the world.

If grasping is an externalist success, what further commitments may follow about its epistemic value? In my view, it’s quite clear that we should care about whether we get the world right and therefore whether we arrive at external epistemic success. Beyond the debate about what makes belief justified, the internalists and externalists have a different

⁴¹ For more on the ‘aha moment’, see Grimm (2010).

⁴² This is an example that Elgin gives of an understanding without explanation or grounding (2017: 42).

set of concerns and both appear important.⁴³ As we see from the above discussion about accessibilism and mentalism, the internalist might care about the agent's first person perspective, their efforts and perhaps whether we should think of the agent in a positive or negative light (e.g. verdicts about whether they are praiseworthy or blameworthy). The externalist, however, seems to care more about whether the agent has got it right and whether we can say there is something wrong with their epistemic states. For example, we might think that Norman has justified beliefs in cases like *Clairvoyance* or even cases where we hold such beliefs dogmatically, if such beliefs are true 'not as a matter of good luck, of happy accident, but as a function of his attuned capacity to get on to the truth' (Srinivasan, 2020: 398). Both sets of concerns have a place in epistemology. We should also remember that to deny that grasping has epistemic value, because it is an external epistemic success, would be controversial, as this could exclude knowledge from having epistemic value – if knowledge entails true belief, then knowing necessarily requires cognitive contact with the truth that extends beyond one's perspective.

Although suggesting that grasping (as an externalist success) is not itself particularly controversial, my approach to grasp carries more implicit externalist commitments in need of defence. To see this, note that some take interest in understanding due to their dissatisfaction with the epistemic value of knowledge, and aligning grasp with internalist theories of epistemic justification is part of the reason for thinking that understanding has more promise (Kvanvig, 2003: 202). By opting for an externalist approach to understanding, I reject this claim that we should value understanding more than knowledge because of an internalist notion of grasping and suggest instead that we should value understanding more than knowledge due to an *externalist* notion of grasping.

However, this does not amount to a strong externalist commitment, given that it is still open to the epistemologist who endorses my view on grasping to find some other internalist consideration, such as rational belief, to be of more importance than both knowledge and understanding. My claim is just that understanding (specifically grasping)

⁴³ William P. Alston (1993) points out that theorists do not tend to deny that internalist considerations (e.g. accessibility) and externalist considerations (e.g. reliability) are good from the epistemic perspective, but just that they are necessary conditions on epistemic justification. Alston suggests that there are features of the debate surrounding epistemic justification that suggest we should abandon the search for necessary conditions for what is properly considered an epistemically justified belief in favour of an approach that recognises numerous ways that a belief can be evaluated from the epistemic perspective – in other words, we should investigate different epistemic desiderata.

is not the source of this value. In fact, I think internalist considerations to do with evaluating the epistemic efforts of agents, despite their unfortunate circumstances, are important considerations, which I address in my discussion of responsibilist epistemic virtue in Chapters 5, 6 and 7. So, a view of grasping that centres responsiveness to objective normative reasons depends on a relatively uncontroversial claim that external epistemic success can have epistemic value and this does not rely on giving any priority to externalist considerations in epistemology.

2.5 Conclusion

In this chapter, I have defended some relatively uncontroversial commitments about reasons and epistemic success that follow from the claim that grasping facilitates responsiveness to objective normative reasons. I have argued that we must maintain that there are objective facts that count in favour of certain responses, that they are an important part of normative reality, and that grasping is plausibly an externalist epistemic success.

In the next chapter, I argue that we grasp to the degree to which we are in an epistemic position to respond to objective normative reasons, I suggest some conditions for what is epistemically required to respond to objective normative reasons, and I argue for the advantages of this approach.

Chapter 3: Grasping

3.0 Introduction

In Chapter 2, I defended some commitments that follow from thinking of grasping as cognitive contact with objective normative reasons for certain responses. However, I have yet to unpack the nature of grasping. The central aim in this chapter is to make the case that the degree to which one grasps corresponds to the degree to which one is in an epistemic position to respond to objective normative reasons. If this is right, we can see how grasp is relevant to doing right by the environment; grasp facilitates our responsiveness reasons and, as we shall see in this chapter and Chapter 4, this includes rendering vast amounts of complex information intelligible, by selectively emphasising our most important reasons.

Section 3.1 argues for three conditions on being in an epistemic position to respond to objective normative reasons. Section 3.2 argues that degree of grasp corresponds to the degree to which we are in an epistemic position to respond to objective normative reasons. This approach comes apart from some prominent ideas about grasping. Although I do not provide conclusive arguments against other accounts of grasping, Section 3.3 defends my view against some possible objections that arise from these conflicts between my account of grasp and others. Section 3.4 argues for some important advantages of my approach to grasping.

3.1 Being in an Epistemic Position to Respond to Normative Reasons

In Chapter 2, I argued that it is a good thing to be in an epistemic position to respond to objective normative reasons: facts that count in favour of certain responses. To be responsive to our objective normative reasons, we need to be in a certain epistemic position. If we are not in that epistemic position, our objective normative reasons will be to some extent out of our reach and unable to guide us. In *Day's End*, the fact of A's possible injury cannot guide B to avoid flipping the light switch. It would be a good thing if B

became responsive to this fact as the reason that it is and this would require some change in B's epistemic position. I take all of this to be relatively uncontroversial (as established by my arguments in Chapter 2).

What is controversial, however, is the question of what is needed to be in an epistemic position to objective normative reasons. To clarify, I am interested in the sort of epistemic position that consistently moves us to act, believe, feel or otherwise respond to objective normative reasons. That is, the sort of epistemic condition that, in some sense, brings these reasons to the forefront of our minds. Many have discussed epistemic conditions on what it takes to be in *possession* of a reason, but not all who are concerned with epistemic conditions on the possession of reasons are interested in the same epistemic position that I am.⁴⁴ Many philosophers are interested in reason possession because reasons must be possessed if they are to affect the rationality of what we do and believe or justify our actions. For example, Paul Silva writes:

in order for a fact to be a reason that justifies an action or attitude for a given agent it must be a fact that is *possessed* by that agent. There are a range of views of what it takes to possess a fact in a way that allows it to function as a reason for an agent (2021: 2925).

What it takes for a reason to function in our reasoning, such that it may justify, explain, or rationalise our responses, can fall short of the epistemic condition that interests me and that is at issue in Chapter 1.

Recall that Chapter 1 argued that agents, in cases like *Animal Suffering* and *Anthropocene*, can *improve* in their ability to do the right thing by improving their grasp. It seems to me that there can be a gap between being in an epistemic position with respect to a reason *r* that is sufficient for *r* to function in one's reasoning (such that, for example, ignoring *r* would be *irrational*) and being in an excellent epistemic position with respect to *r* that is *effective* in consistently enabling responsiveness to *r*. Firstly, this latter position seems to be a matter of degree; it's plausible that our epistemic position may facilitate responsiveness to *r* as a reason to a higher or lesser degree. For example, S in *Animal*

⁴⁴ As mentioned in Chapter 1, some think that the epistemic condition on reason possession is knowledge (Littlejohn, 2017; Williamson, 2002), others think it is being in a position to manifest relevant know how (Lord, 2018: 97-124), and others think it involves awareness (Silva, 2021).

Suffering improves their reason-responsiveness when they grasp by perceiving the video, but they might still have been irrational not to act on their existing knowledge of their reasons not to contribute to unnecessary animal suffering. To the extent that philosophers discussing epistemic conditions on reason possession are interested in defending claims about an epistemic position that consistently brings about responsiveness to normative reality, then I might be in debate with them in what follows, but if their claims concern what at minimum is required for a fact to function in one's reasoning or bear on the rational status of beliefs and actions, then I am not necessarily in disagreement.

I will argue for three conditions on whether we are in an epistemic position to respond to the normative reasons presented by some portion of the world X: we need (i) a representation of X that preserves reason-giving force of features of X, (ii) an ability to treat these features as the reasons they are and (iii) a non-accidental connection to our reasons.

3.1.1 Condition (i): Mental Representation

Condition (i) requires that we have a mental representation of X that preserves reason-giving features of X. To make any sense of what it means to be in an epistemic position to respond to reasons, we need some way to demarcate some part of the world and thereby the reasons it presents. For example, we can be in a position to respond to normative reasons, if any, that a given proposition presents. We can also be in an epistemic position to respond to normative reasons that are presented by a wider domain, such as a subject, event, a series of events, state of affairs, a person, an action, or an explanation. So, being in an epistemic position to respond to normative reasons presented by some part of the world X will involve a mental representation of X.

I will not provide an account of mental representation, but some examples of mental representations of X include the content of thoughts, beliefs, and acceptances that pertain to X. This mental representation need not be conscious or even potentially accessible to the conscious mind. However, it is important that the relevant mental representations can be evaluated according to whether they offer an accurate representation of X. This is understood in contrast to mental representations, such as the content of desires and

imaginings, that may reference a part of the world, but are not evaluated in terms of whether they fit the world.

Often, how well a mental representation provides cognitive contact with reasons is a matter of how accurate it is, e.g. true belief in facts that present normative reasons. However, relevant mental representations may tolerate some inaccuracies, as long as the reason-giving properties of X are preserved and sufficiently emphasised. For example, a representation might include idealisation or an oversimplification, but still share features with the world that provide reasons for inference to conclusions or reasons for action. As I will go on to argue in Chapter 4, deviating from the truth can sometimes improve our reason-responsiveness, by selectively emphasising reasons and downplaying or falsifying irrelevant features of the world.⁴⁵ Whilst some degree of accuracy is required to preserve reason-giving features, our ability to respond to normative reasons does not necessarily improve with accuracy.

At this point, we can see that mental representations improve cognitive contact with reasons in complex ways. For example, we may need more than just belief in a fact to appreciate its reason status. Sherlock has solved the case and tells Watson that the footprints are made by Red Wing Iron Rangers, but Watson is still confused (Lord, 2018: 97). Sherlock sighs and tells Watson that the cabby wears Red Wing Iron Rangers, revealing to Watson the status of this fact as a reason to think that the cabby did it. In order for Watson to see that this fact as a reason to infer that the cabby is guilty, Watson needs to know that the cabby wore Red Wing Iron Rangers. Whether we are in an epistemic position to respond to normative reasons presented by a given proposition will often depend on what we believe about other relevant things. Moreover, reasons tend to interact with one another and so, to preserve the force of a given reason, other reasons might need to be included in one's representation as well. For example, the fact that I am thirsty counts in favour of taking a sip from my glass, but the fact that there is poison in the glass far outweighs this reason. This means believing important reasons, such as the fact there is poison in the glass, will improve our epistemic position to respond to reasons presented by X more than other less important reasons. In other words, we could not simply estimate

⁴⁵ This will be based on Elgin's (2017) account of how a false representation can provide epistemic access to the world.

one's position to respond to reasons presented by X by counting the number of reasons included in one's mental representation of X. We must consider inclusion of other relevant facts that help us appreciate reasons (like in the Sherlock case) and reason interactions (like in the poison case).

3.1.2 Condition (ii): Ability

Condition (ii) is required to avoid over-intellectualisation and regress problems. Consider Errol Lord's example of taking a fact as a normative reason. We can infer that Julia is not at the Rose and Crown if she is at the Royal Oak (2018: 117-118). We only need to be sensitive to these facts, and their logical relationship, and we can make inferences without background beliefs about logic or beliefs about what reasons we have. Young children and animals appear to be able to respond to various reasons, without having beliefs about logic and without even possessing the concept of a reason (Parfit, 2011: 118).

We should also be wary of regress. We could believe the proposition that Julia is at the Royal Oak, a proposition that states this relationship (if Julia is at the Royal Oak she is not at the Rose and Crown) and the consequence that Julia is not at the Rose and Crown, and still fail to recognise the first proposition and the second proposition stating the relationship as a reason to believe the conclusion. We might believe that Julia is not at the Rose and Crown for an unrelated reason, e.g. because Julia does not like the Rose and Crown. We do not solve this issue by believing a proposition that reiterates the relationships between all these propositions because we can still believe that Julia is not at the Rose and Crown because she doesn't like it there, and so on to regress. To avoid regress, we need some kind of disposition, competence or ability to treat a fact as a reason.

To better understand what is going on here, let's investigate why this regress arises. Notice that reasons are always embedded in relationships. The fact that there is poison in the glass is a reason for me not to drink from it because this fact stands in a relation to me as an agent (who has an interest in not being poisoned) and the action of taking a sip. The fact that Julia is at the Royal Oak is a reason to believe she is not at the Rose and Crown because this fact stands in a relation to me as a believer and the inferred belief that she is not at the Rose and Crown. To recognise my reasons, I must be sensitive not only to this count-in-favour relationship, but also the way in which this fact counts in favour. This

implicates many other relevant relationships. In the poison case, this might include causal relationships to do with being poisoned. In the inference case, this might include logical relationships between the facts I believe. To be in an epistemic position to respond to normative reasons requires sensitivity to a non-propositional structure of relationships from which reasons arise.

The regress arises because appreciating reasons requires an appreciation of a non-propositional structure and belief is aimed at propositions. Consider Lewis Carroll's (1895) story of Achilles and the Tortoise. In this story, Achilles has premises written in his notebook that logically entail a conclusion. The tortoise challenges Achilles to force him to accept the conclusion, but the story implies he cannot do this by believing propositions that state the relationships between the premises in the argument. The tortoise can always refuse to believe the conclusion. The fact that the tortoise can believe the premises, but not the logical conclusion, shows us that what it takes to see the relationship between these premises, such that we take premises as a reason to believe the logical conclusion, is not exhausted by belief in the premises. Barry Stroud writes that this story does not provide a lesson about logic, but a lesson about belief: 'a list of everything a person believes, accepts, or acknowledges must leave it indeterminate whether any of those beliefs are based on others' (1979: 187-188). It is part of what it is to appreciate a reason that one appreciates the non-propositional structure of relationships in which it is embedded and it is part of what it is to believe that it cannot account for this, because belief is an attitude towards a proposition. To really see the connection between things, such as what Julia being at the Royal Oak has to do with her not being at the Rose and Crown, we need to have an ability that extends beyond believing more propositions (Gardiner, 2012).

At this point, we can see how being in an epistemic position to respond to normative reasons is not just a matter of knowing (although knowledge will of course often be part of the story). First, knowledge requires *true belief*, whereas knowing the truth is not required to respond to reasons. Second, the object of knowledge is a true proposition, but in order to be able to respond to objective normative reasons, we must be sensitive to relationships, and this involves some ability to do with appreciating non-propositional structure. That said, I now discuss an important feature of knowledge that is also required for responsiveness objective normative reasons.

3.1.3 Condition (iii): Non-Accidental Connection

There is more to grasping than having a mental representation of X that shares features of X and an ability to respond their status as reasons. If grasping is about being in an epistemic position to respond to our normative reasons, and these normative reasons are facts, what is required is a non-accidental connection to the facts themselves and not just a mental representation that happens to cohere with them. John Hyman (2010: 410-414) argues that truly believing is not enough to be guided by a fact, we need knowledge for the fact itself to be what guides us. Otherwise, it is only *as if* we respond to the facts.

To make this more concrete, consider Gettier cases, in which there is a justified true belief that does not amount to knowledge because its truth is in some way accidental.⁴⁶ Say that I look at a broken clock, that I am justified in thinking is functional, at a time of day when it happens to read correctly as 2pm.⁴⁷ In this case, I believe truly it is 2pm and I am justified in believing this, but I do not know it is 2pm. We might think I cannot be guided by the fact it is 2pm in this case, because I do not know this fact; it is just an accident that I am guided by belief content that happens to correspond to the facts. We need a non-accidental connection to the facts for them to be capable of guiding us, so whatever accounts for the lack of knowledge in Gettier cases also accounts for the shortcomings in our epistemic position to respond to normative reasons in these cases. What precisely is required for this condition to be met is of course the subject of much disagreement in epistemology and settling this question is beyond the scope of this thesis.

Some distinguish between Gettier luck and environmental luck, which may make a difference here. In the famous barn façade cases, I am in a region with many fake barns, and I could have easily been looking at a barn façade, but fortunately I form the true belief

⁴⁶ Gettier cases arise when a justified true belief is true by chance (Gettier, 1963). In these cases, the belief is both true and justified, but not an instance of knowledge. We might consider at this point if there is a justification condition on responsiveness to objective normative reasons. Firstly, it's unclear what is meant by a justification condition, given that epistemic justification is contentious to the point where it is questionable that theorists are talking about the same thing. I therefore do not rule out the possibility that there is a conception of epistemic justification that turns out to be required for responsiveness to facts as reasons. However, meeting conditions specified by prominent internalist theses, such as reflective accessibility, do not seem to be required for responsiveness to facts. This is not to say that reflective accessibility is not *ever* required to respond to certain reasons, but that it is not always present.

⁴⁷ The stopped clock case has the paradigmatic features of a Gettier case and is from Bertrand Russell's *Human Knowledge: Its Scope and Limits* (1948).

that ‘there is a barn’ after seeing one of a few genuine barns.⁴⁸ My belief is true and justified, but it is a matter of luck that my belief is true. For what it is worth, I am inclined to say that I would be in an epistemic position to respond to the fact that there is a barn as a reason. It seems this fact can guide me, given that I am fortunate to be looking at a real barn. What is so problematic about Gettier luck is that it interferes with the connection between the fact and the belief, such that you cannot respond to the fact itself. The same is not true for environmental luck, because although your belief may easily have been false, and might be ‘lucky’ in this sense, there is no comparable interference between belief and fact. If all this is right, then either we can know in case of environmental luck (and intuitions to the contrary are misguided), or knowing and being in an epistemic position to respond to reasons come apart in cases of environmental luck; this latter view would say that while knowledge is not present in either Gettier luck or environmental luck cases, we can still be in an epistemic position to respond to objective normative reasons.

We need to meet these conditions (i), (ii) and (iii) to be in an epistemic position to respond to objective normative reasons. This is plausibly a gradable success given that many reasons might be presented by X and the extent to which our epistemic position can facilitate our responsiveness to them comes in degrees. I will say more on this and how it relates to grasp and understanding in what follows. As I have stated above, I am reluctant to describe the epistemic position to respond to reasons in terms of knowledge, because reason-responsiveness does not require true belief and it requires ability that is not necessarily present for propositional knowledge. However, given that I maintain that a non-accidental connection to the facts is required for responsiveness to facts, at a certain point, this just looks like a matter of terminology. Whether one is willing to claim that this epistemic success is a kind of knowledge depends on whether one associates knowledge more strongly with true belief or non-accidental connections to facts.

In Section 3.2, I argue that the degree to which we are in an epistemic position to respond to objective normative reasons corresponds to degree of grasp.

⁴⁸ This example is originally by Carl Ginet and first published in Alvin Goldman’s article ‘Discrimination and Perceptual Knowledge’ (1976: 772-773).

3.2 Grasp Improves with Reason-Responsiveness

As briefly discussed in the thesis Introduction and Chapter 1, some epistemologists claim that understanding is distinguished from knowledge in virtue of a ‘grasping’ element that is supposed to go beyond knowing. Consider the following case:

Faulty Wiring: A parent and their young child witness a house fire. They both come to know, from a fire officer’s testimony, that the house burned down because of faulty wiring. The parent understands why the house burned down, whereas the young child does not; the child lacks a conception of how the faulty wiring brings about the fire that is required for understanding why the house burned down.⁴⁹

The idea is that the parent has something that the child does not; the parent can see how the faulty wiring causes the fire because they can grasp how elements of this explanation fit together, whereas the child is lacking this grasp. Some argue that understanding can be explained in terms of knowledge (Kelp, 2015; Sliwa, 2015; Khalifa, 2017). This thesis does not aim to directly tackle these objections, but I do demonstrate how my account could answer to them in Section 3.4.1. Although it is controversial that understanding comes apart from knowledge in virtue of a grasping element, among those who do accept this claim, there is much interest in the nature and value of grasp. Michael Strevens goes as far as to suggest that the question of the nature of grasping (or apprehension) at issue in understanding ‘is perhaps the deepest in all philosophy’ (Strevens, 2013: 511).

This section defends the following claim about the nature of grasping: we improve our grasp of some part of the world, when we improve our epistemic position to respond to normative reasons. I argue that grasp distinguishes understanding from knowledge because while knowledge is to do with affirming what is the case and proper connection to the truth, understanding is to do with being able to respond well and proper connection to our normative reasons. Before I begin this section, I should acknowledge that many varieties of understanding are distinguished in the understanding literature, including objectual understanding (Elgin, 2017; Kvanvig, 2003), understanding why (Grimm, 2014; Hills, 2016), and understanding how (Bengson, 2017). While we may associate grasp with all these varieties of understanding, accounts of grasp are often developed within the context

⁴⁹ The example of understanding why one’s house burned down is originally from Brogaard (2005). It is developed by Pritchard (2008) and discussed in Elgin (2017: 40-41), Grimm (2014), and Sliwa (2015).

of a specific variety. Just as I have argued that we have no initial reason to think that the sense of grasp in my cases in Chapter 1 are the same sense of grasp at issue in the understanding literature, we have no reason at this stage to think that the same sense of grasp is at issue in every variety of understanding. In what follows, I make the case that, even if there are different senses of grasp that correspond to these different varieties of understanding, responsiveness to reasons is what ties them all together. I therefore formulate my thesis about grasping in general terms, rather than focusing on a particular variety of understanding.

Given that grasp is a gradable epistemic success, in the sense that we can grasp to a higher or lower degree, I suggest the following thesis about degrees of grasp.⁵⁰

Gradable Grasping (GG): the degree to which S grasps X corresponds to the degree to which S is in an epistemic position to respond to the normative reasons presented by X.⁵¹

I now expand on GG, arguing for its initial appeal.

3.2.1 Why Normative Reasons?

An obvious question asks what degrees of grasping have to do with normative reasons. I argue that the initial appeal of GG is that this thesis makes sense of a few compelling ideas about grasp: that grasp is about using information, that the object of grasp is non-propositional and that grasp involves making sense of the world.

First, there are a few intuitive ideas that lead to the thesis that grasping is to do with responsiveness to objective normative reasons, as demonstrated by the following argument:

⁵⁰ It also appears coherent to talk about outright grasping attributions, i.e. S grasps X. However, it's unclear what determines these outright attributions and it's likely that whether an outright grasping attribution is appropriate depends on the context. On my view, this is to do with whatever degree of grasp is good enough in a particular context to facilitate response to some important set of reasons. Also, it's likely that to properly understand outright grasping attributions, we need an account of degrees of grasp. For these reasons, I have chosen to focus on grasp as a gradable success.

⁵¹ This is typically the form that objectual understanding takes: 'S understands X', such that X is comprehensive subjective matter or domain (e.g. S understands Epistemology). However, this thesis is intended to be adapted for other varieties of understanding, e.g. S understands why X.

- (1) The degree to which we grasp corresponds to the degree to which we are in an epistemic position to use information well.
- (2) The degree to which we are in an epistemic position to use information well corresponds to the degree to which we are in an epistemic position to respond to objective normative reasons.
- (C) The degree to which we grasp corresponds to the degree to which we are in an epistemic position to respond to objective normative reasons.

I will say something in support of premise (1) and (2), before explaining how degrees of grasp may relate to degree of responsiveness to objective normative reasons. I am using the term ‘information’ because this commonly used term better captures the common-sense nature of this argument. However, in the context of what I have said above, information just means the content of our understanding, which will include the propositional and non-propositional elements of one’s mental representation.⁵²

In support of premise (1), grasping is often thought to be to do with using information well.⁵³ Daniel Wilkenfeld argues that ‘understanding x involves possessing a representation of x that could be manipulated in useful ways’ (2013: 997). Catherine Elgin sees grasping as a kind of know-how; an ability ‘to wield commitments that bear on the topic – how to draw the inferences and perform the actions that the understanding licenses’ (2017: 56). Alison Hills says that when we grasp the relationship between propositions ‘you have the relationship under your control. You can manipulate it. You have a set of abilities or know-how relevant to it, which you can exercise if you choose’ (2016: 663). Notice that this is perhaps what is at issue in my cases *Animal Suffering*, *Anthropocene* and *Plastic Pollution*, where the sense of grasp relates to the agent’s ability to use information to do the right thing for the environment. Grasp does not entail that we do *in fact* use information in certain ways (we could, for example, become distracted by another task). Grasp is a measure of how much command we have over information, such that we could use it, and

⁵² Elgin (2017: 46) appears to use the term information in a similar way to refer to the truths, the relationship between them and other elements of an account that figures in a person’s understanding. Kvanvig also suggest that understanding requires the grasp of a ‘large and comprehensive body of information’ (2003: 192).

⁵³ Michael Strevens, however, argues that grasp is all detection and not any further ability to use information (Strevens, 2024: 753).

this seems to be another way of saying that grasping involves improving our *epistemic* position to use information well.

Premise (2) says that the degree to which we are in an epistemic position to use information well improves with the degree to which we are in an epistemic position to be responsive to objective normative reasons. In support of premise (2), there is a compelling interpretation of what it means to be in a position to use information well that corresponds to what it means to be in an epistemic position to respond to objective normative reasons. This interpretation is that to be able to use information well is to be able to mobilise information in accordance with reality, such that one can respond to the world in the right way or in a way that is favourable in a given context. If this is right, then this is to respond to objective normative reasons. In *Day's End*, B may think they are using information well and they might be using available information well, but there is an important sense in which they would use information better if they were in a position to respond to the fact that A may be injured if the light is switched on. This interpretation allows you to say intuitive things, such as B may be advised to respond differently given that they do not know about A's injury and if B gets closer to using information in accordance with reality, they are improving in their ability to use information. If we accept premise (1) and (2), we can conclude that grasp improves with being in an epistemic position to respond to objective normative reasons.

Second, many think that, in contrast with knowledge, the object of understanding is non-propositional (Grimm, 2014: 339; Kvanvig, 2003: 192; Zagzebski, 2001: 242). Georgi Gardiner (2012) argues that grasping has to do with seeing relationships between beliefs and, in the interest in avoiding the regress I discuss in Section 3.1.2, grasp cannot be a matter of believing propositions. GG accommodates these ideas. According to GG, grasping is a feature of understanding and reasons are necessarily embedded in relationships. We can know propositions in isolation, but we cannot appreciate the reason status of a proposition without appreciating relevant relationships from which a reason arises. This aligns my view with the prevalent idea that, while the object of propositional knowledge is, of course, propositions, there is a sense in which the object of understanding is non-propositional (even if propositions are also important).

Third, we might expect understanding to provide a deeper insight into the world than knowledge, in that to understand is to make sense of the world or render the world intelligible. GG can accommodate this idea as well. Becoming closer to our reasons is a way of making sense of the world because it is to recognise what the world *means for us*, i.e. what response is appropriate. To grasp, then, is to render the world intelligible because it involves mental representations of the world that selectively emphasise one's reasons and bring them to the forefront.

It seems that GG can explain some intuitive ideas about understanding. Although, it remains unclear exactly how degrees of grasp relate to degrees of understanding. It seems plausible to me that when grasp improves, understanding improves. However, it might be that understanding can improve without improvement in grasp since other measures of success might also be relevant to understanding, e.g. proximity to truth and comprehensiveness. I take grasp to be something like the 'core' of understanding because it is at least one source of its distinctive and superior value, as I argue in Section 3.4.

3.3 Other Accounts of Grasping

Although the previous section looked at the initial appeal of GG, it is important to note that GG may depart from some prominent ideas about understanding, including that understanding is not inherently normative, understanding is essentially related to reflective accessibility, that understanding is immune to epistemic luck, and that grasping can be made sense of by characterising a specific cognitive ability or way of representing the world. My aim in this section is to defend my departure from these ideas, as well as to highlight how GG can accommodate features of other accounts.⁵⁴ I start with some objections that arise from general ideas about understanding. Then, I look at specific claims about grasping developed in the understanding literature.

⁵⁴ Often, accounts of grasping are developed the context of a particular variety of understanding and, although we may be hopeful that an account speaks to the nature of grasping in general, I am mindful in discussion that this is not always the aim.

3.3.1 Is Understanding Always Normative?

Understanding is not typically thought to be inherently normative, since it seems that we can often understand things that are not obviously normative, such as the arrangement of raindrops on my window. There would be a serious objection to GG if we could grasp things that do not present any objective normative reasons; if we can grasp something to a high degree, even though it presents no normative reasons, it is unclear that degree of grasp corresponds to degree of one's epistemic position to respond to normative reasons.

However, an understanding of the arrangement of raindrops includes many non-obvious normative reasons, by which I mean facts that are not paradigmatic reasons and are therefore difficult to spot. Most notably, the arrangement of raindrops presents reasons to feel bored or to direct one's attention elsewhere. We can grasp such matters to a high degree, but this just means arriving at a very good epistemic position to appreciate that what we grasp is unworthy of our attention. This strikes a balance between the ensuring that it is possible to grasp not obviously normative things to a high degree, but at the same time acknowledging this would be poor quality grasp in some sense because the reasons in question are not important.

3.3.2 Varieties of Understanding

Other potential objections are directed at my claim that GG can be adopted with respect to grasp in different varieties of understanding. GG seems to work quite well for objectual understanding; grasp improves with cognitive contact with normative reasons presented by an object. GG also seems to work quite well for practical understanding; grasp in practical understanding concerns responsiveness to practical reasons about how to do something. However, it's least clear how GG applies to explanatory understanding.

The claim that grasp in explanatory understanding concerns responsiveness to objective normative reasons might seem questionable. For example, perhaps grasp in explanatory understanding corresponds to an appreciation of relevant explanatory relationships, such that the more we appreciate the relationships, the more we grasp. However, I shall now argue that, upon examination, grasp in explanatory understanding also concerns reason-

responsiveness because underlying normative reasons determine the explanatory relationships that matter to understanding.

To see this, recall that grasp is supposed to be about using information and GG makes sense of this in terms of reason-responsiveness; to improve in one's ability to use information, in a given context, is to improve in the ability to respond to the reasons it presents in that context. Returning to explanatory understanding, such as the understanding of why the house burned down in *Faulty Wiring*, consider also James Woodward's comment that it is no mystery why we should be interested in uncovering causal relationships, 'if these are relationships that are potentially exploitable for manipulation' (2005: 7). If we understand how one causal factor depends on another, then we are in a position to use this information to intervene and manipulate outcomes. Often, the causal relationships that figure in an explanation that produces understanding are those that we *should* manipulate relative to some practical aim, e.g. in *Faulty Wiring*, the causal relationships that matter might be relative to the aim of preventing faulty wiring and thus avoiding future fires. In other words, the causal relationships that figure in explanatory understanding uncover practical normative reasons.

In another interesting set of cases, the reasons underlying an explanation are purely epistemic, such as facts that count in favour of further beliefs or feeling that one's curiosity has been satisfied. The explanation that matters to the understanding of an agent that is simply curious about why some event occurred is the explanation that reveals the fact that counts in favour of taking their curiosity to be satisfied. In general, the explanatory relationships that figure in understanding reveal something about the world that matters to us.

So, there seem to be normative reasons underlying which relationships feature in explanatory understanding, but it is yet to be established that degree of grasp *corresponds* to the degree to which we are in an epistemic position to respond to these underlying reasons. We can see this, however, from the context-sensitivity of explanations. Consider a case where a server hits a table and spills coffee (Grimm, 2014: 330). Many factors contribute to an explanation of why the coffee spilled, e.g. that the server hit the table, that the server did not sleep well, that they applied for the job 5 years ago, and that they turned up to work this morning. What we take to be the relevant cause depends on context. For the

server, perhaps poor sleep is the relevant cause of the coffee spill because this is what matters to prevention of future mistakes. This fact is a salient reason for action. For a neutral witness of the spill, the relevant cause might be that the server hit the table. This fact satisfies their curiosity about why the spill occurred. Whether grasp improves depends on what reasons we have and how well the explanation we accept brings them to the forefront. As I shall go on to argue, my view is that many accounts of grasp, that are developed within the context of explanatory understanding, provide important insights into the nature of grasp because they aim at something that contributes to the command of explanatory relationships (and therefore responsiveness to the underlying normative reasons).

Finally, another objection concerns the distinction between grasp of a fact (which we might think of as ‘propositional understanding’ or understanding that) and grasp of a body of information, which may include practical, explanatory and objectual understanding. GG says that grasping in objectual understanding increases with being in an epistemic position to respond to normative reasons presented by its object. As above, grasp in explanatory understanding improves with being in an epistemic position to respond to (potentially non-obvious) normative reasons underlying the explanation, such as reasons to make inferences or take action to manipulate outcomes. However, GG seems to say that grasp of a fact concerns being in an epistemic position to respond to the reason that fact presents in a certain context, e.g. for a particular person in a particular circumstance. This means that propositional understanding always derives from understanding the relevant body of information from which reason status arises. We therefore risk eroding what seems like an intuitive distinction between grasp of a fact and grasp of a wider body of information.

Firstly, it’s not clear to me that the distinction is entirely eroded just because propositional understanding collapses into something more like objectual understanding of the relevant relationships from which a reason arises. There might be contexts in which focusing on understanding of a proposition is philosophically interesting because it serves a particular function in our epistemic lives. For example, we might cite grasp of a fact in explaining why someone acted a certain way (e.g. citing their recognition of a particular reason), justifying our actions, or explaining an argument. This might be enough to make the distinction between grasping a fact and grasping a body of information worthwhile. However, if we cannot maintain this distinction, this is not a tragedy for GG, especially

since the sense in which it does erode the distinction is not problematic. Where grasping a fact is important, it does seem to be to do with moving towards an understanding of the wider significance of that fact, i.e. locating it within a wider state of affairs, and therefore it makes sense for grasping a fact to be derivative.⁵⁵

3.3.3 Accessibility and Reflection

As explained in Chapter 2, Section 2.4, understanding is often thought to be an internalist notion and is, in particular, associated with the internalist thesis (*accessibilism*) that says that factors that determine justification for beliefs must be in some way accessible to the believing agent upon reflection (Zagzebski, 2001: 246; Pritchard, 2009: 39; Kvanvig, 2003: 202). When we understand, we can “grasp” how pieces of information hang together, e.g. in *Faulty Wiring* the parent can see how the causal factors interact, and some think that this affords access to whether we understand or what we understand. In contrast, GG says that grasp corresponds to being in an epistemic position to respond to objective normative reasons and this does not plausibly require that our reasons are accessible to us upon reflection, provided that we can respond to our reasons due to implicit beliefs and unconscious grasp that is not even potentially accessible. Moreover, I do not think that the structure of relevant relations, in which our reasons are embedded, needs to be accessible either. If this is right, grasping success bears no essential relationship to what is accessible upon reflection and, on this point, departs from accessibilist accounts of grasping.

We have already seen in Chapter 2, Section 2.4 that there are good reasons to question that grasping affords access to whether we understand. For a person to understand the world, they need to have an accurate picture of the world, e.g. in *Faulty Wiring*, we do not understand the cause of the house fire if we falsely believe it was caused by an unattended flame. Similarly, we cannot understand the cause of the house fire, if our belief that it was caused by faulty wiring is a lucky guess. If you agree that what one grasps is something in the world (and grasp is *about* the world in this sense), then you might be inclined to say that whether one grasps is not accessible upon reflection.

⁵⁵ Emma C. Gordon (2012) argues that propositional understanding cannot be distinguished from other varieties of understanding.

In Section 2.4, I also criticised the idea that grasp affords access to how pieces of information are related and, therefore, provides reflective access to what we understand. While this might be true of conscious grasp, we can see from the commuter case not all grasp is conscious. The commuter grasps the subway system because they have a feel for where to go and where to switch lines even if they do not have reflective access to the relevant relationships. I also suggested that achieving reflective access might be counterproductive, by interfering with the speed with which the commuter can exercise their understanding. Whether you think that *what* one grasps is accessible upon reflection may depend on what you think about the commuter case and this idea that accessibility can sometimes get in the way.

Despite the fact that I reject accessibilist views, GG can still explain their appeal. The phenomenology of knowledge, often thought to be belief, involves an appreciation for the way things are and this might say something about the nature of knowledge: that it connects us in some way with the truth. In the same vein, perhaps the phenomenology of understanding, plausibly the ‘aha moment’, reveals something about grasping. In order to grasp, we must be sensitive to relevant relationships from which our reasons arise and, therefore, we need to organise and systemise how pieces of information are related to each other accordingly. It makes sense that ‘what it is like’ to grasp will involve having access upon reflection to how things fit together and cohere with each other.

3.3.4 Epistemic Luck

Some think that grasping is, in some sense, internal to our mental lives and therefore immune to epistemic luck in cases like the following:

Faulty Wiring Gettier Luck: You approach someone who looks like a fire officer and ask what caused the fire. You are told that the cause is faulty wiring. Although it is true that the fire was caused by faulty wiring, the person you asked is not a fire officer, but someone on their way to a fancy-dress party.⁵⁶

⁵⁶ This specific case is contributed by Duncan Pritchard (2010: 78). Note that Pritchard does not think understanding is present.

In this case, many attribute understanding, even though your beliefs about the fire are only accidentally true.

There are different possible explanations for why we might attribute understanding in the presence of luck and many refer to grasping. Kvanvig argues that, once we meet truth requirements on understanding, its distinctive features are ‘internal to cognition. It is the seeing or appreciating of explanatory and other coherence-inducing relationships’, whereas when we think of knowledge our focus turns to ‘the possibility of fortuitousness, of accidentality, of being right but only by chance’ (2003: 198). Alison Hills (2016: 663) argues that if you understand why p, and q is why p, then you grasp the relationship between these propositions, such that you have a set of abilities relevant to this relationship – this set of abilities, called *cognitive control*, is discussed in more detail in Section 3.3.6. On Hills’ view, the tie between beliefs required to manipulate the relationship between propositions and the tie at issue in knowledge, which concerns the formation of true belief without luck, are separate properties and understanding does not have both ties (2016: 674). Strevens suggests that for a true belief to be knowledge ‘in it must have a certain degree and kind of security, whereas in order for true belief to attain the status of grasp it must have a certain degree and kind of *perspicacity*’ (2024: 750). There is intuitive pull to the idea that successful grasp concerns one’s ability to see or manipulate relationships between the content of our beliefs and how we arrived at these beliefs is not important.

Given the intuitive pull to the idea that understanding is immune to epistemic luck in cases like *Faulty Wiring Gettier Luck* and the notion of grasp as being to do with seeing or manipulating the relationships between our beliefs, we might worry that GG delivers the wrong result. GG says that grasp is not immune to epistemic luck, because it is a relation to reasons, and that grasp is not present in *Faulty Wiring Gettier Luck*. I have argued that, in order to respond to our objective normative reasons, we need to be non-accidentally connected to them, such that we can respond to the facts themselves and not just mental representations that happen to cohere with them. I do not take myself to have a conclusive argument against other accounts of grasp or the verdict that grasp is present in *Faulty Wiring Gettier Luck*. Instead, my aim here is to defend GG against this case as a counter example to my view, by offering an alternative interpretation of what is going on here.

Consider an alternative explanation for why we might attribute grasp in *Faulty Wiring Gettier Luck*. There is genuine grasp in this case, but instead of grasping the cause of the fire, what you grasp is either the content of the party-goer's testimony (which would involve non-accidental connection to the testimony) or the content of your beliefs about the fire (which would also involve a non-accidental connection). If this is right, then this would be a special case of a more general phenomenon. In the special case, what is grasped are features of a *representation* of the world (the content of testimony or your mental representation), rather than the world itself. But presumably we can also know things about representations. Why are we willing to attribute grasp, but not knowledge? Perhaps the reason for this is that we are drawn to the distinctive abilities associated with grasping. In this case, there is an 'aha moment' where these abilities are realised, and this might explain why we notice grasp of representations more than knowledge of representations.

There are some reasons to prefer this alternative explanation. First, note that there are competing intuitions *against* the verdict that grasp is present in *Faulty Wiring Gettier Luck* and these intuitions support this alternative explanation. There is something strange about saying that you grasp the cause of the fire in this case, given that your representation of the fire is based on the partygoer's testimony, and this has nothing to do with the fire. The partygoer simply happens to be right by chance. We might suspect the claim that grasp in this case is really *about* the fire. To see this, consider a different case, in which the partygoer says something false, and you wrongly believe that the fire was caused by an unattended flame as a result. In this case, you do not grasp the fire *itself* (you don't even believe anything true about the fire), but you might grasp some false representation of the fire according to which it occurred due to an unattended flame. The only difference in *Faulty Wiring Gettier Luck* is that the partygoer happens to be right by chance. I find it hard to see why we should conclude there is a grasp of a representation (rather than the fire) in the case of false belief and to conclude that there is grasp of the fire itself in *Faulty Wiring Gettier Luck*, just because the truth value of the partygoer's testimony has changed. This fact that the partygoer's testimony is true in the case where you are fortunate enough to be right is still nothing to do with the fire itself. In both cases, the fire itself has no impact on the representation of the world, and ability to see relationships that constitutes your grasp, and we could say that's why you do not grasp the fire itself in either of these cases.

One might object that I have gone wrong by conflating understanding with grasp. Perhaps understanding is not present in *Faulty Wiring Gettier Luck* (because the content of your understanding is not about the fire in the relevant sense), but grasp is present because ‘grasping’ refers to the cognitive abilities you exercise in understanding. This move would deliver the result that grasping is immune to epistemic luck, even if understanding is not.

One point against this move is that it deflates the importance of grasp. To appreciate this, consider what Grimm says about aligning understanding with an ‘internal’ conception of grasp that is at issue when we achieve a subjective sense of understanding, i.e. when we can figure out how aspects of a representations fit together irrespective of whether a representation is true:

if one allows that the sort of subjective, internal grasp described here is a cognitive success, then cognitive successes plausibly begin exploding all over the place, e.g. believing in accordance with the evidence and recognising defeaters are similar successes to do with figuring out how things seem from the inside (Grimm, 2012: 112).

To say ‘internal’ or ‘subjective’ grasp is all there is to understanding is a problem because we can no longer see what is so special about understanding; it turns out these sorts of cognitive successes are very commonplace. In the same vein, identifying *grasp* with these internal cognitive abilities makes it hard to see what is so special about *grasping*. If grasping is supposed to be a particularly important epistemic success, perhaps GG has more promise than views that identify grasping with cognitive abilities related to piecing things together in your mind.

Before I conclude this section, it is important to recognise that my account can preserve some features of accounts that take understanding to be subject to epistemic luck. GG can say there are cognitive abilities that are distinctive to grasping, and that are not present in knowledge, because these abilities may fulfil the ability condition on what it takes to be in a position to respond to objective normative reasons. This leads me to the next few subsections, in which I highlight some potential compatibilities with accounts that make sense of grasp by characterising particular abilities and ways of representing the world.

3.3.5 Strevens – Recognition

Michael Strevens (2024: 741) seeks the meaning of grasp as it operates in the claim ‘to understand why a phenomenon occurs is to grasp a correct explanation of it’, although we might be optimistic that this ability illuminates other senses of grasp.⁵⁷ Strevens’ recognitional account of grasping ‘equates a thinker’s grasp of a property with their ability to recognise its instances’ (2024: 751). Recognition is an ability to identify or ‘pick out’ entities that belong to a certain kind. The accuracy and scope of recognition corresponds to the degree of grasp; for example, you improve your grasp of “dog-hood” if you can more reliably distinguish between dogs and other things and you also improve your grasp if you can exercise this ability for a wider variety of dogs.

Some important points of agreement with between GG and Strevens’ account could be that recognitional ability can have a role in our appreciation of reasons. Consider Strevens’ comment that recognition is ‘simply inference to a conclusion of the form, “This is one of those”’ (2024: 751). We can find two kinds of cases where recognitional ability is a salient condition on being in an epistemic position to respond to reasons. The first is the case where some particularly weighty objective normative reason is itself a reason to make an inference to the conclusion “this is one of those”, for which recognition will clearly be the salient ability in our assessment of grasp.⁵⁸ Another kind of case is where, in order to respond to some particularly weighty objective normative reason, we need to recognise that some entity is of a particular kind. Consider *Animal Suffering*. Perhaps it is an important part of S’s grasp of why causing animal suffering unnecessarily is wrong is that S recognises animals as being sentient beings.⁵⁹ Moreover, insofar as recognition is important to appreciating explanatory relationships, then it will also contribute to one’s epistemic position to respond to the normative reasons that underly them. If my arguments

⁵⁷ Strevens appears to be open to this possibility, but does not attempt to argue for this claim (Strevens, 2024: 750).

⁵⁸ I wonder if this might be what happens in Strevens’ cases (2024: 742-743). The student who doesn’t think much about the abstract property of effective population, but memorises formulas and their application to answer questions in their exams, is missing grasp of relevant explanations and the student who can recognise the abstract property of effective population as a component of the relevant explanations does grasp them in virtue of this recognitional ability.

⁵⁹ Some might think recognition is fundamental to responsiveness to reasons given that, on some level, we must recognise *r* as the reason it is in order to respond to *r* as a reason. See Kurt Sylvan’s (2015: 600) argument for objective reasons-sensitive competences, according to which, for reasons to be apparent to us, we need to be disposed to treat considerations as objective reasons of a relevant kind only, if true, they would be objective reasons of that kind.

about explanatory understanding in Section 3.3.2 are right, then we can see how recognition is important to grasp according to GG.

An important difference between GG and the recognitional account is that grasp is not identified with recognitional ability if this ability is thought to correspond to accuracy and scope of recognition, i.e. whether we can reliably distinguish entities of a kind and for a wider variety of that kind. This is because grasp corresponds to the extent that we are in an epistemic position to respond to reasons, which is not always a matter of distinguishing entities in this way. GG accommodates the relevance of recognition by arguing that this ability constitutes grasp to the extent that it contributes to our responsiveness to reasons.

3.3.6 Hills – Cognitive Control

Alison Hills (2016: 663) associates the following abilities with grasping: (i) following some explanation of why p given by someone else, (ii) explaining why p in your own words, (iii) drawing the conclusion that p (or that probably p) from the information that q, (iv) drawing these conclusions about similar cases, (v) given the information that p, giving the right explanation q, and (vi) giving the right explanation for similar cases. Hills argues that this set of abilities, called cognitive control, is normally required for understanding why, but does not rule out that other abilities might also be included. Firstly, some of these abilities are directly related to treating something as a reason, such as drawing conclusions, which might be thought of as taking one's belief in a fact as a reason to infer a conclusion. These abilities may contribute to the ability condition on what it takes to respond to an objective normative reason.

Another point to make is that cognitive control is required to have command over explanatory relationships. So, if I am right that explanatory understanding is determined by underlying objective normative reasons, e.g. to manipulate outcomes, make inferences or feel that one's curiosity is satisfied, then we can see how achieving cognitive control is important to reason-responsiveness. To make this clear, consider also what Hills says about explanations. When you ask for an explanation, such as why is lying wrong, your question could request a number of things. Why is lying wrong when telling the truth is not? Why is lying wrong when not saying anything is acceptable? Why is lying worse than misleading? Hills writes:

what it is to explain why lying is wrong – and so what it is to understand why lying is wrong – may therefore be quite different in different contexts, depending on the background assumptions, expectations for the type of explanation and implicit contrast at issue (Hills, 2016: 664).

What counts as an explanation for why lying is wrong (and an understanding of why lying is wrong) can depend on context, notably what we expect from the explanation and what sort of inferences one hopes to make about the wrongness of lying. It's tempting to make sense of this in terms of objective normative reasons for beliefs about lying or actions that may follow.

Following Section 3.3.2, then, cognitive control is relevant to grasp on my view because having command over explanatory relationships is required for responsiveness to the reasons underlying explanations. Importantly, these abilities are not identified with grasp (according to GG), but they are important to how grasp is achieved.

3.3.7 Grimm – Modal Grasping

Grimm's (2014) modal account also focuses on explanatory understanding, arguing that grasping is about seeing or appreciating modal relationships, by drawing upon grasping/seeing metaphors in the literature on *a priori* knowledge. We move from knowing that $2 + 3 = 5$, e.g. by accepting results from a calculator, to grasping that $2 + 3 = 5$ when we not only assent to this fact but appreciate its necessity. Grimm makes a similar point for causal dependencies. He argues that one way to know cause is to assent to a proposition such as "the fire was caused by faulty wiring", but another sense of knowing cause requires more than this; we must appreciate how the faulty wiring brings about the fire, by seeing how certain changes may or may not result in other changes.

As above, if appreciating modal relationships is important to grasping explanations, then we can apply the same argument that this ability is important to responding to objective normative reasons that underly these explanations. The logical or causal relationships that we must appreciate in order to improve grasp are those that improve responsiveness to normative reasons. For example, in appreciating our reason to infer that Julia is not at the Rose and Crown because she is at the Royal Oak, we must manifest an ability to appreciate

the necessary logical relationship between these propositions and this requires an appreciation of its necessity. This is part of what it means to see the reason status of this fact. In the case where the server hits a table and spills coffee, the server recognises the fact that they did not sleep well as the reason that it is by appreciating the relevant causal dependencies that brought about the spill (Grimm, 2014: 330).

3.3.8 Bourget – Phenomenal Account

David Bourget (2025, 2017) argues that we can think of grasping in terms of phenomenological consciousness, or the ‘what it’s like’ element of our mental states that the agent experiences, such as colour, emotional feeling, conscious thoughts and imagery. The idea is that when S moves from non-grasping to grasping, e.g. in a case like *Animal Suffering*, they undergo phenomenal experiences that they did not previously have.

While GG does not require that phenomenal experiences are involved in grasp, the role of phenomenal consciousness in grasping seems to be to provide access to reasons. In a case similar to *Animal Suffering*, Bourget claims that an agent’s experience of a slaughterhouse can improve their grasp of why they shouldn’t eat meat. The phenomenal view explains this because grasping via experience can lead to ‘behavioural changes that are not based on new information yet are rational because it results in representing our reasons in experiential ways that make them more apt to play the role they rationally ought to play’ (2025: 119). It seems to me, then, that the role of phenomenal consciousness in appreciating our reasons is precisely how it is relevant to cognition.

Another possible counter example to this idea may be a case where grasp seems to be directly to do with phenomenal consciousness in a way that does not improve our responsiveness to reasons. For example, we might say that a person does not grasp redness if they have never experienced redness (Bourget, 2017: 286). They might know propositions about what is red and what is not, but they do not grasp such propositions. There are at least two responses to this. The first is that our ‘grasp’ of redness is a different sense of ‘grasp’ and does not target the epistemic success in question in the understanding literature. This is not particularly devastating for GG, at least not as devastating as if GG could not say anything about explanatory understanding, for example. However, there is a

second possible response that can explain grasp of redness in terms of objective normative reasons.

We can explain grasp of redness in terms of normative reasons if we can say that what is gained, epistemically, by experiencing redness is to do with the reasons that redness provides. It seems to me that experiencing redness might enable responsiveness to reasons that could easily be missed by those who have never experienced redness, such as a reason for thinking that bright red would be an unusual choice for the icing of a birthday cake. In the same way an experience of a blue note in music might be a reason to stray from the key in a piano piece and perhaps it would be a reason you couldn't appreciate without hearing a blue note. Of course, we can still maintain that it is possible to achieve cognitive contact with these reasons without experiencing red or hearing a blue note (and this would constitute grasp according to GG), but then I would be tempted to conclude that the phenomenal experience itself is not actually an improvement in grasp because it does not contribute anything to our epistemic position. The reason we are hesitant to resist this comes from the fact it would be such an unusual case. It seems to me that the intuition behind the idea that experiencing redness is an epistemic improvement comes from the fact that, via an experience of redness, we become better able to respond to the range of reasons redness may present in different contexts.

3.3.9 Bengson – Noetic Account

John Bengson (2015) argues that, with a change in understanding, there is a change in conception. A conception of a thing, on this view, is 'how one conceives or thinks, or is somehow inclined to think, of it' (2015: 651). This includes: (i) an attitude of conceiving of, (ii) an object of the conception and (iii) content that characterises its object in a certain way. In Bengson's example, we might conceive of the environment as something sacred. To use the example of *Anthropocene*, S conceives of the present geological epoch as the Anthropocene or the last of a species as an endling.

Conceptions have many properties and Bengson emphasises four that are particularly important to understanding: correctness (the content characterises the object as having the features it does have), completeness (the content characterises the object as having its central features), internal coalescence (identifying pertinent substantive connections

between an object's central features), and external coalescence (consistency with other relevant conceptions). Understanding some part of the world X, on this view, involves conceptions that are to some degree correct, complete, and internally and externally coalescent. Bengson calls such conceptions *noetic*.

One thing that it is important to see here is that Bengson, much like others, takes understanding to be context-sensitive, writing:

whether one qualifies as possessing understanding may vary from context to context, depending upon what is central, pertinent, or relevant in that situation (e.g. a botanist's conception of cocoa beans with a content of the sort described above may qualify her as understanding cocoa beans at an academic conference, but not in a kitchen populated by gourmet chefs whose task is to make dessert) (Bengson, 2015: 654).

The explanation for this, on my view, would be that the reasons that apply to the academic are different from the reasons that apply to the chefs. This is why some conceptions provide understanding in one context and not another.

With this in mind, we can see the relevance of noetic conceptions in our responsiveness to objective normative reasons. The argument goes that the way we think about things, or are disposed to think about them, affects how responsive we are to the reasons that they present. Some conceptions characterise objects in a way that is helpful in emphasising our reasons, such as the conception of ending in *Anthropocene*, which arguably emphasises the vulnerability of the last living being of a species and evokes the sadness of extinction events in a way that emphasises the moral reason to mitigate against them. If we think the conceptions that improve grasp are the ones that emphasise the most important reasons, then much of Bengson's noetic view could be accommodated by GG.

3.4 Advantages

The previous section did not provide conclusive arguments against other accounts of grasping. Instead, I defended the manner in which my account departs from them and suggested some possible compatibilities. Unlike some of the other accounts I discussed, I

have not developed grasping by characterising a particular ability or way of representing the world. This might be posed as an objection to my view: it is not substantive enough. However, we are now in a position to see how this could be conceived as an *advantage* of my view. This is an advantage because GG offers a substantive thesis about how degree of grasp is measured, but it is a thesis that allows multiple ways of successfully grasping, some of which are captured by other accounts of grasp in the literature. Because of this, GG (i) identifies a sense of grasp that is general enough to capture what unites different senses of grasp and (ii) can explain what is right about several other accounts. If you agree with me that grasp aims at objective normative reasons, and you also agree that features of grasping discussed by other accounts have potential to contribute to our appreciation of these reasons in different ways, then the following question arises: Why pick one? Perhaps there are a multitude of abilities and representations that can improve our epistemic position to respond to reasons, in which case GG explains what is compelling about all these other accounts of grasp.

In what remains of this section, I present some further advantages of GG, which is that GG provides insights into the distinctiveness of understanding and the value of understanding.

3.4.1 Distinctiveness of Understanding

GG provides insights into the distinctiveness of understanding. To expand on my claims in Section 3.1, I argue that grasping does not reduce to propositional knowledge because (i) knowledge requires true belief and grasp does not and (ii) taking something as a reason involves abilities to apprehend non-propositional structure of relationships that do not reduce to knowledge. The former claim will be discussed in depth in Chapter 4, but here I will focus on the latter claim. In clarifying the sense in which grasp is a distinctive feature of understanding, I highlight a potential advantage of GG: it can tell a plausible story about a grand distinction between knowledge and understanding.

If GG is correct, understanding must be distinct from propositional knowledge in order to avoid the regress problem discussed in Section 3.1. Being in an epistemic position to respond to reasons requires something beyond knowing propositions and, therefore, propositional knowledge is not sufficient. However, if this is all there is to say then we

might worry that GG does not preserve much of a distinction between knowledge and understanding, as understanding is basically just well connected, or non-propositional knowledge. Consider that the abilities associated with appreciating a non-propositional structure of relationships will also often be required to know propositions because, in order to believe a proposition, we must have some possession of the concepts involved and this requires an appreciation of relationships that underpin this concept (Khalifa, 2017: 51-79). For example, in order to really believe that the fire was caused by faulty wiring, I must possess the concept of ‘because’ and apply it appropriately in this context, which seems to require that I have some command over the relevant explanatory relationships. Perhaps the more I know, the better my grasp must be. If so, knowledge and understanding are not particularly separable.

While it may be the case that the abilities associated with grasp are often required for instances of propositional knowledge, it’s still non-obvious that they stand and fall together:

Faulty Wiring Bias: Another witness of the fire happened to install the wiring; due to their employment, they know a lot about what faulty wiring is and the potential fire hazards. However, their ability to see how the faulty wiring might have caused the fire is affected by their resistance to the idea that they could have had a role to play in the fire.

The knowledgeable witness cannot make appropriate causal inferences due to biases that interfere with their ability to respond to facts, but not their belief in them. While beliefs that explicitly depict explanatory relations, e.g. p because q , often imply corresponding grasp, it’s possible in principle to separate the success of affirming what is the case and the success of having the cognitive abilities required to respond to what is the case. There are factors that impact what we are willing to believe, but there also seem to be factors that undermine the abilities related to seeing the relationships that underpin our reasons, and these are separable. Maybe we can have biases that dispose us to make the wrong inferences or apply incorrectly to new cases, like the knowledgeable witness, even if we really do believe all the relevant truths.

This does not solve the more serious problem, however, as it's clear that the distinction between understanding and knowing is still minimal in comparison to the stark contrast of the parent and child in *Faulty Wiring*. Christoph Kelp (2015) argues that, if one has comprehensive and maximally 'well-connected knowledge' of some phenomenon, then one has maximal understanding of that phenomenon, where knowledge in propositions is well-connected if the support relations one knows to obtain between features of an account of some explanation is reflected in the basic relations that obtain between one's beliefs. From what I can tell, this deals with the regress discussed in Section 3.1.2. If the realistic cases where relevant background knowledge and grasp come apart seem uncommon and potentially quite insignificant, we might begin to question the explanatory benefit of making a grand distinction between knowing and understanding, instead considering understanding to be just well-connected knowledge.

Here is a possible grand distinction between knowledge and understanding provided by GG. When we acquire more knowledge, it does not matter which facts each piece of knowledge is tied to. In contrast, because of its grasping element, an increase in understanding is partially defined by this question of what facts mean to us, i.e. whether they license certain responses. For example, our understanding of the cause of the fire increases as we arrive at reasons for important inferences about what happened and preventative measures, where our understanding of Epistemology increases when we arrive at compelling reasons to revise our views or redirect our inquiries. This provides an explanation for the distinctive epistemic value of understanding. Knowing is plausibly of epistemic value because of its relation to the truth, whereas understanding is of epistemic value because of its relation to normative reasons.

GG can therefore justify a theoretically important distinction between understanding and knowing. Although knowing and understanding will often stand and fall together, because we may need to know the truth to respond to normative reasons, knowing the truth is not always required. As I noted in Section 3.1.1, we can respond to normative reasons via a representation of the world that is not strictly true but preserves reason-giving features. Similarly, knowing may often lead to understanding, because we may know facts that present normative reasons, but this is not always the case. We can acquire more knowledge without an increase in understanding, where we acquire knowledge that does not present us with reasons. Knowledge may also undermine understanding, when it distracts us from the

most important reasons. On my view, understanding and knowing are distinct epistemic states that differ in their measure of epistemic value and may support or undermine one another. This last point brings me to another advantage of GG, which relates to its contribution to debates about the value of understanding.

3.4.2 The Value of Understanding

Many think that understanding is more valuable than knowledge. We should take particular interest in the following claim:

Value Claim (VC): understanding has a distinct epistemic value that is superior to knowledge, in virtue of its grasping element.

VC says that understanding has a distinctive epistemic value that is superior to knowledge and that this is explained by the presence of ‘grasping’. Note the force of the following intuition: if we take understanding to come apart from knowledge, we prefer to understand than to know.⁶⁰ If understanding is more valuable than knowledge, then it has something knowledge doesn’t that explains why this is.

It would make sense if grasping explains the superior value of understanding. Unlike other candidates for distinctive features of understanding, which appear to be ways in which understanding requires *less* than knowledge, e.g. that understanding has less stringent truth conditions or that it is immune to epistemic luck, grasp is supposed to go *beyond* knowing. It is supposed to offer something that knowledge does not offer. Because of this, grasping is a good place to look for a potential source for the additional value of understanding.

I have offered some reasons to think that the *Value Claim* is true. There are also good reasons to think, if it is true, it is a powerful claim in epistemology that has some important implications. An interest in understanding is often driven by a dissatisfaction with the value of knowledge. Kvanvig argues that knowledge is not more valuable than any of its subparts combined and therefore does not have distinctive value, and that understanding appears to be more valuable than knowledge, so attending to the relationship between these epistemic

⁶⁰ This is an intuition shared by others (Kvanvig, 2003: 203; Pritchard, 2010: 74).

states ‘can give us hope in our pursuit of a special and unique value for epistemic achievements, even though we have had to give up such hope regarding the cognitive achievement for knowledge’ (2003: 186). Elgin argues that, unlike knowledge, understanding can explain the epistemic successes of science. If epistemology wants to accommodate the success of science, it should ‘shift its center of gravity from knowledge to understanding – in particular, to what has been called *objectual understanding*, the understanding of a topic or subject matter’ (Elgin, 2017: 3). VC has the potential to validate the strong claim that understanding is not just equally deserving of our attention in epistemology, but *more* deserving of our attention. We should seek to explain VC further.⁶¹

In particular, we need to explain what it is about grasping that affords understanding its superior value. It would be an advantage of an account of grasping if it could explain (i) how grasping distinguishes understanding from knowledge and (ii) how it contributes to its superior value. My approach to grasp offers a novel and particularly compelling answer to these questions about VC. To demonstrate this, let’s consider possible explanations that follow from other accounts I have discussed.

As we have seen above, many think of understanding (and, in particular, its grasping element) as an internalist notion, which means that explanations of the *Value Claim* that naturally follow from these accounts carry internalist commitments. However, the more we think of grasp as an externalist notion, e.g. that might require truth or other features associated with knowledge, the less convincing it is that grasp contributes something that we do not already have by knowing. This is no conclusive argument against accounts that offer internalist answers to the *Value Claim* or accounts on which grasp is closer to knowledge. However, this does seem to uncover three options: (1) endorse some version of internalism, (2) reject the *Value Claim*, or (3) a novel account of grasp that is not internalist, but distances from knowledge.

An accessibilist account of grasping could explain VC because accessibility is not present in knowledge, but we would need to further argue that what is accessible is of superior value. One argument might be that accessibility is particularly important because this is what we can hold an agent responsible for. This argument is at the centre of contentious

⁶¹ For an overview of these discussions about the value of understanding, see Grimm (2012: 103-104).

debates between internalists and externalists about epistemic justification. Recall Bonjour's (1985: 41) clairvoyance case, in which Norman reliably arrives at true beliefs via clairvoyance, but has no reflective access to any reason for thinking he has this power. While Bonjour focuses on the notion of what is responsible for Norman to believe from his perspective, others who deny accessibility theses might have a different set of concerns in evaluating belief, such as whether belief is appropriately connected to the truth. Another possible argument would be that understanding would be more resistant to scepticism than knowledge (Zagzebski, 2001: 247). These possible arguments for the additional value offered by accessibilist grasp require us to share the sorts of motivations and intuitions of internalists. As noted in Chapter 2, Section 2.4, the difference in internalist and externalist positions often reveals significant disagreements about what matters from the epistemic perspective and the appeal of an accessibilist defence of VC will depend on where one stands on this issue.⁶² Remember that, to explain VC, we would not only have to argue that internalist considerations matter, but that they are *more* important than externalist considerations.

Moreover, any view that says grasp is subject to epistemic luck would carry some form of internalist commitment in their defence of VC. Hills argues that understanding is valuable for its own sake, by analogy with true belief; true beliefs are valuable because they accurately reflect the way things are and are therefore a 'mirror of nature' (Hills, 2016: 679). Much like true belief mirrors the world in content, Hills argues that grasp also mirrors the world in form. When we grasp the world the dependencies between our beliefs mirror dependencies in the world. Hills (2016: 680) acknowledges that, given the costs of understanding, whether we should prefer to understand than to know depends on context, such as whether testimony is available and whether one needs the right answer quickly. Nevertheless, this approach can offer a good explanation for VC that says that grasp is usually more valuable, given that it offers mirroring of the world in both content and form.

However, this explanation of VC carries an internalist commitment that is revealed once we attend to the sense of mirroring at issue. The sort of mirroring in content and form, as

⁶² This most obviously applies to stronger accessibility conditions, but might still be relevant to other accounts that centre reflection. For Elgin, successful understanding involves accounts that 'the community of inquiry can, on reflection, endorse' (2017: 66). An account like this may also have the resources to explain VC, but on internalist terms.

present in *Faulty Wiring Gettier Luck*, is missing interaction with the part of the world that it is about. Your beliefs about the cause of the fire and their structure do not reflect causal dependencies in the world, *because* of the way the world is in this case; they happen to cohere with the way the world is. This means that to take this sort of success to be what makes understanding more valuable than knowledge would be prioritise properties of mental states that happen to cohere with the world over those that are affected by what happens in the world in the way that knowledge is. It is an interesting question to what extent we should take issue with this internalist commitment. I only point out that it constitutes an internalist commitment.

Some accounts of grasping could offer an account of its value without internalist commitments. On Grimm's modal grasping account, this is possible because modal relationships are relationships in the world and it is implausible that we can achieve grasp of them via lucky guesses and deception, we need some non-accidental connection to how things are in the world. However, Grimm (2014: 338) welcomes the idea that knowledge of propositions, such as 'the house burned down because of faulty wiring', may entail a degree of grasp of the relevant causal relationships. This is because it seems difficult, and perhaps impossible, for the child in *Faulty Wiring* to genuinely assent to a proposition that invokes 'because', without having some grasp of how the fire depends on the faulty wiring. Importantly, Grimm takes understanding to be a form of non-propositional knowledge. A view on grasp that is tied closely to knowledge in this way could undermine the importance of VC. It makes sense that externalist-leaning approaches to grasp end up looking a lot like knowledge because the conditions that make grasp externalist, e.g. truth and non-accidental connection to the world, also bring grasp closer to knowledge.

It is a point in favour of GG that it can provide a clear explanation of the distinctive and superior value of understanding without internalism. My view explains VC because it says grasp has a distinctive value due to the fact it aims at objective normative reasons. This is valuable because normative reasons are invariably parts of the world that mean something to us, in that it is a good thing when we are in a position to respond to them. This view can also explain why grasp contributes to the superior value of understanding. Insofar as it is more valuable to be in contact with normative reasons than it is to know a fact, whatever that fact may be about, GG can explain the intuition that we should prefer to understand than to know without internalist commitments.

3.5 Conclusion

In this chapter, I have argued for some conditions on what it takes to be in an epistemic position to respond to objective normative reasons and I have made the case for the thesis that the degree to which one grasps corresponds to the degree to which one is in an epistemic position to respond to objective normative reasons.

This chapter suggested that grasping requires a mental representation of features of the world that preserves their reason-giving force. While this often involves true belief or knowledge of facts that present reasons, truth is not always required to preserve reason-giving force. In Chapter 4, I elaborate on how representations of the world might deviate from the truth in a way that not only preserves reasons, but actually improves our responsiveness to them.

Chapter 4: Understanding and Truth

4.0 Introduction

So far, this thesis has focused on the argument that understanding is distinct from knowledge because it has an additional grasping element that goes *beyond* knowledge. But many think that there is a sense in which understanding requires *less* than knowledge. Some argue that, unlike knowledge which is commonly thought to require true belief, understanding is not subject to such stringent truth requirements (Kvanvig, 2003; de Regt, 2015; Elgin, 2017). This chapter argues that the approach to grasp defended in Chapter 3 validates this claim about understanding.

In what follows, Section 4.1 argues that we can become responsive to objective normative reasons despite (and sometimes because) we accept an account of the world that deviates from the truth. If this is right, there are two important upshots. The first is that grasping, and therefore understanding, does not require true representation and does not necessarily improve with accuracy. This validates a non-factive view of understanding, according to which understanding can be constituted by a primarily or centrally false account of the world. The second upshot is the strengthening of my arguments against veritism in Chapter 1. I argued in Chapter 1 that truth is not the sole fundamental epistemic aim or concern and that it is not an important epistemic aim or concern relative to grasping. However, truth can maintain a central place in epistemic normativity if true representation turns out to be necessary for grasping. I will show that true representation is not necessary for grasping and therefore that truth does not have this central role. The remaining sections tackle some objections to my view. Section 4.2 defends my view that understanding is non-factive against those who argue that, upon examination, the content of our understanding is true. Section 4.3 defends my view against the objection that we can ultimately explain the epistemic value of the falsehoods in the content of our understanding as a means to true belief. Finally, Section 4.4 defends my view against the objection that we can explain the value of falsehood in understanding in terms of ‘truth-oriented’ values, i.e. values that bear some essential relation to truth. I conclude that, insofar as understanding is valuable

because it involves grasp, its value does not require, and is not explained by, reference to true representation.

4.1 How Exemplification Enables Reason-Responsiveness

Understanding differs from knowledge in virtue of its relationship to truth. The more we know, the more truth we believe. In contrast, a higher degree of understanding does not necessarily get us closer to the truth. This is because understanding centrally concerns being in an epistemic position to respond to reasons and true representation is not required to be in an epistemic position to respond to reasons. This section provides a clearer idea of how this works, by building on the notion of *exemplification* found in Elgin's work on understanding.

4.1.1 Elgin and Understanding

Elgin (2017) argues that any adequate epistemology must be able to explain the epistemic value of scientific understanding and veritist epistemology cannot achieve this. Veritism is defined by Elgin (2017: 9) as 'any position that takes truth to be necessary for epistemic acceptability' and that assesses epistemic practices in terms of whether they are truth conducive. Contrary to veritism, scientific understanding is an important epistemic success, despite the fact it often deviates significantly from the truth. Our scientific understanding of the world is enriched by idealisations, models, thought experiments, curve-smoothing, or abstractions and stylised facts that falsely represent the world. Because of this, Elgin argues that we must either reject veritism or take scientific understanding to be defective or non-cognitive (2017: 16).

To get a clearer idea of what is going on here, consider Elgin's claim that we can acquire understanding of the world by accepting an account, where an account is construed as a 'constellation of mutually supportive commitments that bear on a topic' (2017: 12). These commitments might include a range of beliefs, acceptances, and evidential commitments. An account represents the world in some way, and thereby must refer to the world in a way that answers to the facts, but Elgin argues that an account of the world need not be true to figure in our understanding.

There are at least two ways in which an account might deviate from the truth and constitute scientific understanding. First, scientists in the past could acquire understanding from theories that were false and are now outdated. Elgin thinks that epistemology should recognise the understanding scientists had in the earlier stages of this process (2017: 60). Second, science employs accounts of the world that do not even purport to be true and that are not improved by accuracy. For example, idealisations and thought experiments are centrally or primarily false, but they can still be indispensable to scientific understanding. For understanding to be an epistemic success, it must provide epistemic access to the world, but this is always via true representation. If not true representation, something else must tether an account to the portion of the world it pertains to, such that when we accept the account, we can understand the world. Elgin argues that this tether to the facts is exemplification.⁶³

4.1.2 What is Exemplification?

Exemplification is the relationship that holds between an example and the thing that it is an example of. To borrow Elgin's terminology, call whatever is doing the exemplifying an *exemplar* (2017: 179). A sample card can serve as an exemplar of a particular colour. The sample card exemplifies this shade of yellow in the following way: it refers to this shade of yellow by instantiating this particular shade of yellow and downplaying other properties, such as being rectangle or being made of paper (Elgin, 2017: 187-188). In doing so, the sample card affords epistemic access to this shade of yellow.

Exemplification provides epistemic access to the world by making a feature of the world salient. The sample card makes the shade of yellow salient to the agent who interprets the sample card properly. Another important point for my purposes is that exemplification does not just highlight what is already salient, it can also make previously unnoticed features of the world apparent to us. We can come to view some feature of the world as salient because of the way that it is exemplified in our account. An exemplar can lead an agent to find a particular feature of the world significant by virtue of being better equipped to recognise it and view it in a certain way. In Elgin's words, exemplification 'is not just a vehicle for

⁶³ For more on exemplification, see Elgin (2017: 183-203).

conveying salience; it is a source of salience' (2017: 187). When we come across the same feature, it now appears salient to us in a way that it did not before. This makes exemplification useful when we encounter unfamiliar objects.

We should take great interest in the role of exemplification in our understanding of unfamiliar things, such as beings, processes or other entities that exist on very large or very small scales. Recall that part of the epistemic problem in Section 0.1 is that to understand the environment, is to understand unfamiliar things that are at the centre of environmental discourse and that are often difficult to grasp due to our lack of direct or extensive experience with them. Some argue that we are often out of touch with the entities and processes at the centre of environmental issues and this presents a serious epistemic barrier in our understanding of present environmental challenges (Jamieson, 2007; Sterman, 2011). It is notable that exemplification explains why it is helpful to make comparisons with objects that are closer to our relative size. It is, for example, common to demonstrate how big planets are relative to each other by using spherical objects closer to our own size, e.g. a beach ball, a tennis ball and a pea. We might compare the size of the Great Pacific Garbage Patch to twice the size of Texas. In this case, Texas twice exemplifies the size of the Great Pacific Garbage Patch and, by virtue of being a more familiar reference point, emphasises its size. We might also compare accumulation of CO₂ in the atmosphere to the flow in and out of the bathtub (Sterman, 2011). When there is more water flowing into the bathtub than there is leaving, we can have movement both ways and still have bathtub that is gradually filling up. The bathtub can exemplify the accumulation of CO₂ in the atmosphere and the dynamics of sources of CO₂ and sinks. It does so by instantiating the feature of some substance flowing into a finite space faster than it is able to exit and then downplaying other properties. In these cases, we are understanding unfamiliar, large-scale changes and processes via familiar exemplars.

Exemplification does not fundamentally concern true representation and, as Elgin argues, falsehood can contribute to the efficacy of an exemplar. For example, each data point on a graph represents an individual truth, e.g. the temperature at a particular time, but a graph will often include a smoothed curve that does not pass through each data point, but reveals a pattern in the data. The curve does not represent the individual data points, and therefore does not represent the truth about the relevant relationship between the time and temperature. Instead, the curve on the graph exemplifies this relationship, by instantiating

some approximation of a rate of change that also applies to the temperature increase and drawing attention to it by excluding anomalies. In Elgin's words 'to accommodate every point would be to abandon hope of finding order in most data sets, for jagged lines and complicated curves mask underlying patterns and regularities' (Elgin, 2017: 24). Smoothing the curve affords us epistemic access to the rate of change, precisely because it deviates from the truth.

4.1.3 What are the Success Conditions?

If we do not necessarily improve our understanding by getting closer to the truth, and successful exemplification has something to do with this, we must identify what makes it such that exemplars, like the bathtub, the size of Texas twice, or the curve on a graph, contribute to understanding. What makes an exemplar effective? It is at this point that my view comes apart from Elgin's internalist answer to this question.

Elgin argues that, to contribute to understanding, an exemplar must not be misleading:

A mining inspector extracts air samples from a mineshaft to find out something no one yet knows about the distribution of gases in the mine. If the samples are properly taken, he has reason to believe that the gases his samples exemplify are typical of the gases in the mine. A pollster questions potential voters about their attitudes towards immigration. If the sample is suitably chosen and the questions well framed, she acquires new information about public opinion on the topic. There are, of course, no guarantees...if the samples are well taken and the background theory is well founded, there is epistemically good reason to project to a wider class of cases (2017: 191).

On this view, whether an exemplar is epistemically acceptable turns on whether it is responsibly chosen; this aligns with Elgin's view on epistemic responsibility as the responsible management of one's epistemic commitments to achieve reflective equilibrium. This is an internalist-leaning position on what makes a good exemplar, according to which what matters is whether we have good reason to think an exemplar is not misleading relative to our epistemic commitments.

We could also construct an externalist-leaning position about what makes a good exemplar based on my account of grasping as being in an epistemic position to respond to

objective normative reasons. I have argued that objective normative reasons are facts, independent of the perspective of an agent, that count in favour of a certain responses. I propose that whether an exemplar is effective concerns whether it preserves the force of our objective normative reasons, as defined in Chapter 2.

Here is how I think this should go. To achieve responsiveness to objective normative reasons presented by some portion of the world X , the mental representation we have of X must have some degree of accuracy – as argued in Section 3.1.1. However, the mental representation of X does not have to truly represent the facts that present reasons. We do not need true beliefs about facts that objectively count in favour of certain responses in order to respond to them. Consider this case. An agent views a graph with a very steep line upwards that represents rise in global temperatures and they may achieve cognitive contact with normative reasons to infer certain beliefs (e.g. “climate change is very serious”), act in a certain way (e.g. reducing their emissions) or feel certain emotions (e.g. shock or fear). This graph does not truly represent the normative reasons for these beliefs, actions and emotions, because these are the facts about global temperatures, i.e. the data points, rather than the curved line on the graph. As before, the graph exemplifies some approximation of the relevant rate of change.

The graph preserves normative reasons for beliefs, actions and feelings via exemplification. This is possible because the steep line shares a property with a part of the world, i.e. a range of a rate of change that rise of global temperatures might also fall into, and this property is reason-giving. The rate of change presents normative reasons for the agent to respond in certain way and this is made salient by the graph because the graph instantiates this property and downplays other properties in a way that, when properly interpreted, helps the agent recognise their reasons. This would be an externalist-leaning position on exemplification because whether the exemplar misleads us is not relative to our epistemic commitments, but a matter of whether our account exemplifies a property in a manner that preserves and emphasises the reason-giving force of objective facts.

Despite this important disagreement with Elgin, I endorse her view on another dimension of evaluation of our exemplars, which is to do with how well an exemplar can be interpreted. In order to be interpreted properly, an exemplar must selectively emphasise the property it refers to. Elgin suggests that exemplification can fail when the exemplar is

difficult to interpret. For example, if the relevant property is camouflaged, overshadowed or not present enough in the chosen exemplar (2017: 191-192). Often, whether an exemplar can be properly interpreted depends on the context in which something is made an example of. In the case of the sample card, the context is the accepted system according to which each sample card serves as a symbol of a range of colours around the shade appearing on the card, all close enough to be included under that shade, e.g. ‘jonquil yellow’. In other contexts, exemplars can be produced spontaneously just by treating something as an example, e.g. by stating you are using an example or gesturing.

With this in mind, we can see that whether an agent can interpret an exemplar is, to some extent, relative to their circumstance. Consider the previous example of the agent who improves their understanding of the size of the Great Pacific Garbage Patch by learning that it is twice the size of Texas. We could imagine another agent who is unfamiliar with the geography of the United States, for whom this might not be a good exemplar. Instead, their understanding improves in virtue of learning that the Great Pacific Garbage Patch is three times the size of France. For one agent ‘Texas twice’ selectively emphasises the property of being roughly the size of the Great Pacific Garbage Patch, but for another agent it does not, and this is due to whether they have relevant experiences to interpret the exemplar properly.

Another way that exemplification may go wrong is if an exemplar is only accidentally effective. For example, in a case where an agent falsely believes that Texas is the size of France and then mishears that the GPGP is three times the size of Texas. In this case, the agent makes two errors; the first leads them astray, but the second fortunately leads them to a representation of the world that gives them a sense of the size of GPGP. What is worrying about this case is that the representation of Texas three times appears to preserve the reason-giving force of the size of the GPGP, in a way that suggests that grasp is present. However, upon examination, this is just a Gettier case that resembles the stopped clock case (discussed in Section 3.1.3), in which we make the mistake of looking at a clock that is broken, but by chance we are led to the truth because we happen to look at the clock at the time of day that it is correct. So, it turns out that the intuitions that lead us to think that the agent might grasp via Texas three times in this case are the same intuitions that lead others to think we grasp in *Faulty Wiring Gettier Luck* (Section 3.3.2). To make this connection clear, consider a case that is directly parallel, where someone impersonates a

lecturer and, despite having no knowledge about the changing climate, draws a line on the whiteboard that might have been an effective exemplar for the students to grasp rising global temperatures. This case is like *Faulty Wiring Gettier Luck*, in which you happen to acquire an accurate explanation for the cause of the fire from a partygoer. I have already suggested an alternative approach to these cases, according to which the genuine grasping in this case is a grasp of a representation that happens to cohere with the facts, rather than the facts themselves. For the same reasons, we should maintain that exemplification is part of grasp to the extent that it is a tether to the facts and therefore ‘exemplars’ that are only accidentally connected to the facts, such as the exemplar of Texas three times in this case, cannot constitute genuine grasp of the GPGP itself.

If all this is right, then we can achieve cognitive contact with objective normative reasons by accepting an account of the world that exemplifies properties of the world that present these reasons, even if this account is primarily or centrally false. This vindicates the view that understanding is non-factive. To situate this view within the literature, attend to the distinction between the view that understanding is quasi-factive and the view that understanding is non-factive. Some think that understanding is quasi-factive because, although the content of our understanding may tolerate falsehoods, the content of our understanding must be mostly true and the falsehoods in question must be peripheral.⁶⁴ Those who think understanding is non-factive, as I do, argue that the content of understanding can contain central falsehoods.⁶⁵

Moreover, the role of exemplification in understanding shows that it is not just that our account of the world can tolerate deviations from the truth and still preserve reason-giving force of these properties of the world, but also that it can be *preferable* to include falsehoods.⁶⁶ This strengthens the anti-veritist view put forward in Chapter 1. We do not necessarily improve in our epistemic position to respond to reasons by believing more truths or by accepting accounts of the world that get closer to the truth. Often, we can better respond to our reasons by deviating from the truth. We are finite beings with limited

⁶⁴ Some argue that the content of understanding can contain *peripheral* falsehoods (Kvanvig, 2003; Mizrahi, 2012).

⁶⁵ Some argue that the content of our understanding can contain *central* falsehoods (de Regt, 2015; Elgin, 2017).

⁶⁶ This supports Elgin’s point that often accounts of the world might offer more valuable understanding in virtue of being false (2017: 16).

time, attention and resources, and we are also the sort of beings for whom some features of the world matter more than others, because of the reasons they present to us. In this context, we benefit from accepting accounts of the world that selectively emphasise the properties of the world that present the most important reasons, even where this comes at the cost of truth. This selective emphasis on our most important reasons is how to best to achieve an epistemic relation to the world that guides us to respond to the world in the right way.

In what follows, I address some objections. I have grouped these objections in a way that best highlights their specific threat to my claim that GG vindicates the non-factive view of understanding and undermines veritism. The first objection attempts to show that the content of understanding is always true, e.g. an approximate truth or counterfactual truth, by challenging the claim that true representation is not required for reason-responsiveness. The second objection is to say that false accounts are only valuable to understanding as a means to relevant truths. This presents a particular threat to my account, if exemplifying reason-giving properties contributes to valuable understanding by preserving reasons to believe truths. The third objection says that, even if accounts that figure in understanding are not always true or valuable as a means to truth, the epistemic value of these accounts will nevertheless be truth-oriented. This means that their epistemic value always bears some essential relation to truth, i.e. because what is of concern is approximate truth, truth in models, or because representations include ‘something true’, such as a true pattern or property. The third objection presents a particular problem for my externalist account, which says that the force of objective normative reasons, which are facts, must be preserved in an account that provides understanding. If an account preserves facts in this way, it seems there must be some truth to it. I discuss each objection in turn, arguing that none of them successfully demonstrate that true representation is required (or even preferable) for grasping. This chapter therefore shows that my take on exemplification, as a tether to our reasons, can act as a defence against objections to non-factive understanding.

4.2 Truths in the Content of Understanding

Some claim that we can explain the epistemic value of idealisations and thought experiments, because they express or are partially constituted by a belief in some other

relevant truth. The strategy is to show that the content of our understanding is true, or at least that the central or important parts of the content of our understanding are true.⁶⁷ One point in favour of this view is that false accounts of the world often do not afford understanding because we take them to be true. When scientists make use of idealisations or philosophers make use of thought experiments, they do not believe in these falsehoods. It seems, then, that the source of their epistemic value is not the falsehoods explicitly included in the account, but some other relevant truth that these scientists or philosophers *do* believe.⁶⁸ Once we keep this in mind, perhaps we can see that the role of falsehoods in understanding does not undermine veritism after all. We could construct an argument here that undermines my claim that we can have an exemplification tether to our reasons that does not turn on truth. This section discusses possible candidates for the relevant true belief (or set of true beliefs) and denies these candidates can account for the role of exemplification. I argue that deviating from the truth is often indispensable in our responsiveness to normative reasons and falsehoods are both part of the content understanding and have direct epistemic value virtue of this.

4.2.1 Approximately, p

Some accounts deviate from the truth in a way that makes the world intelligible, but they are still approximately true, such as the example of ‘curve smoothing’. The curve on the graph represents the world inaccurately, because the curved line does not run through the data points and therefore does not represent something that obtains in the world in the same way as the data points do. I have suggested that the curve provides epistemic access to our reasons because the line changes in roughly the way that the temperature changes, to a degree of accuracy that preserves reasons for action, beliefs and feelings that follow from

⁶⁷ Here are some examples of this general approach. Michael Strevens (2017) argues for the *difference-maker approach*, which says that an explanation does not speak on the entire causal history of an event, only the parts of this history that make a difference to whether or not the event happened. On this approach, idealisation may deliver understanding by falsifying factors that do not make a difference. Insa Lawler (2021) argues for the *extraction view*, which says that the content of understanding is true information that is extracted from an account of the world that may include falsehoods. On this view, certain falsehoods are important because they capture something true about the phenomenon of interest and therefore enable our understanding, but these falsehoods are not an element of the content of our understanding. Kareem Khalifa (2020) argues for *inquisitive veritism*, which concerns true answers to relevant questions. This approach attempts to show that understanding often involves true answers to relevant questions (or means to such true answers, which is a different strategy I cover in the next section).

⁶⁸ Scientists will be aware of deviations from the truth, such as false assumptions in idealisations, and so they cannot believe in the idealisations themselves (Mizrahi, 2012: 245).

the rise in global temperatures. The veritist can say that this sort of inaccurate representation of the world expresses a relevant truth of the form ‘approximately, p’.

To make this point, we should distinguish between (i) the true claim ‘approximately, p’ and (ii) the false but approximately true claim that ‘p’ (Khalifa, 2020: 947). The veritist might then claim that while scientists state false claims (ii), they intend to communicate true claims (i). It is, after all, built into the practice of ‘curve smoothing’ to acknowledge that the curve itself does not connect the data points and therefore does not strictly represent the data. When interpreting the graph, we recognise that the curve only approximates the data itself. Perhaps we can argue from here that it is not the falsehood that provides access to the reason, and features in the content of our understanding, but the true belief in ‘approximately, p’.

While it’s true that stylised facts, abstraction and curves on graphs etc. support ‘approximately, p’ beliefs and this partially explains why they are epistemically valuable practices, this fact does not fully explain the role of these falsehoods in our understanding. On my view, this is because ‘approximately, p’ beliefs cannot explain how the agent who accepts these false accounts of the world becomes capable of responding to their reasons. We can imagine an agent that has beliefs about approximate global temperatures, but who is not guided by these beliefs to feel shock or to act, in the way that they would be guided by the visualisation of the graph. Such an agent may require the mental representation of the graph to be guided in this way, because the exemplification relation makes their reasons for shock, and reasons for action, salient in a way that ‘approximately, p’ beliefs do not. Moreover, even in agents who see the graph *and* believe ‘approximately, p’, it may not be affirmation of the truth involved in the ‘approximately, p’ belief that moves them to act in accordance with their reasons.

The belief in ‘approximately, p’ will not always be what is most important in the content of an agent’s understanding, because this belief is not always the aspect of one’s mental representation of the world that consistently guides them to respond to their reasons. Recall the case of *Animal Suffering*, from Chapter 1. In this case, S is moved by experiencing the video of the cow, not the belief that they are morally required to avoid contributing to unnecessary animal suffering. One interpretation of this case is that S’s beliefs about their moral obligations, while true, are somehow unintelligible, but the cow’s individual

suffering makes them intelligible by being a good example. The cow's suffering selectively emphasises reasons not to contribute to unnecessary animal suffering. In this case, we can imagine that S regularly remembers the image of the cow and it is this representation of the world that guides S to consistently respond to the relevant moral reasons. We might say something similar about the curve on the graph. We can imagine that the agent moved to shock after viewing the graph will continue to be guided by the curve on the graph, rather a belief of the form 'approximately, p'. If it is the graph, and its use of exemplification, that consistently facilitates reason-responsiveness, then it is the specifics of this representation that are directly and indispensably operative in one's understanding, including its departures from the truth.

4.2.2 Counterfactual Truths

Some accounts of the world appear to be part of our understanding despite the fact they are not even approximately true, e.g. idealisations, models and thought experiments that rest on false assumptions. However, this could be because they tell the truth about what would be the case if certain false assumptions did obtain. The veritist might say that, in such cases, the content of one's understanding is a counterfactual truth.⁶⁹ I argue here that, similarly to 'approximately, p' beliefs, it is not belief in counterfactual truths that enables reason-responsiveness. Therefore, my view on understanding is resistant to this objection.

Consider the thought experiment, *Last Human*, which invites us to see the value of nature independently of its relationship to humanity, by constructing a scenario where the last human on Earth is given the power to destroy what remains of nature (Sylvan (Routley), 2008: 7). The counterfactual truth in this example would be something like 'if the last human on Earth had the chance to destroy what remains of nature, it would be wrong for them to do so'. This also comes out true, assuming the thought experiment does what it is supposed to.

⁶⁹ Collin C. Rice argues that models can produce understanding despite the fact they do not accurately represent the phenomenon one understands, but that the understanding produced is 'constituted by a set of *true* beliefs about the counterfactual relevance and irrelevance of various contextually salient factors' (2016: 88). These true beliefs about what is possible are then integrated in a more comprehensive account of the phenomenon. Note that Rice still supports a quasi-factive view of understanding, because these true beliefs are part of a larger account that is only mostly true (2016: 88).

The first thing to say about these counterfactual truths is that they are truths about what is possibly the case. In order for truths about what is possibly the case to constitute understanding about what *is* the case, they have to at some point become about the actual world. This is unproblematic by itself. In *Last Human*, there is an inferential pathway from the conditional claim ‘if the last human were to destroy what remains of the natural world, it would be morally wrong’ to something like ‘nature has some moral value independent of humanity’. *Last Human* gets rid of any interfering factors that might prevent us from making this verdict in a more realistic scenario. If we destroy a part of nature in the actual world, it remains a possibility that another human might now, or at some point in the future, be affected by this and we could explain the moral wrong of this action by appeal to this risk. *Last Human* establishes that something of value would be lost if nature were destroyed, even if no humans would be affected by this. If the moral rules are the same in the last human possible world and ours, we can easily make inferences about morality from this thought experiment.

In order to undermine veritism, what we really need to show is that belief in counterfactual truths and comparisons with the actual world are not doing the work in providing epistemic access to reasons. On my view, the thought experiment *Last Human* provides cognitive contact with reasons via exemplification, by instantiating the properties that present moral reasons for certain actions and inferences about morality and downplaying other features. The thought experiment is therefore an aspect of our mental representation that enables reason-responsiveness, precisely because it is inaccurate. Much like the example of the graph, if the thought experiment itself is the aspect of the mental representation that provides cognitive contact with moral reasons to think that nature has independent moral value from human interests, then it is the thought experiment that provides crucial epistemic access to the world and not a counterfactual truth. Even if underlying counterfactual truths explain the efficacy of the thought experiment, it is this specific mental representation that makes reasons intelligible to the understanding agent, including its falsehoods. The falsehoods do direct and indispensable work in the agent’s understanding, by making the agent’s moral reasons salient to them, and should therefore be considered part of the content of their understanding. In the next section, I address another objection that says the value of falsehoods in understanding, such as idealisations and thought experiments, could be part of the content of our understanding, but they are only epistemically valuable as a means to truth.

4.3 Understanding as a Means to Truth

Another objection attempts to explain the epistemic value of falsehood in our understanding in terms of being a *means* to true beliefs. The difference between this strategy and the previous strategy is that the veritist might admit that falsehoods can be included in our understanding, and that they can even be a central and important part of our understanding, but only because they are a means to truth. This objection poses a particular threat to my account because my arguments in Section 4.1 rest on the idea that we can understand via a false account of the world if the account in question preserves or emphasises the reason-giving force of objective facts. There seems to be a readily available way to talk about this in terms of means to the truth, which is that, by achieving this contact with our reasons, we arrive at reasons to believe truths.⁷⁰ This section argues that the epistemic value of non-factive understanding is not exhausted by its value as a means to true beliefs.

It could be argued that accounts of the world that deviate from the truth are only epistemically valuable as a means to true belief. By selectively emphasising some features of the world over others, we are in a better position to appreciate important truths. Recall that the agent who sees the curve of on the graph can infer that “temperatures are rising quickly”, “climate change is a serious issue” etc. Similarly, the agent who considers the *Last Human* thought experiment can infer truths about the value of non-human entities independently from human interests and about what they ought to do.⁷¹ We could argue that the falsehoods in these accounts are instrumentally valuable as a means to true belief and, even worse, that my explanation only supports this, because the reasons that matter to grasp in these cases are reasons to believe truths.

I do not deny that falsehoods in our understanding are epistemically valuable as a means to true belief, because understanding will often be valuable, in part, because it preserves reasons to believe the truth. However, it is implausible that falsehoods in our understanding

⁷⁰ It seems that such reasons could be an inferential or non-inferential basis for true belief. We might think that inferences to true belief are made in accordance with logical or probabilistic relationships and therefore an account can preserve reasons for such inferences, by including or implying propositions from which we can infer truths. A non-inferential basis for belief might license true belief in accordance with certain perceptual experiences or memories, for example, and an account might include or provoke such experiences and provide basis for true belief this way.

⁷¹ For a similar argument about idealisations in science, see Khalifa (2020: 951).

are valuable solely as a means to true belief because the manner in which an account deviates from the truth has additional epistemic value beyond being a basis for actual true beliefs or even being a potential source of future true beliefs.

Consider two students who both attend the same environmental philosophy lecture. The lecturer begins by drawing a simple line on the whiteboard to demonstrate just how steep the increase in average global temperatures is. The lecturer has found this more understanding-conducive than sharing complicated statistics or showing an actual graph. Both students at time t_1 accepted an account of the world that includes the drawn line of rising global temperatures and this leads them to the true belief that temperatures are rising quickly. At t_2 , the second student forgets all about the drawing but retains the true belief. My view is that the student who remembers the drawn line could be in a better epistemic position than the student who forgets. We might see evidence of this as the students go on to live their lives and the student who remembers the line, and the reasons it makes salient, is guided to make choices that are better for the environment than the student who forgets, *even if they believe the same things about the environment*. My explanation for this is that the drawing makes reasons salient to the student who remembers via exemplification and its success in doing so is separable from how much truth they believe.

This result accords with the observation in Section 4.1 that accounts of the world may deviate from the truth in a way that enables reason-responsiveness. By deviating from the truth, an account can selectively emphasise what is important and de-emphasise what is not. In doing so, an account of the world can directly figure in valuable understanding. Crucially, if we want to be able to respond to our reasons, this is the sort of epistemic standing we should aim for as finite beings with limited time, resources and cognitive ability. That is why we often prefer the graph, the approximations and the idealisations over the truth, and why their value extends beyond an instrumental role in producing true belief.

4.4 Truth-Oriented Values

Some think that Elgin's examples do not undermine veritism because the value that falsehoods contribute to our understanding still centrally involves the epistemic value of

truth. Linda Warenski (2021: 2722) suggests that the veritist should admit of a ‘rich array of truth oriented values’, where a truth oriented value is construed as bearing an *essential* relation to truth. So perhaps the veritists can claim that, while exemplification complicates matters, understanding is still epistemically evaluated in terms of truth and deviations from the truth can be explained in terms of non-epistemic values. This section argues against versions of this view.

4.4.1 Approximate Truth

One argument is that some falsehoods are of epistemic value because they are approximately true. This is distinct from the argument in Section 4.2.1 about ‘approximately, p’ beliefs, because the agent only needs to accept an account of the world that is approximately true. They need not have a belief of the form ‘approximately, p’, or take any stance on the approximate truth of the account they accept. For example, the veritist might say about the example of the graph that the curve reveals a pattern in the data because it approximates the data (Warenski, 2021: 2724). The idea here is that generalisations are of epistemic value when they are approximately true for cases of interest (Warenski, 2021: 2724). The objection is that if epistemic evaluations of falsehoods in understanding are about approximate truth, then these evaluations still bear an essential relationship to the epistemic value of truth.

One immediate worry is that truth is not central to this evaluation after all, because accuracy is not always better. Where approximations are evaluated in terms of how close they get to the truth, truth is clearly central to this evaluation (Le Bihan, 2021: 5418-5419). But, some approximate truths are not evaluated in terms of their proximity to the truth, but because they enable responsiveness to reasons, which threatens the claim that truth is necessary and central to their epistemic evaluation. To deal with cases like this, the veritist can say that where inaccuracy is preferable, this is due to non-epistemic considerations. This preserves a veritist sentiment: from the epistemic perspective, being closer to the truth is always preferable, but there can be non-epistemic considerations that make approximate truth better overall.

In other words, the veritist must explain the degree of acceptable inaccuracy by appeal to *practical* (or perhaps moral) concerns. The veritist then says that truth is always necessary

and central from the *epistemic* perspective. This move is important because veritists can and do acceptably relativise the epistemic value of truth to practical matters, e.g. the corresponding goals to believe truth and avoid error tend to be relativised to matters of interest (Alston, 2005: 29; Goldman, 1999: 88-89). The veritist can also frame this point in terms of reason-responsiveness; an account that is approximately true can have epistemic value when approximate truth is enough to preserve practical reasons, in a given context. The veritist says that an account can permissibly deviate from the truth if, by doing so, practical reasons are made salient. This ensures that what is left out of approximately true representations will not be practically significant to the epistemic agent that accepts or believes them and therefore truth is still of central concern to *epistemic* evaluation.

While it is often valuable to exemplify properties that preserve the force of practical reasons, it is also possible to deviate from the truth in a way that exemplifies properties of the world that are interesting from the purely epistemic perspective. As I argued in Chapter 1 Section 1.4, epistemic value arises from cognitive contact with normative reasons, including purely *epistemic* reasons to take an interest in some topics over others, to feel one's curiosity is satisfied and to cherish our understanding of more interesting topics. Imagine a curious beginner who is rightly interested in some area of science for its own sake. It is better for the curious beginner to accept an account of topics in this area that is only approximately true, but that better captures these epistemic reasons to feel curious or to feel that their curiosity is satisfied, than an account that is more accurate and discourages their interest. What the curious beginner needs is an account that makes these purely epistemic reasons salient. This is not simply a matter of practical constraints, e.g. time constraints, because the beginner may have time and capacity to improve the accuracy of their account of the world, where doing so would fail to capture their purely epistemic reasons (e.g. putting too much focus on the wrong things, in a way that is boring or confusing). What determines the range of appropriate representations in this case is whether interesting parts of the world are made salient and the normative source of the deviation from the truth does not reduce to truth relative to practical considerations. It is sometimes *epistemically* better to deviate from the truth, because doing so makes what is interesting salient.

4.4.2 True in a Model

Another argument is that falsehoods in idealisations and thought experiments have epistemic value due to claims that are true (or approximately true) within the model. Warenski writes that these accounts are ‘tools for investigating a particular topic that is true in a model...the veritist does not commit to believing something false’ (2021: 2726). We can use what is true in the model as an analogy to reason about the actual world. Thought experiments can also help us reason about the actual world (2021: 2727). With *Last Human*, we consider a false scenario and then reason about whether the value of nature is independent from human interests. This objection is distinct from the arguments in Section 4.2 about belief in counterfactual truths; the idea is not that the value of these falsehoods is exhausted by a true belief about counterfactual truths, but that the epistemic evaluation of idealisations and thought experiments in understanding bears an essential relationship to truth, because they are evaluated in terms of being true within a practically useful model or false scenario.

The problem with this objection is that anything could be ‘true in a model’. We can make many true statements about a relationship between a falsehood and what follows from it. What is interesting about idealisations and thought experiments is that these ‘truths in a model’ reveal something about the actual world. In order for Warenski’s argument to be effective, whether truths in the model are the right sort to figure in understanding has to be a matter of *non-epistemic values*, such as being true in a practically useful model. Otherwise, it is hard to see how truth *essentially* explains why these accounts are valuable from the epistemic perspective.

If the epistemic evaluation of truths in a model is relative to non-epistemic considerations in this way, then the same argument from Section 4.3.1 applies again. It is not always practical considerations that determine which inaccuracies are appropriate. What matters is whether certain reason-giving features of the world are made salient and sometimes this can be a matter of selectively emphasising purely epistemic reasons, such as what is interesting from the epistemic perspective. In a case where the agent considers an idealisation or a thought experiment to understand a part of the world for its own sake, the account should exemplify properties that emphasise what is interesting, even if this requires deviating from the truth. Our epistemic evaluations of understandings that involve

idealisations and thought experiments are not essentially a matter of practically useful truth, because it can sometimes be epistemically better to accept falsehoods.

4.4.3 'Something True'

Finally, there is a sense in which exemplification must require truth. There is 'something true' about false accounts that provide understanding, otherwise they would not reveal anything about the world. This is also part of the definition of exemplification; in order to exemplify some property, e.g. a causal pattern, an account must instantiate that property and refer to it in a way that implies it exists in the world. Moreover, there must be something true about an account that delivers understanding on my view because, in order to successfully grasp, one's representation of the world must preserve the reason-giving force of relevant objective facts and this requires some degree of accuracy about these facts. Understanding involves truth in this sense, but this claim cannot be used to defend an explanatory veritist thesis about epistemic value.

When we examine what 'something true' means here, the idea is just that any epistemically valuable understanding must involve representations of the world that refer to something that obtains in the world. In the context of my approach to grasp, this will be a reason-giving property. However, we shall see that this permissive notion of 'something true' does not adequately explain the epistemic value of either knowledge or understanding. It is unclear how the veritist can maintain that our epistemic evaluations centre around this minimal sense of 'something true', if this notion of 'something true' cannot explain what is epistemically valuable about either of these important epistemic states.

The notion of 'something true' cannot explain the strict true representation required for knowledge. As I detailed in Section 1.2.2, knowledge is thought to be factive, in the sense that it implies true belief. Moreover, belief is non-factive because we can believe falsehoods, but there is still something wrong with believing p if p is false. This is why, from our own perspective, there is something wrong with saying 'I believe p , but not p ' and we are subject to rational requirements that forbid certain contradictions, e.g. belief in ' p and not- p '. In other words, belief seems to be held to a standard of correctness, such that a belief is correct when it is true (Shah and Velleman, 2005). If knowledge requires true

belief, and a belief in something only approximately true is false, approximate truths cannot be known facts.⁷²

If these claims about what it takes to believe the truth are right, then there is a strict sense of true representation involved in knowledge that is a much higher bar than the ‘something true’ claim made above about understanding. To understand something via exemplification, what is minimally required is that there is something in the account that is in the world and that the account refers to that thing to imply that it is in the world. A false belief can also do this. Even a false belief in a proposition that ‘elephants do not have trunks’ includes and refers to something that obtains in the world, elephants, in a way that implies it is in the world. However, there is something wrong with this belief because it is false and it does not amount to knowledge either because it is false. The minimal sense in which exemplification requires truth cannot explain these evaluations about belief or knowledge and therefore cannot explain what is good about knowing.

The sense of ‘something true’ at issue in exemplification cannot explain the value of understanding either. In Chapter 3, I argued that understanding is distinguished from knowledge and true belief because it involves a grasping element that is construed in terms of being in an epistemic position to respond to objective normative reasons. We cannot explain the epistemic value of grasping in understanding with the notion of ‘something true’, because accounts of the world that include and refer to something true can be entirely irrelevant to reasons or may fail to selectively emphasise them. So, the minimal sense in which exemplification requires truth does not explain its contribution to understanding either.

The notion of ‘something true’ might be necessary for knowledge and understanding, but it does not explain the epistemic value of either. At this point, it’s hard to see how the notion of ‘something true’ is even among our epistemic aims or concerns. We therefore cannot vindicate veritism by saying that all good epistemic states have ‘something true’ about them, where this simply means they refer to something that is in fact in the world.

⁷² For discussion of knowledge with approximate truth conditions, see Bricker (2025).

Another important point arises here from the distinction between belief and acceptance. Some take the fact that the scientist does not *believe* the idealisation, thought experiment or approximation as evidence that it is not the falsehood in question that has epistemic value, but the true beliefs that might be inferred or otherwise extracted from these falsehoods.⁷³ In contrast, I take the distinction between acceptance and belief to *support* the view that epistemic value is not essentially related to true representation. A belief might aim at truth in the sense that is correct if it is true and incorrect if it is false. By contrast, acceptance is a policy of taking *p* to be the case in one's reasoning and it aims to facilitate some kind of cognitive activity (Cohen, 1995; Elgin, 2017). Acceptance tolerates false contents because it does not aim at truth in the same way that belief does; an acceptance is correct not if it is true, but if it serves its role in our reasoning. An acceptance might be faulty if it fails to fulfil this role in our reasoning, e.g. if it ends up conflicting with the rest of our account or leading us to a false belief. It seems, then, the measure of success for an acceptance is not whether it has true contents, but whether it preserves contextually relevant normative reasons to respond in a certain way. If the distinction between belief and acceptance marks the distinction between an epistemic value tied to truly representing the world and the epistemic value tied to whether we preserve its reason-giving features, then the fact that scientists accept rather than believe in models turns out to count in favour of the view that the measure of understanding is about whether we arrive at reasons, rather than how close we can get to the truth.

4.5 Conclusion

This chapter has argued that we do not need to truly represent the world to be in an epistemic position to respond to normative reasons and that we can arrive at this position by deviating from the truth. I have also argued that, insofar as I am right that one's understanding improves with one's epistemic position to respond to objective normative reasons (as argued in Chapter 3), my arguments vindicate the non-factive view of understanding and strengthen arguments against veritism.

I now recap the last four chapters and make it clear why an adequate epistemology for the environment should take great interest in understanding. Chapter 1 argued that truth is

⁷³ Insa Lawler (2021: 6872) argues that the fact scientists are aware of such falsehoods suggests that they are not part of the content of their understanding.

not the sole fundamental epistemic aim or concern because it cannot explain the value of grasping and that it is not an important epistemic aim or concern, because truth inherits its importance from its contribution to grasping. These issues come to the forefront in developing an adequate epistemology for the environment. Chapters 2 and 3 provide a more concrete approach to grasping, as an epistemic position that enables responsiveness to objective normative reasons. Chapter 4 argues that grasp does not require true representation and does not necessarily improve with accuracy. An important upshot of Chapters 1-4 is that, insofar as grasp is a distinctive feature of understanding, understanding is of more interest to an epistemology for the environment than knowledge.

An adequate epistemology for the environment should focus on the epistemic condition that enables agents to do right by the environment. To resolve the epistemic problem introduced in the Introduction, Section 0.1, this epistemic position must render the complex reality of environmental issues intelligible in a way that facilitates the right response to the world. As promised in the Introduction, I can conclude that grasping is a good candidate for this sort of epistemic position because when we grasp the world, we arrive at a representation of the world that selectively emphasises facts that objectively count in favour of responses. In contrast to knowing and believing the truth, which may be less important and potentially counterproductive in this context, to understand is to make sense of the world, by bringing important reasons to the forefront of our minds.

Chapter 5: Being a Good Environmental Thinker

5.0 Introduction

In previous chapters, I have argued that understanding is of importance to an epistemology for the environment. At this point, one might reasonably complain that constructing an account of grasp and its value doesn't directly help epistemic agents hoping to navigate environmental issues, especially given that we do not necessarily have reflective access to whether we understand. An epistemology for the environment would be amiss if it had nothing to offer those seeking an understanding of environmental issues.

In this chapter, I attempt to resolve this concern, by turning my attention to character traits that make for a good environmental thinker. We might call this Environmental Virtue Epistemology. I motivate this turn towards virtue epistemology and my specific approach. Section 5.1 argues that Environmental Virtue Epistemology can provide guidance to those who want to become a good environmental thinker and may also provide a basis of responsibility for how we think about the environment. Section 5.2 unpacks what makes a virtue epistemic in this context. Section 5.3 clarifies what makes a virtue environmental. Section 5.4 addresses the question of whether being epistemically virtuous qualifies as an epistemic aim or concern and how it relates to grasping. Section 5.5 introduces the plan for Chapter 6 and 7.

5.1 Environmental Virtue Epistemology

Environmental Virtue Ethics approaches ethical questions that arise from our relationship with the environment by focusing on the character traits of the agent (Cafaro and Sandler, 2005; Sandler, 2007; Treanor, 2014). This research programme concerns the character traits that further a proper relation to the natural environment. So far, Environmental Virtue Ethics primarily attends to the moral virtues and has mostly overlooked the epistemic

virtues, even though epistemic virtues are likely to be crucial to a comprehensive and successful environmental ethics.⁷⁴

I have two main reasons for turning my attention to the virtues, both of which stem from the fact that whether understanding is achieved is not always accessible from the agent's perspective and is often out of their control – this is a consequence of endorsing an externalist-leaning account of understanding. Firstly, it could be objected that Chapters 1-4 are not practically relevant for the environmental thinker, because whether they achieve understanding is beyond their control. If so, I have not provided guidance or advice about how to conduct inquiries about environmental issues. Secondly, whether an environmental thinker achieves understanding does not settle questions to do with whether, and in what way, they are responsible for how they think about the environment and properly subject to attitudes such as praise or blame. I now address these points respectively, arguing that progress can be made on these questions by looking into the virtues.

5.1.1 Guidance

In this section, I argue that an adequate epistemology for the environment should provide guidance in our epistemic lives, by theorising about virtues and vices.⁷⁵ Some think that epistemology aims to analyse epistemic concepts like justified belief and knowledge, rather than to provide guiding advice for epistemic agents. As noted by Robert Audi (2013: 129), this historical focus might be part of the reason why it is commonplace to distinguish between meta-ethics and normative ethics, but less commonplace to make an analogous distinction between 'conceptual epistemology' and 'normative epistemology'. However, recently there has been much more interest in epistemology that tackles inquiry and focuses on the agent, such as virtue and vice epistemology.⁷⁶

⁷⁴ For more arguments that epistemic virtues are valuable to Environmental Virtue Ethics, see Sue P. Stafford's (2010) paper on this topic.

⁷⁵ Some think that the notion of action guidance is important to assessing whether we have arrived at the right picture of norms, because it is important that the normative is able to guide us. John Gibbons (2013: 132) claims that genuinely normative reasons must be capable of guiding us. See also Wedgwood (2002b). To be clear, my interest in guidance is motivated by the idea that epistemology should have something to say about how to conduct our epistemic lives. This is distinct from the idea that the *normative* must be action guiding or that the most important aspects of the normative are guiding.

⁷⁶ Roberts and Wood note that virtues have the potential to humanise and deepen the discipline of epistemology, despite the fact virtue epistemology was initially in service of answering old problems.

In my view, epistemology should have something to say about how to become a good thinker and an epistemology for the environment should have something to say about how to become a good environmental thinker. Thus far the guidance produced by Chapter 1-4 is quite limited. I have not said much other than the good environmental thinker seeks understanding or perhaps that the good environmental thinker should become responsive to their objective normative reasons. These recommendations may be correct, but they are too abstract to provide meaningful advice that could have a bearing on the complexity of our intellectual lives. I suggest that individuating and characterising the virtues can provide this meaningful advice.

To see how virtue epistemology can provide meaningful guidance, notice that virtue terms are more descriptive than other normative concepts, such as epistemic duty and epistemic justification, which lack descriptive content.⁷⁷ Virtue terms such as being close-minded, careless, or untrustworthy, are ‘not only normative terms, conveying a negative evaluation, but indicate the way in which the believer acted improperly’ (Zagzebski, 1996: 20). Virtue and vice terms together provide a richer description of the way in which things have gone well or badly and this has potential to guide our intellectual lives.

By characterising and individuating the virtues, we can construct a framework that helps us navigate our intellectual lives.⁷⁸ Here are some reasons to think this, all of which naturally follow from widely held ideas about the virtues. First, many think that being virtuous involves *practical wisdom*, i.e. good judgement required for the exercise of any given virtue. We might think that practical wisdom is improved by individuating virtues from one another, because by identifying the distinct role that each virtue may have in our intellectual lives, we become better able to recognise the circumstances in which each virtue is appropriate. Second, many think that virtues are human excellences that are established through practice over one’s life (e.g. acquisition and revision).⁷⁹ It seems that thinking about the virtues and how they are individuated, when they might be applicable

⁷⁷ G. E. M. Anscombe (1958) makes the point in ‘Modern Moral Philosophy’ that some moral terms indicate blanket wrongness, whereas others have the benefit of a richer description. See also Bernard Williams’ (2006: 129-130) distinction between thick and thin moral concepts.

⁷⁸ Other virtue epistemologists have also expressed this sentiment that getting clear on the nature of epistemic virtues might provide guidance on how to become a good thinker (Roberts and Wood, 2007: 2-3; Baehr, 2011: 26-27).

⁷⁹ Many think virtue is a matter of practice. Baehr writes that while a person’s upbringing or community can impact whether they are virtuous, ‘virtues are to a significant extent a product of their possessor’s repeated choices or actions – choices or actions that are under their possessor’s voluntary control’ (Baehr, 2011: 27).

etc, can provide guidance in one's practice. For example, an agent might think to themselves "remember to be honest", and because honesty describes a specific way to do things right, this encourages and aids practice in a way that "remember to do the right thing" does not. Many also recognise that virtue is acquired through imitation of virtuous people. Individuating and characterising virtues may help us recognise them in others and thereby know who to learn from. Third, individuating and characterising virtue and vice is indirectly guiding insofar as our social environment impacts the development and maintenance of virtuous character and virtue theory can help us bring about virtue-conducive social environments.

5.1.2 Moral Responsibility and Related Attitudes

Following my articulation of the epistemic problem in the Introduction, Section 0.1, we face many epistemic difficulties in our attempts to understand environmental issues. As a result, it is plausible that much environmental harm is a result of misunderstanding or lack in understanding. In my preferred terminology, the objective normative reasons for actions that prevent environmental harms are often out of our reach. Sometimes this is because non-normative facts are difficult to understand, e.g. because this harm occurs at a spatial or temporal distance from us or because we need certain scientific expertise. Other times it is difficult to recognise the moral significance of non-normative facts we already know. The moral questions that arise from environmental issues are themselves complicated, such as intergenerational justice and moral significance of non-human life. Dale Jamieson (2014: 148-151; 2007: 475-477) suggests that the moral challenges raised by climate change deviate from the paradigmatic moral scenario, in which an individual intentionally harms another individual, both individuals and the harm are identifiable and this occurs close together temporally and spatially. This means that issues raised by issues like climate change are outside the realm of common-sense morality.

In light of the epistemic problem, it is plausible that much environmental harm is a result of failing to appreciate normative or non-normative facts and this raises the question of when factual and moral ignorance related to the environment is culpable (Levy and Brownstein, 2021: 10; Robichaud, 2017). In this sub-section, I argue that character virtues play a central role in explaining responsibility for environmental harm and the

appropriateness of a range of reactive attitudes, such as praise, blame, indignation, condemnation, resentment or gratitude.

Virtues can be part of who we are and we might think we are responsible for beliefs, actions and feelings that express who we are. However, the sense in which we are responsible for being virtuous depends on whether responsibility centres agent *control* or what is *revealed* about an agent.⁸⁰ The agent control approach to responsibility is motivated by the idea that, in order to be held responsible for some action or attitude, the agent must have some degree of control over that action or attitude, such that things could have been different. In contrast, the revelatory sense of responsibility roughly says we are ultimately responsible for features of ourselves that reflect something important about who we are (sometimes referred to as our *quality* of will).⁸¹ Therefore, we are also responsible for any action or attitude that reflects who we are, even if we have no control over it.

I briefly explore the role of virtue in these three broad approaches: (i) control responsibility only, (ii) revelatory responsibility only, or (iii) a middle ground between them, perhaps by admitting of different senses of responsibility. For each of these approaches, virtues will be an important part of the story. I argue that, at the very least, there is a minimal sense in which we are responsible for whether we are a good person or not, even if this does not license reactive attitudes such as praise and blame, and the virtues that reflect who we are (such as character traits) are what make us good people. I take this to mean that we should endorse either option (ii) or (iii), but the sense of responsibility in question does not license the full range of reactive attitudes.

If we go for option (i), the virtues can have an explanatory role in how we are responsible for some action or attitude, but only when they stem from our choices, e.g. whether we choose to engage in self-improvement, or who we choose to listen to. However, we do not appear to have full control over whether we are virtuous, as our traits can be influenced by aspects of our moral education and our present environment. Consider those who have become close-minded due to indoctrination. Such agents may not be in control of the acquisition of this vice and they may also be ill-equipped to reverse

⁸⁰ Alexander A. Guerrero (2017: 208) argues that it is useful to attend to this distinction to understand moral responsibility, and its relationship to praise and blame, in cases of epistemic difficulty.

⁸¹ Some examples of this view include Arpaly and Schroeder (2014), Hieronymi (2014) and Strawson (1962).

their indoctrination, precisely because they have acquired the vice of close-mindedness (Battaly, 2019: 25). Another point, that is applicable to a wider range of agents, is that we all have some virtues and lack others because of non-normative, neutral factors of our life experiences (Jones, 1999: 65; Jones and Schroeter, 2012: 223). Some lives and experiences, such as specific professions, might lead us to excel in some areas at the cost of others and this generates another concern that, even when our virtues do stem from choices that we have made, they might be good choices or neutral choices. If there is no possible path of improvement that could lead us to be fully virtuous, many of the choices that have led us to acquire and practice some virtues over others are neutral. I take it that we would only be held responsible for virtues or vices that stem from choices that are not neutral in this way, which may limit the virtues or vices we can be held responsible for even further.

Although an approach to moral responsibility that centres control is intuitive in many ways, one significant disadvantage to this approach is that intentionally acting against one's principles is rare (and tracing an action back to akratic actions is difficult) and so this view is vulnerable to scepticism about moral responsibility (Rosen, 2004). This view could justify a revisionary approach to our moral practice, meaning that we are not responsible for most wrongdoing. Additionally, a control-only approach to responsibility is even more likely to generate scepticism about moral responsibility for environmental harm brought about in ignorance because, as I have already pointed out, it is often difficult to achieve relevant understanding and this makes the voluntary choice to cause environmental harm even more rare.

We avoid these sceptical worries by endorsing a revelatory-only approach, a possible advantage of option (ii). If we are always responsible for virtues and vices that manifest something relevant about who we are, then we will be responsible for environmental harm that is an expression of vice or lack of virtue, even when this harm is brought about in ignorance or when these traits do not stem from morally relevant choices.

However, the revelatory-only approach has some disadvantages. The first is that many reactive attitudes that we expect to follow from holding someone responsible seem inappropriate, unless the agent has control. I am inclined to think that some particularly

negative reactive attitudes are inappropriate, such as the ‘angry attitudes’ described by Susan Wolf as:

emotional attitudes that involve negative feelings toward a person, arising for a belief or impression that the person has behaved badly toward oneself or to a member (or members) of a community about which one cares and which tend to give rise to or perhaps even include a desire to scold or punish the person (2011: 336).

Other attitudes such as disappointment or a more objective and detached sense of blame might be more appropriate (Mason, 2015: 3052-3056). Conversely, perhaps very positive responses are inappropriate responses to good deeds by mistake, such as praise and reward, but gratitude might still be appropriate when someone does something good unintentionally. If we go for option (ii), then we must explain why we are inclined to think some attitudes are inappropriate without control or provide good reasons for why these intuitions are misguided.

It is important to note, however, we can at the same time have a negative reactive attitude towards an agent and recognise the tragedy that we do not always have control of who we are. This is not a tragedy, however, that undermines the fittingness of certain attitudes that imply some sense of responsibility (Hieronymi, 2014: 30). We must also make sure not to understate the fact that an individual’s responsibility for their character is not the whole story. Since unjust circumstances can lead to vicious character, other individuals (or a wider group that they are a part of) can also be responsible for shaping an individual’s character.⁸² Perhaps we can simultaneously recognise injustice, while also maintaining attitudes indicative of responsibility, such as disappointment and blame.

Another possible disadvantage of the revelatory-only approach is that we do not yet have a clear idea of what it is for an action or attitude to manifest something important about who we are in a way that grounds responsibility. I will not discuss in detail what might be required for virtue to reveal something about an agent, as this is too far beyond the aims of this thesis. I won’t attempt to settle the question, for example, of what it takes for a disposition to be *close enough to the agent* to reflect badly on them that will include

⁸² For a discussion of the relationship between injustice and epistemic character, see Fricker (2007) and Kidd (2020).

implicit bias among epistemic vices, but exclude impaired vision (Battaly, 2019). I do agree, however, that this problem needs to be solved and it counts against option (ii) if we cannot solve it.

Finally, option (iii) seeks a middle ground that inherits the advantages of both agent-control and revelatory senses of responsibility, while avoiding the disadvantages. One way to do this might be to admit of both senses of responsibility, but to claim that some reactive attitudes are only fitting for the control sense of responsibility, but not for revelatory responsibility in the absence of agent control. If this approach is right, then it may be possible to avoid sceptical results that follow from the control-only approach and at the same time justify the claim that some reactive attitudes are inappropriate in the absence of control.

I conclude this section with the claim that, even if it turns out that reactive attitudes such as angry blame are inappropriate in the absence of agent control, at the very least there is some limited sense in which we are ‘responsible’ for some virtues, i.e. character traits. We are responsible for them because they reveal something about who we are. I will discuss this in more detail in what follows. Here I claim that, even if reactive attitudes we normally associate with responsibility are inappropriate in many circumstances, there is a sense in which we can evaluate an agent based on who they are, in a way we cannot evaluate an agent due to the quality of their vision, for example.

Despite the potential for environmental-epistemic virtues to guiding advice or explain responsibility for our attitudes to the environment, this topic has been neglected in *Environmental Virtue Ethics*, with some exceptions (Kawall, 2018; Stafford, 2010). I now unpack what makes a character trait an environmental-epistemic virtue.

5.2 Epistemic Virtues

Epistemic virtues are qualities of an agent that bear an important relationship to something of epistemic value. The epistemic value in question is traditionally thought to be true belief or knowledge, but many virtue epistemologists have also taken interest in a wider range of epistemic goods, including understanding (Zagzebski, 1996; Riggs, 2003;

Roberts and Wood, 2007; Baehr, 2011). The kind of relationship that holds between virtues and these epistemic goods is the question that prompts the divide between *reliabilists* and *responsibilists* about epistemic virtue. The reliabilist thinks that epistemic virtues are reliable cognitive faculties, including perception, introspection and memory, and often define knowledge in terms of these virtues (Sosa, 1991; Greco, 1994). The relationship to epistemic goods like truth, on this view, is reliability; epistemic virtues produce true belief in the virtuous agent more often than not.

In contrast, the responsibilist thinks of epistemic virtues as analogous to moral virtues and, as such, character traits that the agent has some degree of responsibility for (Montmarquet, 1992; Zagzebski, 1996; Baehr, 2011).⁸³ Some examples include intellectual humility, open-mindedness, and intellectual autonomy. For the responsibilist, the relationship that the virtuous trait has to epistemic goods must ground a sense in which we are responsible for having the trait. Based on what I have argued in Section 5.1.2, this means that the relationship between trait and epistemic good must at minimum explain how it reveals something good about a person if they are in possession of the trait.

I will not settle the question of how exactly we should conceive of the relation to epistemic goods that is at issue in responsibilist virtue, but my view is that this will be along the lines of Robert Adams' conception of 'being for the good':

being for is meant to cover a lot of territory. There are many ways of being for something. They include: loving it, liking it, respecting it, wanting it, wishing for it, appreciating it, thinking highly of it, speaking in favour of it and otherwise intentionally standing for it symbolically, acting to promote or protect it, and being disposed to do such things (2006: 15-16).

To be for an epistemic good, is about being the sort of person who has or is disposed to have these positive attitudes towards it. Another way to talk about this in the context of epistemic virtues, includes Baehr's *personal worth* view, according to which the basis of personal worth is when a person is 'positively oriented toward or "loves" what is good and is negatively oriented toward or "hates" what is bad' (2011: 97). Another approach, from

⁸³ Lorraine Code (1984: 39) distinguishes her approach from Sosa's virtue reliabilism, calling it 'virtue responsibilism', because having these traits is what it takes to be a responsible thinker.

Alan T. Wilson, is that the epistemically virtuous have a deep motivation for cognitive contact with reality, where deep motivation has the following features: (i) it is sufficiently persistent, rather than fleeting, (ii) sufficiently strong to move an agent to action and (iii) robust, such that it is not easily overcome by other considerations (2017: 1041). As I have already stated, I do not endorse a specific interpretation of what grounds responsibility for character virtues, but the answer seems to be in the realm of these ideas about what it means to be *for* what is epistemically good.

I follow the responsibilist tradition, which is in line with the aims of this chapter for three reasons. First, unlike many reliabilists, I do not attempt to explain epistemic concepts such as knowledge, understanding and justified belief in terms of epistemic virtues. I am interested in mapping out the epistemic virtues in a way that does not commit to the claim that epistemic virtues are the foundation of epistemic normativity or conceptually prior to the evaluation of epistemic states (Roberts and Wood, 2007). This goes against the idea that virtue epistemology involves a reversal of virtue before justified belief akin to putting moral virtue before the right action. In other words, I do not consider epistemic virtues to have a central role in the conceptual analysis of epistemic concepts. Instead, I consider them in their own right.⁸⁴

Second, one of my stated reasons for theorising about virtues is to provide a means to evaluate environmental thinkers, specifically in terms of responsibility for actions and attitudes, and I am therefore interested in character traits that reveal something important about an agent. Following Baehr, reliability conditions would make virtue subject to luck in a way that would prevent virtue from revealing much about a person's character:

whether we actually succeed, say, in bringing about the good consequences or states of affairs at which we aim is often very largely a matter of luck. It often depends on the cooperation of other rational agents and on a congenial environment: factors that are often substantially, if not entirely, beyond our control. But surely it is mistaken to

⁸⁴ We might understand this in terms of the distinction Baehr (2011: 11-12) makes between *Conservative Virtue Epistemology* and *Autonomous Virtue Epistemology*. *Conservative Virtue Epistemology* takes virtues to have either a central or secondary role in solving problems in traditional epistemology, e.g. Zagzebski (1996) argues that the virtues play a central role in an account of knowledge. *Autonomous Virtue Epistemology*, however, considers virtues independently from problems in traditional epistemology, e.g. Roberts and Wood (2007) map individual virtues and do not aim to make progress on traditional epistemological questions.

think that an individual's goodness or badness qua person might depend on whether she is lucky enough to have the cooperation of other agents or her environment (Baehr, 2011: 97).

Following from my discussion in Section 5.1.2, a person does not even have the minimal sense of revelatory responsibility just in virtue of the fact that they either have achieved or failed to reliably achieve understanding and so reliabilist virtues are not relevant to the aims of this chapter.

Third, responsibilists have more guidance to offer than reliabilists. Reliabilists are interested in cognitive faculties that are often not in our control, e.g. memory and perception, and therefore cannot provide us with much guidance. While we do not control everything that influences our character, there are many opportunities to cultivate responsibilist virtues. For example, we can improve responsibilist virtues when we imitate exemplars of virtue, practice virtuous behaviours, reflect on our own virtues and vices, and advocate for virtue-conducive environments.

None of this is to say that there is no place for reliabilist virtue in epistemology, that they are not valuable, or that reliabilist virtues are not real virtues. Some think we do not have to choose between the responsibilist and reliabilist virtues and I have not ruled this out. Heather Battaly (2014: 183), for example, points out that both reliabilist and responsibilist virtues make us excellent thinkers. It is just that, for the purpose of this thesis, it is responsibilist virtue that matters.

5.3 Environmental Virtues

Environmental virtues are character traits that aim at environmental flourishing. We need to investigate environmental virtues because the traditional moral (and epistemic) virtues are not all that matters to a proper relationship to the environment, particularly in the context of present environmental challenges; we might need to add to the list of virtues, modify existing virtues, or eliminate virtues. We have reason to think that traditional epistemic virtues are also subject to this criticism, because environmental issues also have a bearing on our epistemic lives. In accordance with the approach in Section 5.2, I do not

attempt to explain the right action via the environmental virtues, but instead identify and analyse the environmental virtues in their own right.

In a similar vein to the discussion of virtue epistemology, what makes for excellent environmental thinker can vary, because the manner in which environmental considerations are relevant to virtue varies. There are at least three distinct ways a virtue might be environmental, identified by Ronald Sandler (2007: 42-43). *Environmentally responsive* virtues involve responding to environmental entities (e.g. landscapes and ecosystems). Some virtues might not involve responsiveness but are still justified by environmental considerations (e.g. the value of organisms and landscapes). These are *environmentally justified* virtues. *Environmentally productive* virtues promote environmental goods and values (e.g. support flourishing of life). Often a virtue will fit into all three. For example, wonder is characterised by a feeling of awe and wonder seems to be virtuous when we experience a feeling of awe as *responsiveness* to the right things; perhaps we should wonder at sounds of wildlife in a forest, but we should not wonder at city lights and pollution it causes (even if it is aesthetically pleasing). Wonder might also be *justified* by an environmental consideration, such as the value or worth of the environmental entities towards which we should feel awe. Finally, wonder motivates us to protect and seek information about these entities and, therefore might also be *environmentally productive*. It is worth noting that not all environmental virtues obviously emphasise our responses to natural environments because the status of an environmental virtue might be relative to the context of present environmental challenges (Treanor, 2014: 60). We might, for example, take virtues surrounding when we should defer to others (e.g. intellectual autonomy) to be environmental virtues, if they promote environmental goods and values.

I support this pluralistic approach because I am interested in epistemic virtues that are relevant in environmental contexts, rather than a set of distinctively environmental virtues. There might be a set of virtues that are distinctively environmental in the sense that they concern one's 'being for' the environment, but I am interested in epistemic virtues that are especially important in overcoming distinct and particularly profound epistemic difficulties that arise from our relationship to the environment. We should not restrict our attention to virtues that involve 'being for' the environment because agents can possess epistemic

virtues in a way that is important to being a good environmental thinker, even if they are not an expression of ‘being for’ the environment in particular.

So far, I have unpacked what counts as an environmental-epistemic virtue and demonstrated why this set of virtues is particularly worthy of our attention. There are many important tasks for Environmental Virtue Epistemology. As I detail in Section 5.5, Chapter 6 and 7 identify and characterise some individual environmental-epistemic virtues, in order to make progress on the aim of providing guiding advice and a basis of responsibility for how one thinks about the environment. Before this, the next section discusses the status of being epistemically virtuous as an epistemic aim or concern, situating the next few chapters in the context of arguments in Chapters 1-4.

5.4 Being Epistemically Virtuous as an Epistemic Aim or Concern

I now discuss where being epistemically virtuous fits within the structure of the epistemic aims or concerns discussed in Chapter 1. Recall that Chapter 1 argues that truth is not the only fundamental epistemic aim or concern because this view cannot explain the status of grasping as an epistemic aim or concern. I also argued that, if truth is a fundamental epistemic aim or concern, it is not important relative to grasping because where truth is important, this is due to its role in grasp. There are different ways we could make sense of being epistemically virtuous as an epistemic aim or concern: (i) being epistemically virtuous is more fundamental and grasp is derivative of being epistemically virtuous, (ii) grasp is more fundamental and being epistemically virtuous is derivative of grasp, (iii) being epistemically virtuous and grasp are independent epistemic aims or concerns and neither derives its status as an epistemic aim or concern from the other, (iv) being epistemically virtuous is not an epistemic aim or concern because it is non-normative. While I cannot settle this issue here, I make my perspective on the relationship between aspects of this thesis clear by stating my reasons to prefer (iii).⁸⁵

⁸⁵ Note that my arguments in the following section are not meant to be decisive. My aim is to outline some different ways to think about the relationship between grasping and being epistemically virtuous. In order to establish more decisive arguments about this relationship, I would have to explore relevant arguments in much more depth.

Consider first that (i) is true: being virtuous is more fundamental than grasping and grasp derives its status as an epistemic aim or concern from being epistemically virtuous. As we have seen, some argue that virtues explain epistemic standings such as knowledge, but this does not appear to be a good approach in light of my commitments about both grasp and responsibilist virtues. My approach to grasp is externalist in the sense that to grasp is to achieve cognitive contact with facts that are objective normative reasons and I have argued this is of value because it is valuable to appreciate facts that objectively speak in favour of certain responses. On the other hand, I have taken an internalist approach to responsibilist virtues, by arguing that they are not subject to success or reliability conditions. This means that my take on responsibilist virtues will not have the resources to explain what is epistemically valuable about grasp as a relation to objective normative reasons.

If (ii) is true, then things are the other way round and grasp more fundamentally explains being epistemically virtuous, such that the status of being epistemically virtuous as an epistemic aim or concern derives from grasping. There is a story to tell that says that to grasp is to be in cognitive contact with objective normative reasons and to be epistemically virtuous is to strive for cognitive contact with objective normative reasons, and therefore being virtuous is instrumental to achieving grasp or is part of what is required to achieve grasp. However, there are good reasons to question this. First, it's not clear that my view of epistemic virtues makes them instrumentally valuable in virtue of their capacity to bring about grasping. This is because it seems to be nevertheless good in itself to strive to be a good thinker. Even if we fail catastrophically to achieve grasp, it's still good that we are the sort of person who is properly oriented towards what is valuable about grasping. Secondly, the value of being for what is epistemically good is not swamped by the achievement of grasp. We may achieve grasp in a non-virtuous way and we are missing something valuable as a result. So, being epistemically virtuous seems to be an epistemic aim or concern that does not derive from the more fundamental aim of grasping.

If (iii) is correct, then neither grasping nor being epistemically virtuous derives status as an epistemic aim or concern from the other. They may be independently valuable. Perhaps whether we achieve epistemic standings like grasping is a separable normative realm than considerations about whether one is a good thinker. This is my preferred approach, as we then avoid the serious problems raised by (i) and (ii). I will remain neutral on the question

about whether it is grasp or being epistemically virtuous that is of more importance – as established in Chapter 1, Section 1.1 this is a separate issue to which is fundamental.

If (iv) is right, then being virtuous is non-normative. This would be to argue that being virtuous does not fall within the realm of epistemic normativity. Although I'm not sure exactly what this would look like, we might think that being a good thinker is not genuinely normative in the way that grasping and responsiveness to objective normative reasons is. Perhaps normative reality consists of objective normative reasons and being virtuous is just trying your best, or perhaps being virtuous is being for our epistemic aims or concerns and it is incoherent to count being virtuous as *among* our epistemic aims and concerns. I am, however, sceptical of this approach because it seems to me that it matters from the epistemic perspective that we are the sort of person who stands for whatever is of epistemic value in the ways I have discussed.

5.5 Plan for Chapter 6 and 7

In this chapter, I have argued that it is an important part of Environmental Virtue Epistemology to individuate and characterise the environmental-epistemic virtues. In the following chapters, I provide two demonstrations of Environmental Virtue Epistemology. Both virtues discussed in what follows are supposed to be responsive to the epistemic problem faced by the environmental thinker: the need to process vast amounts of complex information. Given that the first few chapters of this thesis argue that this epistemic problem suggests the particular importance of understanding, we might suspect that something has gone wrong if the discussions in Chapters 6 and 7 do not reference understanding. In the following chapters, we see how a focus on understanding (and my particular account of grasp) influences how we should think about these environmental-epistemic virtues.

As argued in this chapter, we might expect to find either *new* virtues or *revise* familiar virtues through the environmental lens. Chapter 6 discusses a 'new' virtue that appears to become particularly relevant in this context. In order to process vast amounts of complex information about the environment, the environmental thinker must see the bigger picture. In Chapter 6, I develop an account of outward-lookingness and its role in big picture

thinking. Chapter 7 concerns the familiar virtue of honesty. I argue in Chapter 7 that honesty is important to avoid deceiving oneself, or others, about environmental issues and the epistemic problem of processing vast amounts of complex information demands revisions to how we think of honesty as an environmental virtue. By discussing these virtues, I provide an example of both the development of new virtues (where the task is primarily to individuate and characterise a distinct virtue and its contributions to guidance and responsibility) and the revision of familiar virtues (where the task is primarily to recognise how environmental demands might influence debates about the nature of familiar epistemic virtues).

Another reason to focus on these virtues in particular is that they cover different aspects of our epistemic lives. While outward-lookingness primarily concerns seeking new knowledge and understanding, honesty primarily concerns retaining existing knowledge and understanding. By discussing virtues that differ in this respect, this thesis covers a wider expanse of the epistemic lives of environmental thinkers.

5.6 Conclusion

In this chapter, I have explained and motivated a specific approach to Environmental Virtue Epistemology that may provide guidance to the environmental thinker and a means to evaluate an agent's efforts to be a good environmental thinker, efforts for which one may be responsible. In the next chapter, I apply this approach to virtues that contribute to big picture thinking.

Chapter 6: Outward-Lookingness and Seeing the Bigger Picture

6.0 Introduction

To be able to ‘see the bigger picture’ seems to be an important part of being a good environmental thinker and we should take interest in characterising virtues that contribute to this aim. In Section 6.1, I clarify the role of outward-lookingness in seeing the bigger picture. In Section 6.2, I distinguish this trait from other environmental-epistemic virtues that contribute to big picture thinking: curiosity, open-mindedness, far-sightedness and holistic thinking. In Section 6.3, I present a problem that might lead us to question whether character traits that contribute to big picture thinking are environmental virtues and I explore some potential solutions.

6.1 Outward-Lookingness

In Chapter 5, I argued that individuating and characterising the environmental-epistemic virtues is an important task for epistemology. It’s helpful to begin by thinking about the challenges we experience in our inquiries about the environment and the traits that can help us overcome them. One challenge for the environmental thinker is my focus in this thesis. To understand the environment is to grapple with a vast amount of complex information and, because the environmental impact of our actions extends far beyond our immediate surroundings, this information is not obvious to us from simply living our lives. It is not usually possible for us to come to know and understand environmental impact passively. We must actively seek to ‘see the bigger picture’. If we are unable to see the bigger picture, we will fail to recognise how we contribute to environmental harm.

Relatedly, the account of grasp defended in Chapters 1-4 says that degree of understanding relates to the importance of normative reasons, but evaluations about degrees of understanding is relative to an *object* of understanding that is already

demarcated.⁸⁶ As a result, degrees of understanding do not necessarily reflect whether we have arrived at our most important objective normative reasons. We can have a high degree of understanding that only involves the unimportant reasons, if the object of that understanding only presents unimportant reasons to begin with, such as the arrangement of raindrops on my window. In these cases, the most important normative reasons that we arrive at, I have argued, are normative reasons to feel bored or direct our attention elsewhere. Some topics, objects or explanations are more worthy of our understanding than others, due to being morally and practically relevant or more interesting from a purely epistemic perspective. This raises the issue of what we should seek an understanding of, which cannot be settled by understanding attributions alone (since I have construed understanding attributions as relative to an already stipulated topic, object, explanation etc.). It seems to me, that good environmental thinkers will seek understandings that target the bigger picture and therefore capture the deeply interconnected and potentially spatially and temporally vast nature of environmental issues.

Many virtues plausibly contribute to seeing the bigger picture, e.g. curiosity, holistic thinking, open-mindedness, and far-sightedness. Due to the limited discussion of these virtues, as both epistemic virtues and environmental virtues, their contributions to seeing the bigger picture are at risk of being diluted or misunderstood. In effort to individuate and characterise an important set of environmental-epistemic virtues, I first theorise about outward-lookingness and then, in the following section, I distinguish outward-lookingness from other virtues that contribute to seeing the bigger picture.

It helps to start with some examples for how we might think of ‘being outward-looking’ as virtuous, particularly in the sense of being a good environmental thinker.⁸⁷ Notice that to look outward is essentially relative to the position of the agent looking outward. Roughly, to be outward-looking is to look beyond oneself and one’s own life. One sense of being outward-looking concerns what is spatially ‘outward’ from our location. This is environmentally virtuous in cases where the environmental impact of an agent’s actions disproportionately affects parts of the world spatially distant from their own.

⁸⁶ See Chapter 3, Section 3.3.1 in particular.

⁸⁷ The methodology for identifying and distinguishing the virtues is to reflect on virtue terms in ordinary language (Roberts and Wood, 2007: 81).

One can also be outward-looking in the sense of looking outward from the present point in time. This can include recognising the long-term consequences of our actions, such as future impact of present-day emissions. However, it is not just considerations about the future that matter to being a good environmental thinker. Simon P. James (2015: 78) points out that the environmentally virtuous person will be humbled by natural historical timescales and will realise how momentous it is when a species goes extinct, thus ending the evolutionary history that produced it. In order to exercise (or even activate) this humility, the environmentally virtuous person must properly consider what it is in the past. It seems then that we are not just interested in balancing the present against the future (as we shall see, this is what matters to far-sightedness), but generally looking outwards from our own point in time. Again, this appears to be about looking beyond oneself and one's own life.

The above cases concern looking outward from one's spatial or temporal location, but looking outward can also describe looking beyond what is not obvious from one's surroundings in another sense. Consider the agent who lives in a town with a large green space that is home to wildlife. We might think that the agent who extends their attention beyond the built environments in which they live their daily lives, and is therefore aware of why we should oppose building on this green space, is outward-looking in some sense. This is true even though these species are close by and subject to a present threat. We might think that this is a case of being outward-looking because the agent actively extends their attention outside what is immediately obvious to them from just living their life.

Being outward-looking centres around an agent's ability to look beyond themselves and, as a result, extend their inquiries beyond what is immediately obvious from their surroundings. Being outward-looking seems to resist a "what I see is all there is" bias that stems from overemphasising or giving undue priority to our own perspective. In each case presented above, this is environmentally virtuous because environmental considerations are not likely to be present or obvious from our immediate surroundings.

To develop this thought more, we can conceive of outward-lookingness as a balance between two extremes. On one extreme, being outward-looking contrasts with being too inward-looking; when we are viciously inward-looking, we focus too much on ourselves and our lives, such that we unjustifiably neglect everything else. On the other extreme, we

could be viciously outward-looking when we consider what is beyond ourselves and our immediate surroundings to the extent that we neglect ourselves or what is obvious in our immediate surroundings, e.g. we are entirely focussed on the long term and fail to recognise what will happen in our own life tomorrow. It is important to recognise the excess of outward-lookingness. Otherwise, it is unfair to those who are experiencing times when their attention ought to be on themselves and those around them.

Outward-lookingness has an epistemic dimension because it involves actively extending our inquiries beyond ourselves and our lives. Following my arguments in Chapter 5, Section 5.2, outward-lookingness is a responsibilist epistemic virtue where the outward-looking agent is properly oriented toward what is of epistemic value, like knowledge and understanding. It's easy to conceive of outward-lookingness in this way, where an agent extends their inquiries beyond what is most relevant to their own lives as a result of 'being for' the epistemic good, such as a love of understanding or motivation to acquire understanding.⁸⁸

6.2 Other Virtues

This section distinguishes outward-lookingness from other virtues that are closely related, in that they also aim at bigger picture thinking. Epistemic virtues have different functions or roles in our epistemic lives because they help us overcome different challenges that we might face, and the environmental-epistemic virtues help us to overcome the challenges that are particularly profound in environmental contexts. I assume then that an epistemic virtue is distinct from other virtues if it plays a different role in our epistemic life, e.g. by helping us to overcome a different challenge or helping us overcome the same challenge in a different way.⁸⁹ It is important to individuate the virtues because doing so provides guidance in our epistemic lives and our evaluations related to responsibility. My aim is to defend my claim that outward-lookingness is a distinct virtue against some possible objections and, in the process, further clarify its nature.

⁸⁸ Outward-lookingness may also be a moral virtue where it manifests being for what is of moral value.

⁸⁹ Baehr (2011: 18) groups virtues in terms of how they help us meet the demands of inquiry.

6.2.1 Curiosity

Curiosity as an epistemic virtue shares some important qualities with outward-lookingness as an epistemic virtue.⁹⁰ Firstly, curiosity appears to be to do with motivating inquiry, including motivating us to learn about what is beyond our own lives. Curiosity has an affective component that prompts the question “what else is there beyond what I already know and understand?” and motivates us to look outward. Secondly, curiosity can be spontaneous, in the sense that it is independent of our prior interests, and this can lead us to extend our attention to what is not already directly relevant to ourselves or our lives. Thirdly, being virtuously curious is important to being a good environmental thinker for similar reasons that being outward-looking is. Being curious about the natural world will also lead us to actively acquire more knowledge and understanding about the environment that we are unlikely to acquire passively. It could be argued that outward-lookingness follows from what it takes to be a virtuously curious person and, therefore, that being outward-looking reduces to being virtuously curious. However, when we inspect the role of curiosity in one’s epistemic lives, and how it contributes to seeing the bigger picture, there are good reasons to think curiosity and outward-lookingness are distinct environmental-epistemic virtues.

Notice that curiosity involves an emotional or desire-like quality, a feeling of wanting to know or understand about a topic or question that is satisfied when we reach these epistemic states. Being disposed to experience this ‘state’ of curiosity can be virtuous because these emotions and desires have a normative dimension: there is a proper way to experience them. This is about having this feeling of wanting to know about the right things and in the right way (Ross, 2020). If it turns out that being virtuously curious centrally concerns asking the right questions, then there is an argument that outward-lookingness collapses into virtuous curiosity; being curious about the right questions will require that we extend inquiries beyond our own lives. However, there are good reasons to reject this approach to curiosity.

Say that being virtuously curious involves ‘being for’ (or being positively oriented towards) asking the right questions in this broad sense. This means that curiosity is then separated from its distinctive feeling of interest or wanting to know the answer. It seems to

⁹⁰ Others think of curiosity as an epistemic virtue (Miscevic, 2015; Ross, 2020).

me that sometimes we are not any less virtuous epistemically because we do not experience a feeling of wanting to know the answer to the ‘right’ question. For example, sometimes we do not experience curiosity because we are unaware about any specific gaps in our understanding or because the right question is whether we need to go to the shop to get milk and this does not inspire any feeling of interest or desire to know. A more plausible approach to virtuous curiosity, in my view, is that being virtuously curious concerns having an interest in what is inherently interesting, such that the virtue of curiosity regulates the feelings of interest that may motivate inquiry. We can be viciously curious about the wrong things, i.e. uninteresting things, or we can be viciously deficient in curiosity when we fail to want to understand the right things, i.e. the interesting. On this view, being curiosity is not centrally about asking the right questions as such. Instead, being virtuously curious concerns being interested in the interesting.

If this view on curiosity is right, then curiosity and outward-lookingness might help us overcome the challenge of navigating vast amounts of information and the need to see the bigger picture, but these traits offer different remedies. Curiosity regulates whether we have appropriate feelings of interest towards what we should properly find interesting and can therefore help us when what prevents us from seeing the bigger picture is *a lack of interest in what we should properly take interest in*. Outward-lookingness regulates when and the extent to which we extend our inquiries beyond our immediate surroundings and therefore will help when what prevents us from seeing the bigger picture is *a fixation on our own lives*. We can see that being outward-looking does not always follow from curiosity, because there might be infinite interesting aspects about ourselves and our own lives that are appropriate objects of our curiosity. If this is the case, we won’t find the answer to how far we should extend our inquiries outwards by being curious about what is interesting.

Perhaps this argument is unconvincing and I have not considered all the reasons to think of curiosity as asking the right questions. It is worth noting that if outward-lookingness ends up being a form of virtuous curiosity, the other arguments in this chapter still stand.

6.2.2 Open-Mindedness

Open-mindedness is also closely related to outward-lookingness. Most notably, there is a sense in which both of these traits involve being receptive to other perspectives and resisting the temptation to prioritise one's own. We might suspect that being outward-looking is just a form of open-mindedness, where the alternative perspectives in question are those that extend beyond our own lives in the way I have described.

Open-mindedness, however, centres confidence in our own convictions and openness to evidence that contradicts our held views. For example, Nomy Arpaly (2011) argues that the moral concern of the open-minded person prevents them from resisting evidence that is contrary to their belief. On this view, someone might have the epistemic virtue of open-mindedness where they are prevented from resisting evidence due to epistemic concerns, such as a desire to reach the truth. We can also conceive of an excess of this trait, where we do not resist evidence when it is appropriate to do so (King, 2021: 209). To be indiscriminate about what evidence to entertain can waste time or distort reality with obviously false points of view.

Even on Baehr's (2011) view of open-mindedness, which states that to be open-minded does not necessarily involve disagreement or rational assessment, what is centrally at issue is that there is a default standpoint that we must detach from cognitively or transcend in some way to take an alternative standpoint seriously. Consider Baehr's cases where a teacher attempts to push their students from a basic understanding of Einstein's Special Theory of Relativity to an understanding of Einstein's General theory, which will require moving beyond their usual approach. The teacher requires their students to open their minds in the sense that they must be willing to let go of some of their ways of thinking (Baehr, 2011: 145-146). This does not require any disagreement between the two theories in question (one is an extension of the other) and neither does it involve rational assessment, insofar as the students are not aiming to assess the truth of this theory, only to follow and understand it. In general agreement with Baehr, Wayne Riggs writes that being open-minded requires the 'flexibility of mind and character to "talk seriously" some point of view different from one's own' (2016: 22). Even when open-mindedness is not a matter of weighing evidence or assessing the truth of beliefs, it still centrally involves letting go of

previous or default ways of thinking for the purpose of properly considering an alternative view.

While open-mindedness can clearly help us to see the bigger picture, this only applies in case of conflict with our current ways of thinking. In fact, open-mindedness is only relevant when this conflict is the pertinent challenge that needs to be overcome. However, there need not be a particularly challenging conflict between the cognitive standpoint we have to aspects of our own life and what extends beyond it; it's therefore not always necessary for an agent to detach from their own life in order to extend their inquiries in this way. These virtues contribute to big picture thinking in distinct ways; being open-minded helps when there is conflict between our own viewpoint and seeing the bigger picture, whereas outward-lookingness helps resist a tendency to focus on our own lives.

It's important not to understate how much open-mindedness enables outward-lookingness. We can imagine that we must often challenge assumptions about what is important in order to be outward-looking and this requires the sort of transcending and detaching at issue in open-mindedness. However, the fact that one virtue often enables another is not enough to make one virtue reducible to the other or to undermine the importance of individuating them in order to better understand what it means to be a good environmental thinker.

So, open-mindedness is about regulating confidence in the truth of our own beliefs or the priority we afford default commitments that figure in our understandings. What is at stake is the proper consideration of a viewpoint that contradicts some of your existing beliefs or features of your present understanding. By contrast, what is at stake with outward-lookingness is to what extent we look beyond ourselves. This is not so much about resolving a tension that arises from a default assumption that we are right, but instead avoiding the tendency to over-emphasise information that is more readily available.

6.2.3 Far-Sightedness

Another putative virtue that contributes to big picture thinking is far-sightedness, where the agent properly considers what is far away and does not over-emphasise the near. If far-sightedness is a virtue, it operates differently to the outward-lookingness and can therefore

be distinguished. I admit, however, that far-sightedness could overlap significantly with outward-lookingness.

If far-sightedness is a virtue that contributes to big picture thinking, I imagine it would look something like this. The far-sighted agent does not overemphasise what is near (or perceived to be near) and therefore properly considers what is far away. This might play a crucial role for the environmental thinker in their search for understanding of ecosystems, where something far away can play a fundamental role in understanding something nearby. For example, the moon has an impact on life on Earth and those who wrongly disregard potential influence of distant entities, like the moon, might disbelieve or fail to recognise such connections (even if they are a curious, open-minded, holistic thinker, perhaps).

From my discussion in Section 6.1, outward-lookingness seems to demand far-sightedness because looking beyond ourselves in a virtuous way requires that we consider what is spatially or temporally distant. Far-sightedness might end up with significant overlap with outward-lookingness for the environmental thinker. This result is acceptable, unless the role of outward-lookingness is entirely subsumed by and better explained by the virtue of far-sightedness. I now present some reasons why this is not the case.

Firstly, I have some sceptical concerns about whether far-sightedness is virtue. One of my worries is how far-sightedness can manage spatial and temporal distance as if they are the same thing. Outward-lookingness avoids this because it concerns looking beyond one's own life and these considerations will drive us to look at temporally and spatially distant considerations that are not present in our immediate surroundings. However, it is not clear to me that our dispositions to look far away spatially and to look far away temporally will be the same.

Secondly, far-sightedness in the spatial and temporal sense does not deal with the case of wildlife in one's own town or, for example, events that occur on a very small scale. There is another sense of 'distance' here that is not strictly spatial and certainly is not temporal if these events occur in present times. The sense of distance seems to be to do with distance from what we perceive to be relevant and how difficult it is to recognise something passively from our immediate surroundings. I propose that outward-lookingness neatly encapsulates all these considerations in a way that does rest on metaphors with respect to

distance, but instead rests on looking beyond ourselves, including our spatial, temporal position, but also perhaps the scale that we can perceive, without aids like telescopes and microscopes, as well as what we perceive to be directly relevant to how we live our lives, such as what is emphasised by the built environments or social conventions we live with.

6.2.4 Holistic Thinking

An environmental virtue that has many similarities with outward-lookingness is holistic thinking (Treanor, 2014: 60). Holistic thinking is about recognising interconnectedness and properly considering the environmental whole, rather than thinking that parts are separable from the whole. Both virtues aim at the big picture thinking required for the agent to have an awareness of the wider structure of interconnections that they are situated within. The deficiency of both virtues results in a narrowness of awareness and concern. This is a reason to think that these virtues are environmental.

The difference between these virtues becomes clear when we consider the two extremes that they regulate. Unlike outward-lookingness, holistic thinking is about a balance between proportionate consideration of parts and wholes. On one hand, we are deficient in holistic thinking when we abstract parts from wholes. Brian Treanor (2014: 60) suggests that an example of this would be viewing the world as a collection of resources for us to exploit without consequence. In taking this view, we fail to recognise interconnectedness. On the other hand, we have an excess of holistic thinking when we prioritise the whole and only consider individual parts, e.g. individual persons, instrumentally to the whole – a kind of ‘ecological utilitarianism’. On some readings, this seems like more of a moral virtue, that balances moral consideration for individual parts and wholes, but there is an epistemic reading that concerns recognising interconnectedness for the sake of knowledge and understanding. Either way, holistic thinking is about the appropriate recognition of interconnectedness, such that we do not neglect the interdependence of parts and we do not fail to consider them apart from their service to an environmental whole.

In contrast, we have seen that being outward-looking is not about finding the balance of recognising interconnectedness. Too little outward-lookingness therefore does not tend to manifest as abstracting parts from wholes or viewing the world as a separable collection of

resources and too much outward-lookingness does not tend to manifest in the ‘ecological utilitarianism’ described above. Given that these virtues strive for a different kind of balance, we can imagine an agent that has each virtue to a different degree. For example, there could be an agent that is appropriately outward-looking, but is deficient in holistic thinking, such that the scope of their inquiries is far-reaching, but they tend to view individual entities as separate from the environmental whole.

From here we could launch an objection against outward-lookingness as an environmental virtue. Outward-lookingness centres a distinction between the self and what is beyond the self, whereas holistic thinking centres interconnectedness. Given that interconnectedness is crucial to an understanding of ecological structure and threat to ecological stability, it could be objected that holistic thinking is doing all the work here. The objection says that the agent who is outward-looking, but falls short of holistic thinking, does not achieve environmental virtue at all. By looking outward, those who fail to appreciate interconnectedness merely expand the scope of their knowledge and understanding of what they perceive to be resources to exploit. It could even be argued that viewing oneself as ‘distant’ from the environmental impact of our actions runs counter to the development of proper holistic thinking.

Although I accept that the outward-looking agent is limited in how environmentally virtuous they can be if they lack other environmental virtues, like holistic thinking, I maintain that being outward-looking has a distinct function that cannot be reduced to holistic thinking. Even if we view ourselves as completely interconnected, such that all concept of a self as distinguished from the environmental whole is irrelevant, we can still recognise a distance between one physical location and another, or one point in time and another, and identify with one of either. In this sense we can still be outward-looking without recognising a straightforward distinction between the self and the rest of nature.

There is also something missing from holistic thinking that outward-lookingness can contribute. Firstly, holistic thinking carries the commitment that everything relevant to environmental flourishing can in fact be conceived of in terms of parts and wholes. Whereas outward-lookingness does not carry this commitment and could potentially extend to environmental considerations that cannot be conceived in this way, if any did arise. Secondly, even if *all* environmental considerations can be conceived of in terms of

parts and wholes, it is not plausible that an agent is always in a position to appreciate these interconnections or that it would be right to do so. For example, on some occasions the salient feature of a particular action that makes it virtuous is that it came from a place of indifference about how this action affects or otherwise relates to oneself, rather than an attempt to trace interconnections that would reveal some relevance to oneself. It seems to be virtuous in its own way to care about environmental issues regardless of any identified or assumed impact they might have on one's own life. If this is right, there is room for outward-lookingness even though it is often enabled or enhanced by holistic thinking.

To summarise this section, I have distinguished outward-lookingness from some other virtues that also contribute to big picture thinking. In the following section, I present an objection that threatens the status of all these traits as environmental virtues.

6.3 The Intelligibility Problem

In this section, I introduce an objection to the claim that the virtues discussed in this chapter are environmental virtues due to their contributions to big picture thinking. The *intelligibility problem* is that we face many barriers when we attempt to acquire knowledge and understanding about the environment, particularly when we try to expand our attention to the great volume of complex information available, and this calls into question whether virtues that aim at capturing the 'bigger picture' can ever be effective. In fact, we might worry that it is counter-productive to aim for the bigger picture because doing so will only make things unintelligible. After I set up this problem, I explore some solutions.

6.3.1 The Intelligibility Problem

There are several aspects of human epistemic agents that frequently come into conflict with processing the sheer volume, complexity and weight of information most relevant to our grasp of environmental challenges (Kawall, 2009). Humans are finite beings and most of us are constantly managing numerous responsibilities to ourselves and others. On the other hand, paying attention to and understanding environmental issues is difficult. It requires that we dedicate a certain amount of the limited time we have in a day researching, reading news stories, and keeping up to date with what is happening in the

world. Because environmental changes are not noticeable in our immediate surroundings, this will not be a passive affair and will require actively seeking this information.

In actively seeking and processing this information, we are subject to significant *time constraints*. Some have more time to dedicate to environmental issues than others. We are also subject to *cognitive constraints* of various kinds. We will be to some extent susceptible to psychological biases that interfere with our understanding, we can only remember so much, and we only have a certain amount of background knowledge etc. Finally, we are also subject to *emotional constraints* that are separable from cognitive constraints. Paying too much attention to certain content, e.g. disturbing news stories, visual imagery or even our own imperfections, can provide us with brilliant understanding of states of affairs, but may cause emotional harm that is counterproductive or unfair.

Considering the vast amount of complex information about the environment and the various constraints faced by the environmental thinker, the intelligibility problem suggests that big picture thinking virtues like outward-lookingness are not the environmental virtues that I have presented them as. By aiming for the bigger picture, we only render the world unintelligible. With respect to my specific aims in this chapter, advice such as ‘be outward looking’ is, for many, not very guiding and a failure of to be outward-looking rarely explains responsibility for actions or beliefs, e.g. for those who do not have free time due to their circumstances, such as employment or caring responsibilities. This suggests that individuating big picture thinking virtues is not very helpful to the aims of this chapter.

Before I address the intelligibility problem, I want to acknowledge that virtues can be masked by circumstance and failing to arrive at understanding does not always imply absence of virtue (Kawall, 2009: 123). Another important point is that responsibility for the presence of a trait (or the ease with which it can be exercised) does not solely fall on individuals. Changes in social and political environments can help to overcome these epistemic difficulties, such that we can better understand environmental issues. Some suggestions by Kawall (2009: 127) that might be particularly relevant here are transparency for consumer goods and shift in media attention. But even taking this into account, big picture thinking virtues are threatened by the intelligibility problem because the problem appears to arise from a fundamental tension between the constraints of human

agents and the scale of environmental problems. This issue arises even in the face of efforts to mitigate these difficulties.

There are different ways to respond to the intelligibility problem. One option is to take more of an ideal view of the virtues (or, more accurately, their capacity to produce understanding), such that it's okay if the target of big picture thinking is unattainable for human beings. This does not solve the issue, however, because the objection is not that seeing the big picture is unattainable; an unattainable standard can still be guiding, if by aiming for it, we could achieve better than we would have otherwise. The objection instead follows the form of Robin McKenna's (2023) argument for pursuing non-ideal epistemology. It might be that the non-ideal agent is, in fact, *worse off* epistemically if they strive for an ideal sense of big picture thinking. The issue is not that aiming to extend one's inquiries to vast amounts of complex information is unattainable, but that this attempt to see the bigger picture actually renders environmental issues unintelligible.

A better option is to accept that the dispositions and motivations associated with non-ideal big picture thinking are not the same as the dispositions and motivations of the non-ideal agent who strives for ideal big picture thinking. This approach will accommodate barriers to understanding without losing sight of what makes big picture thinking virtues important environmental-epistemic virtues in the first place. In what remains of this chapter, I attempt to present some ideas about how this might work.

6.3.2 Disposition to Make Intelligible

Some ways of presenting information are more effective than others. In general, we can better understand the nature of a phenomenon if we examine a particular instance of it, e.g. an image or an individual's story. Whereas comprehensive information that gives a more accurate and detailed idea of the bigger picture, such as statistics or reports, is difficult to process. These issues were discussed in Chapter 4, where I argued that an account of the world might stray from the truth but nevertheless exemplify reason-giving properties by highlighting the features of the world that present these reasons. For example, in the case *Animal Suffering*, S grasps the moral wrong of unnecessary animal suffering after watching a video that depicts such suffering, because the video provides access to S's moral reason to avoid contributing to suffering.

In addition to this, some ways of presenting information better capture sustained attention. This point is made by Kathie Jenni (2003) in analysis of our attention to moral issues. Using cruelty towards animals as a case study, Jenni describes how cruel farming practices, that would be condemned by the general public if they happened in front of us, are left unchallenged because they are hidden from view. Even those who at some point became aware of these practices, and felt strongly that they were morally wrong, allow the moral force of this finding to fade or become vague in their mind, such that it does not move them in the same way as it once did. Jenni makes a general point about our attention to moral matters that are hidden from plain view; she points out that we experience a strong response to unfamiliar moral matters (like cruelty to animals and climate change) but that it is normally temporary, ‘hours later the sharpness of the imperative is blunted, and few days later it is gone. We lose our focus; our attention is diverted; we forget’ (2003: 281). It is not just about the initial grasping of unfamiliar moral issues, it is also about whether this information captures us in a way that sustains our attention through time. To preserve our understanding of a moral issue, we must continue to feel the weight of this information by vividly representing it in our minds.

Big picture thinking virtues must be sensitive to these tendencies. The outward-looking agent, for example, will take particular instances to be an example of a wider trend. It is possible to experience video imagery of environmental harm and, while appreciating the harm displayed in the film, fail to appreciate the wider implications. It is of central importance that the outward-looking agent who considers particular and intelligible instances in this way can take them as an example of something bigger, otherwise they miss the point. The outward-looking agent will also be disposed to seek intelligible representations of information, such as video imagery. Sometimes this sort of information is not available to us, or we are prone to remembering it less vividly than we first experienced it. The outward-looking agent will also be disposed to rectify this, perhaps by using their imagination or their memory to vividly represent environmental harm in a manner that does justice to its moral weight (Jenni, 2003: 289).

It is important to note that, relative to an individual’s limits, vividly representing environmental harms may be too upsetting. One objection here is that improvement with respect to cognitive constraints carries a cost when it comes to emotional constraints. My

response to this objection is that big picture thinking virtues do not demand that we push ourselves beyond emotional constraints or cognitive constraints. An outward-looking agent, for example, might be more or less vulnerable to the emotional toll that making information intelligible takes and, in these cases, emotional limits should be taken as seriously as other limits in our judgements about character. As such, it's inappropriate to hold agents responsible for a failure of outward-lookingness if the failure to achieve relevant understanding is attributed to emotional constraints, rather than not being positively oriented towards acquiring understanding. There is simply a limit to how much knowledge and understanding big picture thinking virtues can deliver for a particular agent in certain circumstances.

To conclude, to possess big picture thinking virtues involves being sensitive to one's own vulnerability to fail to grasp or fail to remember information if it is represented in a less vivid manner. In order to maintain a grasp of issues that are not present in their immediate surroundings, good environmental thinkers are disposed to seek intelligible and vivid representations, use their imagination, and to remind themselves accordingly. This appears to be a key takeaway from arguments in Chapters 4 that a focus solely on truth and knowledge may overlook.

6.3.3 Order of Priority

Although those who focus on intelligible representations will have more success with understanding environmental issues, the sheer volume of information remains a big problem. This problem is worsened in some ways by the previous section, given that it is potentially more demanding to seek out particular instances or otherwise intelligible representations of information about every single issue. We still need a principled way to limit what sorts of information the environmental thinker engages with, since they cannot engage with it all.

The good environmental thinker needs to be able to properly prioritise what should be included in their inquiries and what should not. Kawall (2009) suggests some 'initial rules of thumb' for how to prioritise information about environmental issues. These rules are that we should prioritise according to (1) the degree to which are personally affected or able to influence an issue, (2) the perceived severity of the issue (and the extent of attention it is

already expected to attract), (3) strength of evidence that this issue is moral in nature, (4) the ease or cost to ourselves if we pursue this information, and (5) our own preferences.

I would add to these criteria. Although it makes sense to prioritise serious moral issues, particularly where they are neglected and we have the power to make a difference to them, there is room to introduce priorities according to the notion of purely epistemic considerations which I discuss in Chapter 1. If there are features of our environment that are of epistemic value to understand just because they are interesting, they should have priority over features of our environment that do not have this property. This would mean that outward-looking thinkers may be disposed to prioritise spatially distant, but intriguing natural phenomena in their account of the world (or a particular part of the world) rather than less interesting natural phenomena that is close by. This may not come under criterion (3), since the issue is not necessarily moral but a matter of purely epistemic value of understanding what is interesting. This would not come under criterion (5) either, because I have argued previously that whether something is deserving of our interest is a separate matter from whether we have an interest in it.

6.3.4 Awareness of Unawareness

Although these ideas go some way in making sense of what the big picture thinking virtues are like, how they guide environmental thinkers and explain responsibility for environmental harm brought about in ignorance, this still cannot be the whole story.

Although prioritising information might be a good way to identify environmental causes that we wish to invest significant time and attention to, many of our individual actions have an environmental impact. If almost everything we do makes a small contribution to a big problem, then perhaps we cannot make a meaningful difference by prioritising in this way. In addition to this, prioritising according to whether we are affected by an issue, whether we can influence an issue and our own preferences can sometimes perpetuate the biases that big picture thinking, particularly outward-lookingness, is supposed to resist. If so, striving to see the bigger picture in this way still risks rendering the issue unintelligible.

My response to this is that there is a way to look outward (and see the bigger picture in general) that is about being aware that we cannot acquire an understanding of many of the

implications of our actions outside of our immediate surroundings. We cannot always digest the information we might need about the spatially or temporally distant impact of our purchases, for example, and perhaps in this case the outward-looking person is constantly confronted with this state of unawareness. The outward-looking person, who is aware that they lack the vivid representations required to grasp their environmental impact, might perceive themselves to be moving through the world surrounded by a fog, within which the potentially severe moral impact of their actions lie. Because the outward-looking person lives in this scary and uncertain reality, they simply limit their participation in activities that have consequences which extend into the fog, e.g. they generally purchase less.

This idea of the outward-looking agent who is constantly confronted with their unawareness validates the significance of barriers that prevent us from truly ‘seeing the bigger picture’ and at the same time secure a way to be virtuous that does not lead the agent astray by making environmental issues unintelligible. This way of being outward looking is not so much about seeking a volume of information that we could not possibly comprehend, but ensuring that we are keenly aware that we cannot comprehend it. This protects us from underestimating or ignoring the environmental harm produced by our actions, harm that may not be obvious in our everyday lives.

6.4 Conclusion

The guiding idea in this chapter was that being a good environmental thinker involves seeing the bigger picture. I have argued that outward-lookingness, alongside other virtues, has a distinctive role in overcoming the barriers faced by the good environmental thinker, by contributing to big picture thinking. In Chapter 7, I will focus on honesty as an environmental-epistemic virtue.

Chapter 7: Honesty

7.0 Introduction

In Chapter 6, I provided a demonstration of distinguishing and characterising new environmental-epistemic virtues, by exploring the virtue of outward-lookingness in relation to other virtues that contribute to big picture thinking. Given that the status of outward-lookingness as a virtue is in question, the previous chapter mostly focused on justifying the distinction between this virtue and others and defending the virtue status of big picture thinking virtues. In contrast, this chapter provides a demonstration of revising a familiar virtue of honesty, via the environmental lens. Given that the status of honesty as a distinct virtue is not the main point of contention, this chapter focuses on the nature of honesty and whether revisions to an account of honesty become apparent in the context of being a good environmental thinker.

Section 7.1 presents some initial reasons to take interest in honesty as an environmental virtue. Section 7.2 argues for some desiderata for an account of honesty. Section 7.3 presents the case for two prominent and highly intuitive claims about honesty: (i) that honesty should be characterised as an avoidance of deception and (ii) that deception is construed in terms of truth. Section 7.4 raises a challenge to both of these ideas, that arises in an environmental context. Section 7.5 argues that honesty involves a motivation to avoid deception when we are rightly trusted not to deceive, where deception may include distortions to understanding that tolerate or even rely on the distortion of truth. Section 7.6 demonstrates some applications of honesty in an environmental context.

7.1 Is Honesty an Environmental Virtue?

In Chapter 6, I argued that outward-lookingness is an environmental-epistemic virtue because it helps us to overcome a challenge in our epistemic lives that is particularly profound for the environmental thinker: the challenge of how to consider vast amounts of complex information on matters that extend beyond our own lives and that may not be obvious from our surroundings. We can similarly argue that honesty helps us overcome a

challenge in our epistemic lives that is also particularly profound when it comes to thinking about environmental issues. Honesty clearly manifests in cases where an agent chooses not to take up an opportunity to gain an advantage over others by deceiving them.

Unfortunately, there is sometimes an advantage to engaging in deception about environmental issues, but possessing the virtue of honesty will prevent an agent from seeking these advantages by deceiving others.⁹¹

Due to its role in our interactions with others, honesty is thought of by some as an ‘other-regarding’ virtue: a virtue that aims at producing knowledge in not just oneself but for others in one’s community (Kawall, 2002: 259). We should recognise other-regarding virtues, alongside self-regarding virtues that help us acquire knowledge ourselves and that have traditionally been of most interest in epistemology, because other-regarding virtues have an important role in producing and preserving knowledge within a community. Honesty can be an environmental virtue because of its role in producing and maintaining knowledge and understanding about the environment within a community, where there may be an advantage to deception. However, the other-regarding element of honesty is not the whole story.

If we can deceive ourselves, then it seems that we can be dishonest or honest with ourselves, and there is a self-regarding element to honesty.⁹² Christian Miller (2021: 59) claims that self-deception is about one’s own beliefs and whether they are formed appropriately, meaning that avoiding self-deception is a matter of *epistemic* honesty. Miller discusses four different ways that we can form beliefs with self-deception. If we desire *p*, we may deceive ourselves about whether *p* in the following ways: if we (i) dismiss what we would otherwise take to speak against *p*, (ii) misinterpret something in support of evidence that *p*, (iii) selectively attend to evidence that *p* and not on evidence against *p*, and (iv) selectively gather evidence for *p* and overlook evidence against *p*.⁹³ If the virtue of

⁹¹ In *Merchants of Doubt*, Naomi Oreskes and Erik M. Conway (2010) discuss the role of manufacturing doubt in scientific findings about global warming. Another example may be ‘greenwashing’, where a company misleads their consumers about the negative environmental impact of their products and services or exaggerates their positive environmental policies.

⁹² There is much debate surrounding self-deception and its moral and epistemic relevance. For example, there are different interpretations for why we should attempt to prevent self-deception, e.g. self-deception negatively impacts autonomy (Baron, 1988), suggests lack in self-control (Jenni, 2003) or involves an inappropriate response to evidence (Nelkin, 2012).

⁹³ These criteria are originally found in Alfred R. Mele’s *Self-Deception Unmasked* (2001: 26-27).

epistemic honesty precludes these forms of self-deception, then there is a sense in which being honest with oneself is an important virtue in guiding inquiry.⁹⁴

Much like deceiving others, we might say there is unfortunately an advantage to be gained from deceiving oneself about environmental issues. It is often difficult and unpleasant to confront environmental reality and to accept the revision of one's beliefs and actions that might follow; denying environmental reality is a way of avoiding that. Some claim, for example, that peoples' beliefs about anthropogenic climate change tend to correspond with their political beliefs because views about climate are formed through motivated reasoning (Druckman and McGrath, 2019). Motivated reasoning describes the way in which prior beliefs and desires interfere with good reasoning, consciously or unconsciously. We might think of this as a form of self-deception, at least to the extent that it is driven by prioritising our prior beliefs and desires over arriving at an accurate picture of the world. The kinds of self-deception that we engage in can get complex, beyond simply denying the basic facts about climate, and may involve compartmentalising abstract beliefs about climate and refusing to allow them to influence our everyday lives.⁹⁵

This section has established some good reasons to take interest in honesty as an environmental virtue, due to both its other-regarding and self-regarding element. I will return to some of these issues in more detail in the final section. It's not initially clear how honesty helps us overcome the epistemic problem of processing vast amounts of complex information, but being honest with oneself and others entails confronting complex environmental realities and, we shall see in the following discussion, this affects how honesty manifests as an environmental virtue. Honesty not only concerns upholding the truth, but also taking care not to distort understanding.

⁹⁴ This argument is strengthened if avoiding self-deception is thought to have a central role in what it means to engage in rational inquiry. See Andow (2023) for an argument that epistemic deontology should be considered as an alternative to epistemic consequentialism because it can explain why some epistemic trade-offs are impermissible in terms of paternalistic self-deception.

⁹⁵ This form of denial is discussed in *Living in Denial* (Norgaard, 2011).

7.2 What Does it Mean to Be Honest?

Despite the fact that honesty is often cited as an example of virtue, some point out that honesty has been neglected in virtue ethics (Miller, 2021: ix-xi; Wilson, 2018: 262).⁹⁶ In order to fill this gap in the literature, Wilson and Miller have endorsed versions of four desiderata for an account of honesty that aim to capture something important about what it means to be an honest person or identify what is at the core of honesty (Miller, 2021: 27; Wilson, 2018: 267). I also endorse versions of these desiderata in my development of honesty as an environmental-epistemic virtue. This section clarifies my take on them.

7.2.1 The Unification Challenge and the Scope of Honesty

There is agreement that being honest involves the avoiding a range of apparently quite distinct behaviours, including lying, misleading, stealing, cheating and promise-breaking (Hursthouse, 1999: 10; Miller, 2021: 7-8; Wilson, 2018). If a person regularly engages in any of these behaviours, then they are not honest. Miller (2021: 20-21) argues that honesty is associated with the following dispositions:

Truthfulness: the disposition to reliably tell the truth for good moral reasons.

Forthrightness: the disposition to reliably avoid misleading by giving a sufficient presentation of the relevant facts for good moral reasons.

Being Respectful of Property: the disposition to reliably respect the property of others for good moral reasons.

Proper Compliance: the disposition to reliably follow the relevant rules in a situation of voluntary participation when they are fair and appropriate and when there are good moral reasons to do so.

Fidelity to Promises: the disposition to reliably keep promises for good moral reasons.

These dispositions concern the avoidance of a range of distinct behaviours and this presents a challenge for an account of honesty.

⁹⁶ There are numerous instances where honesty is assumed to be a virtue or provided as an example of virtue in influential works in virtue ethics and virtue epistemology (Zagzebski, 1996: 60; Hursthouse, 1999: 11; Driver, 2003). However, it has been argued that honesty is not a virtue (Bloomfield, 2022). In *Environmental Virtue Ethics*, some associate honesty with environmental stewardship (Dzwonkowska, 2025: 101; Sandler, 2007: 82).

The challenge for an account of honesty is to explain what each of these apparently distinct dispositions have in common that makes them part of the supposedly unified virtue of honesty. This is what Miller calls the *Unification Challenge*. The *Unification Challenge* demands that either (1) we continue to think honesty is a unified virtue and identify what these dispositions have in common that make them part of being honest, (2) we admit that not all these dispositions are part of being honest and offer an error theory for why some dispositions are mistakenly associated with honesty or (3) we deny that honesty is a unified virtue and explain why honesty is thought to be a distinct virtue alongside others, such as courage and open-mindedness, so much so that it is commonly included as an example of virtue (Miller, 2021: 24-25). It is important that an account of honesty solves the *Unification Challenge* because if we don't explain why we should think of honesty as a unified virtue that is associated with these apparently distinct dispositions, we have said very little about the nature of honesty. I keep the name of this desideratum as the *Unification Challenge*.

Although honesty is associated with these narrow dispositions, and hopefully we can arrive at an account of honesty that explains what they have in common, it is notable that we often use the term 'honest' to describe a person who is good in general. Some evidence of this is that the online Cambridge dictionary includes many words that indicate general sense of morality as synonyms of honesty, e.g. uprightness, honor, guiltlessness, and, indeed, virtuousness (Cambridge Dictionary, 2025). By comparison, other examples of virtue (e.g. courage, open-mindedness, and temperance) do not have words like this listed as synonyms. This would suggest that honesty has a far reach in comparison to other virtues. It would be a benefit of an account of honesty if it had an answer to this deeper question of whether being honest has a particularly wide scope relative to other virtues and why this might be. Although related, this desideratum is distinct from the already named *Unification Challenge*. I therefore call this desideratum *The Scope of Honesty*.

7.2.2 Avoid or Explain Counterintuitive Results

Another desideratum for an account of honesty is to *Avoid or Explain Counterintuitive Results*. It is desirable for an account of honesty to avoid or explain counterintuitive results about who is honest and who is not. The reason we want verdicts about honesty to accord with our intuitions is that counter-intuitive accounts may not capture the role of honesty in

our moral or intellectual lives. If left unexplained, counterintuitive results are evidence that we have gone wrong somewhere.

To unpack this further, recall that in Chapter 6 I individuated virtues by assuming that we encounter various challenges and distinct virtues help us overcome these challenges in different ways. It is in this sense that epistemic virtues have distinct functions in our epistemic lives. Plausibly, we have some working theory about what these functions are that guides us in our verdicts about who is honest and who is not. If an account of honesty contradicts this working theory about the function of honesty, we must provide a principled explanation for why this is.

To demonstrate this, let's consider accounts of honesty that deliver counterintuitive results. Some accounts of honesty under-generate honest verdicts because they are too demanding, e.g. Carr's (2014: 9) account of honesty as a 'the disposition to seek and honour truth' has the counterintuitive result that we can fail to be honest by being incurious. Accounts like this contradict our general understanding of honesty, which is that being honest does not require us to extend our inquiries and uncover new knowledge and understandings, but rather that we uphold the knowledge and understanding we already have and preserve it in our communications with others. Unless we have a good way to explain why we are wrong about this, this counterintuitive result cannot be explained away and this counts against Carr's account (Miller, 2021: 266).

Honesty also appears to be relativised to an agent's perspective and accounts that do not recognise this, e.g. by requiring that an honest agent tells the truth, also deliver counterintuitive results. Consider the following cases:

The Flat Earth Society 1: A member of the Flat Earth Society, Samantha, sincerely believes that the Earth is flat. One day she is asked by a friend about the shape of the Earth, and to keep her own beliefs a secret, Samantha tries to deceive her friend and replies that the Earth is round.

The Flat Earth Society 2: Now suppose instead that Samantha is forthright. She tells her friend that he believes the Earth is flat and has no intention to deceive her friend at all (Miller, 2021: 36-37).

We want to say that Samantha is dishonest in *Flat Earth Society 1*, even though she says something true, and she is honest in *Flat Earth Society 2*, even though she says something false. To be honest is to refuse to deceive others and to be dishonest is to be willing to deceive others; this means that whether Samantha is mistaken when she attempts or refuses to deceive is not relevant to our assessments about whether she expresses honesty. This contrasts with curiosity, perseverance or carefulness, because we can fall short of these traits without intentionally denying or deviating from the truth. An account of honesty should uphold verdicts like this or explain why they are mistaken.

7.2.3 Corresponding Account of Vices

The next desideratum is *Corresponding Account of Vices*. Minimally, an account of honesty must be compatible with a corresponding account of dishonesty. This is clearly a requirement for an account of honesty, but this chapter in particular demands something more in line with what Miller (2021: 26) argues, which is that an account of honesty should not only be compatible with an account of dishonesty but should also help us understand it. This is because theorising about virtue should not be carried out independently from theorising about vices. Moreover, while most agree that dishonesty is a vice of deficiency associated with honesty, there is much disagreement about whether honesty has a corresponding excess vice. An account of honesty should make progress on this question as well.

According to the Aristotelian framework, virtue is a character trait that manifests a balance between a vice of deficiency and a vice of excess. There is general agreement that *dishonesty* is a deficiency of the trait of honesty, in that being dishonest is to be lacking dispositions and motivations associated with being honest. However, there is much disagreement about whether there is an excess vice associated with honesty and, even among those who agree there is an excess vice, its nature is contested. Annette C. Baier (1990: 86) appears to suggest that an excess of honesty is brutal frankness. Miller (2021: 139-141) argues that we should abandon the Aristotelean framework with respect to honesty, a virtue that does not have an excess vice. The apparent candidates like brutal frankness and the vice of ‘telling all’, i.e. sharing too many irrelevant truths, are instead explained in relation to other virtue traits: the former is an excess of the distinct virtue of frankness and the latter a lack of tact. Heather Battaly (2024: 422) denies this conclusion in

her discussion of the Tell All. The Tell All says what they think is true, without intending to deceive, but does so in excess, sharing irrelevant details about what they had for breakfast. While Miller thinks that the Tell All is ‘exemplary’ in the virtue of honesty and falls short of tact, Battaly argues that Tell All is not virtuously honest because they do not have good judgment regarding the right way to be honest in a given context. Note that we can also tell all to ourselves (in an excess of self-regarding honesty), when we endeavour to avoid self-deception to an extent where we keep lots of irrelevant truths before our mind, e.g. in attempt to make sure no relevant evidence is dismissed.

Yet another possible excess vice associated with honesty is ‘indiscretion’. P. Quinn White (2022: 6) argues that honesty, the virtue of telling the truth, and discretion, the virtue of withholding the truth, work together to answer the question ‘when should I be truthful?’. On this picture, we need both virtues to be truthful in the correct way. One way to fail is to practice indiscretion, e.g. by sharing personal information that other people have trusted you with. Indiscretion could be thought of as an excess vice associated with honesty, at least in the sense that the virtues of discretion and honesty work in opposite directions and the dispositions associated with being honest (governing when to share truths) might undermine the dispositions associated with having discretion (governing when to withhold them). This means that being too honest may lead one towards indiscretion. We can also imagine a self-regarding expression of indiscretion; if there are truths that we should not know about, we can express indiscretion when we don’t take the appropriate steps to conceal them from ourselves.

There is clearly a lot of disagreement about the vice of excess associated with honesty. It is important that we answer this question about what characterises a vice of excess associated with the trait of honesty for at least two reasons. First, the answer to this question speaks to what is at the core of honesty and the correct way to think about its function in our epistemic lives. Miller’s picture characterises honesty as an avoidance of dishonesty. Battaly’s picture of honesty regulates the sharing of relevant truths, where the deficiency is to not share enough and the excess is to share too much. Baier’s picture is similar, but sharing too much takes on the character of brutal frankness, meaning part of the function of honesty could plausibly involve being mindful of the impact of certain truths. White’s picture, however, is situated within relationships and trust between people and regulates the sharing or withholding of truths that occurs within in those relationships.

These different views about the corresponding excess vice affect how we characterise the virtue of honesty. Second, recall that my interest in Environmental Virtue Epistemology is to provide guidance for environmental thinkers and means to evaluate when thinkers are responsible for attitudes and actions. In order to achieve these aims, we must understand not only how inquiry can go right, but also how it can go wrong. This is part of the reasoning for the emergence of ‘vice epistemology’, as ‘epistemological guidance will include measures on how to reduce the pernicious influence of the intellectual vices, and being alert to the influence of such vices may be expected to improve human inquiry’ (Cassam, 2016: 161). We should, therefore, strive to understand vice, including vices of excess.

In summary, an account of honesty must at least be compatible with a plausible account of dishonesty. However, it is also preferable for an account of honesty to say something about the nature of corresponding vices. If I’m right that we cannot really have a good picture of what honesty is and how to be honest unless we settle the question of whether we can have too much of it, we should also seek an account of honesty that contributes to the question of whether honesty has an excess vice and, if so, what it looks like.

7.2.4 Explanation of Value

If honesty is a virtue, there must be something about this character trait that explains its virtue status. Being virtuous is about being a good person and, therefore, virtues must bear some relation to something that is of value. Wilson, in discussion of moral honesty, writes that ‘an account of honesty ought to at least point in the direction of an explanation for why honesty is a moral virtue’ (2018: 268). For example, we might think that kindness is a virtue because involves being motivated to have caring relationships. In the same vein, an account of honesty as an epistemic virtue must point to why honesty is an epistemic virtue, e.g. because it bears some relation to epistemic value. The final desideratum, then, is that an account of honesty should provide an *Explanation of Value*, an explanation with certain features that I will set out in this sub-section.

Firstly, an explanation of the value of honesty may not be straightforward, as honesty is relevant to both moral and epistemic values. I have spoken as if moral and epistemic honesty can be distinguished, but it is not clear how we should make this distinction or

whether it is even possible to make a clear distinction.⁹⁷ Assuming that we can make a clear distinction between epistemic and moral values, it is unclear whether honesty should be understood as an epistemic virtue, a moral virtue, both an epistemic and moral virtue at the same time or two distinct virtues, one epistemic and one moral. In the context of this thesis, an explanation for the value of honesty must shed light on this.

Moreover, it is important to reiterate that I am interested in responsibilist virtues and therefore the explanation of value should, in accordance with the arguments in Chapter 5, manifest an important feature of who the agent is that they can be held responsible for. I have suggested that, where honesty is an epistemic virtue, the honest person must be ‘for’ or positively oriented towards what is of epistemic value.

7.3 Honesty, Deception and Truth

There is an intuitive way of thinking about honesty that I go on to reject in Section 7.4, which centres around the following two claims. First, at the core of honesty is an avoidance of deception. Second, deception is fundamentally construed in terms of truth. Even though I go on to reject this view, there are important reasons to discuss these ideas because this attempt to meet the desiderata in the previous section reveals something important about what it means to be honest. Moreover, these are highly intuitive claims about honesty that I reject, so it is important to consider why they are so compelling.

Many characterise honesty as an avoidance of deception, or an intentional distortion or denial of the truth. Some accounts focus on dispositions to avoid deception. Louis M. Guenin argues that the ‘kernel of *intellectual honesty* consists in a disposition of an agent such that when presented with an incentive to deceive in any way, the agent will not deceive’ (2005: 217). Miller (2021) argues that honesty involves being disposed, centrally and reliably, not to intentionally distort the facts as the agent sees them. Others focus on the honest agent’s motivations. Wilson (2018) argues that honesty centrally involves a deep motivation to avoid deception, where a deep motivations are thought to be persistent, strong and robust. In the context of responsibilist approach to virtue, I am interested in

⁹⁷ Louis M. Guenin (2005) provides an in depth account of intellectual honesty. Miller (2021: 111-112) focuses on moral honesty, but includes suggestions about epistemic honesty.

virtues that manifest something important about an agent that they can be held responsible for, and therefore Wilson's approach is most relevant for my purposes. However, what is important to the point I make in this section is that all these accounts characterise honesty as an avoidance of deception.

It is natural to understand deception in terms of truth. Guenin (2005: 200) explicitly defines deception as an utterance that the speaker takes to be false or misleading, where misleading statements contribute to acquisition of false belief, maintenance of false belief, loss of true belief, or prevention of true belief. In other accounts, it is less clear that deception must be made sense of in terms of true and false belief, but accuracy is still central. In discussion of what it means to distort the facts, Miller does not provide an analysis of 'distorting' but writes that 'a close synonym to what I have in mind is captured by misrepresentation: an honest person is someone who is disposed to not intentionally misrepresent the facts, either to herself or to other people' (2021: 30). Insofar as misrepresenting the facts means representing things differently from the way they are, i.e. falling short of true representation, the deception in Miller's account is also fundamentally construed in terms of truth.

The challenge for accounts of honesty that uphold these two ideas, i.e. that honesty is avoidance of deception and deception fundamentally concerns truth, is that they must accommodate all of the narrow dispositions at issue in the *Unification Challenge*. While truthfulness and forthrightness are straightforwardly related to deceiving others about the truth, the connection is less clear for the other dispositions. Hence, some have distinguished between the virtues of *truthfulness* and *justice*, which are both thought to be subsumed under honesty (Roberts and West, 2020). Truthfulness involves upholding the truth, most obviously by telling the truth. Justice involves treating others fairly and delivering what is owed to others; fidelity to promises, proper compliance and respect for property seem to fall primarily into this latter category. It is notable that, despite dividing honesty in this way, Roberts and West still seem to centre truth, writing that: 'while honesty is a two-handed virtue, it's not quite ambidextrous. Rather, honesty's dominant hand is truthfulness' (Roberts and West, 2020: 212). This goes to show how strongly honesty is associated with truth.

Rather than divide honesty in this way, others meet the *Unification Challenge* by referring to truth (Miller, 2021). Truthfulness is most clearly explained, as lies involve deceiving others by saying things that you believe to be false. Perhaps forthrightness can also be explained in terms of truth, because misleading others involves deceiving others, and even though you do not necessarily say something false, you do intentionally lead others away from the truth. With respect to proper compliance, failing to follow the rules is to deny the truth about these rules by acting as if they do not apply and, in some cases, misrepresenting other facts, e.g. presenting yourself as having skill that you do not possess (Miller, 2021: 42-45). If someone routinely breaks their promises, this suggests that they do not make promises with the intention of keeping them, thus deceiving others by leading them to believe something false about your intentions (2021: 49-51). Similarly, stealing often involves acting as though something belongs to us when it does not, which involves distorting the truth about what belongs to whom and the norms of property (2021: 45-48).

However, we can now see that there are some problem cases with respect to some of these dispositions. For example, sometimes we break promises we had every intention to keep, in which case we break promises without deceiving others about our intentions. These are cases where promise-breaking may not amount to dishonesty (Wilson, 2018: 269). Another problem case is the bald-faced thief, who does not attempt to mislead anybody about what belongs to whom. Miller (2021: 47) responds to this by saying that if the thief believes they are doing wrong, they are still intentionally distorting the normative facts about how we should treat others' property (as the thief sees them) and so they are still honest. With respect to a case where the thief is mistaken about these normative facts, and does not see themselves as doing any wrong, we would have to deny that they are dishonest. (Of course, we could explain their failings in terms of other virtues.) The general approach to cases where an agent does not engage in any deception or take themselves to be distorting the truth in any way, is that we are right to associate fidelity to promises and being respectful to property with honesty, but not all instances of stealing and promise-breaking are dishonest. I find this to be an acceptable result. It seems to right to say that these are not instances of vicious dishonesty, but instead failings of other vices to do with fairness, justice and kindness. I conclude that, at this point, it appears to be an advantage of truth-based accounts that they can meet the *Unification Challenge* without having to deny that honesty is a unified trait.

With respect to *Explanation of Value*, an account of honesty as avoidance of deception, where deception is about truth, might explain the distinction between moral and epistemic virtue in terms of the moral and epistemic values associated with truth. In other words, truth has both epistemic and moral value, so if honesty is fundamentally related to truth, then we can explain why honesty is both a moral and epistemic virtue by reference to the moral and epistemic values of truth. An epistemic value of truth might be truth for its own sake; we can say epistemic honesty is present when someone is oriented towards the truth for its own sake.⁹⁸ In contrast, the moral value of truth includes its role in trust and autonomy. We might say, for example, that moral honesty is present when someone is oriented toward avoiding deception as this undermines autonomy and trust. An account of honesty that contrasts with deception seems to have a coherent story to tell about how to distinguish epistemic and moral honesty.

A truth-based account of honesty could also lend initial support to the idea that there is no excess vice associated with honesty. Notice that the Tell All, the person who expresses indiscretion and the brutally frank, for example, all avoid deceiving. So, if honesty is just about *not* deceiving, then all these agents qualify as honest.

In the next section, I challenge both the claim that the core of honesty is avoiding deception and the claim that deception is construed in terms of truth.

7.4 Honesty and Truth

In this section, I challenge the two ideas in the previous section: that honesty is about avoiding deception and that deception is fundamentally about truth. I first question the claim that honesty is an avoidance of deception by arguing that not all deception amounts to vicious dishonesty and therefore honest people may be motivated to deceive in some circumstances. I argue that this is because honesty is more fundamentally about not betraying others or ourselves, by engaging in deception when we are rightly trusted not to. I then argue that thinking of honesty in terms of fidelity to the truth (or not intentionally

⁹⁸ Others have argued that the motivational component of honesty might be the point at which epistemic honesty departs from moral honesty. Miller suggests that epistemic honesty has a motivational component that ‘can be unpacked using notions like *being concerned about confiding the truth*, *loving the truth*, and *caring for the truth*’ (2021: 112).

distorting the truth) delivers counterintuitive results in two sorts of cases. In the first case, an agent deceives by intending to uphold the truth. In the second case, an agent expresses honesty by intending to distort the truth.

7.4.1 Not All Deception Matters to Honesty

This section argues that there is more to honesty than avoiding deception because not all deception prevents a person from being honest. To be clear, the worry is not that honesty is restricted to a certain kind of deception, but that something else is at the core of honesty, rather than deceiving others or leading them away from the truth.

First consider a familiar point that not all instances of deception are objectionable. In fact, sometimes we seem to be required to lie to others, e.g. lying to a murderer about the whereabouts of their intended victim. Another point is that whether deception is objectionable depends on the nature of relationships:

Claudia has an intimate marriage with her spouse, Sally; they are (typically) quite open with one another and are committed to living their lives together. Where Claudia may deceive her co-worker about her cancer, she may not similarly deceive her wife; that would constitute a grave betrayal. Indeed, even if no deception were involved, simply not telling Sally seems objectionable. Where lying is permissible with a co-worker, deception or even mere withholding of information is impermissible with a spouse; and that seems connected to the fact that a cancer diagnosis is (in some sense) Claudia's wife's business but not her co-worker's (White, 2022: 10).

In cases like the above, Claudia is dishonest if she lies to her spouse about her diagnosis and perhaps even if she withholds this information. Whereas, if Claudia withholds this same information from her co-worker, this is not dishonest. Even if she lies about it, this is not obviously objectionable.

One response to cases like this is to say that honesty only requires we generally try to avoid deception and cases where deception is not objectionable are exceptional. A few instances of deception in exceptional circumstances does not make a person dishonest and this is compatible with an account that characterises honesty as avoidance of deception

(Miller, 2021: 136). However, the problem with this response is that it treats the case where being honest is relative to the nature of relationships as the exception, when in fact it is the norm. As argued by White (2022), whether deception is objectionable is always sensitive to the nature of the relationship in which the deception occurs. Relationships have constitutive ends or functions, e.g. colleagues have the aim of fulfilling work tasks and friends might have an intrinsic ‘aim’ of mutual respect or love, and a person can be judged against these aims, e.g. as a good friend or a good co-worker. These ends inform communicative standards that dictate when we should share information with others and when we shouldn’t. This could explain why it is objectionable for Claudia to deceive her spouse but not her co-worker. The nature of the relationship produces communicative standards and this is what explains why we rightly trust some to share certain truths and why we cannot rightly trust others to do the same.

Given that whether deception is objectionable or not is relative to relationships in this way, and these cases are not exceptional, it is not deception as such that the honest person avoids, but a certain objectionable kind of deception. In support of this point, consider cases where lying is obligatory, such as a case where someone else confides in Claudia about their diagnosis and Claudia must lie to avoid sharing this private information with others. We would not judge Claudia to be a dishonest person because of this. To strengthen this point, consider a world where an agent must lie due to unfortunate circumstance, a world where they are regularly morally required to lie because they are surrounded by people who regularly ask about personal information. Being in these circumstances does not make a person dishonest in a way that would constitute a vice, at least not according to the responsibilist account I endorse, because the fact that they must lie does not reflect badly on who they are.

If we take my preferred approach to this adapted case and say that Claudia lying to her co-worker about someone else’s diagnosis does not count against her honest character, then honesty is sensitive to the nature of relationships we are in with other people and moral demands that govern when we should share truths with others. We might suspect that there is something more fundamental that explains why some instances of deception matter to honesty and others do not. My view is that the sort of deception the honest person avoids is deception that constitutes betrayal. We can only betray others when they can rightly trust us to fulfil some expectation and we fail to do so. This is not about whether Claudia can be

judged as *likely* to present information, such that it is wise to actually depend on them to do so; in the first case, Claudia's spouse might ask Claudia a question knowing full well that she will lie, and it would not be wise to believe what is said or act on it, but Claudia is still dishonest in this scenario. There is a sense in which Claudia's spouse should be able to trust her and this is because Claudia should live up to some expectation that they will not lie, even if it is clear to everyone that they will not live up to this expectation. In the case of morally obligated lies, Claudia's co-worker should not be able to trust her to give a truthful answer about someone else's diagnosis, because they are not right to trust that Claudia will do so.

I suggest then that being honest is to do with 'being for' or being positively oriented towards the avoidance of deception when we are rightly trusted not to deceive and at the core of honesty is concern for or respect for this trust, rather than simply an avoidance of deception as such.⁹⁹ The idea is that when we manifest this concern for or respect for trust, we are honest, and when we fall short of this, we are dishonest. This suggests that trust more fundamentally explains the status of honesty as a virtue, rather than an abstracted fidelity to the truth or avoidance of deception.

There are at least two important upshots to the arguments in this sub-section. The first is that, if trust is more fundamental than truth, then it is not clear how we should distinguish between moral honesty and epistemic honesty. We can no longer say that epistemic honesty is matter of valuing truth for its own sake because, on my view, the role of truth in honesty cannot be detached from its role in upholding trust. It is trust that explains when upholding the truth matters to honesty. The second is that my view lends support to an excess vice to do with revealing too much when we are rightly trusted not to do so, whereas I suggested in Section 7.3 that conceiving of honesty as an avoidance of deception lends support to the view that honesty has no corresponding excess vice.

⁹⁹ On this account, to be 'rightly trusted' is supposed to capture the relevant communicative norms and moral demands, but I do not attempt to settle the question of their normative source. I just claim there is a sense in which others are right or justified in trusting us to share some truth with them and when we fail to do so, we fail to be honest.

7.4.2 Truth and Deception

As pointed out in Section 7.3, it is natural to make sense of deception in terms of truth. If this is right, then although honesty centrally concerns trust not to deceive, there is an argument that says that truth is still at its core because *not deceiving* always involves faithfulness to the truth. The issue with this approach is that deception does not necessarily concern truth.

To make this point, I will need to recap Catherine Elgin's (2017: 183-203) argument that we can have epistemic access to the world that does not primarily concern truth, via 'exemplification' – as discussed in Chapter 4, Section 4.1. Recall that an *exemplar* provides epistemic access to the world. A sample card of a particular shade of yellow serves as an exemplar of that shade of yellow. An exemplar affords epistemic access to some property by instantiating that property, e.g. the shade of yellow, and downplaying its other properties, e.g. the shape of the card.

As I argued in Chapter 4, exemplification is a different tether to the facts than true representation. Elgin argues that idealisations, oversimplifications and thought experiments can afford epistemic access to the facts by deviating from the truth, because these representations instantiate properties and, by downplaying other features, emphasise these properties. I also argued that degree of grasp corresponds to the degree to which we are in an epistemic position to respond to objective normative reasons and this can be achieved via an exemplification, i.e. where the exemplar instantiates properties that present normative reasons and downplays other properties, resulting in an intelligible representation of the world that enables reason-responsiveness.

Importantly for this chapter, we can undermine understanding by choosing a bad exemplar. I will now go through a few of Elgin's examples of this (2017: 191-192). A tiger instantiates being striped but does not effectively exemplify being striped when this property is camouflaged by the tiger blending into the background. A bald man who wears a hat is not well placed to exemplify being bald because this property is occluded. A lion's roar might be so loud and terrifying that it does not exemplify its pitch, because the pitch is overshadowed by these other properties. A yellowish green is not a good exemplar for the colour green, because it is a borderline case. A good exemplar must not only instantiate the

important property in a given context, but do so in a way that makes this property salient and is therefore easily interpreted.

If we can opt for a bad exemplar that fails to make the relevant properties salient, we can do so deliberately with the intention of making an exemplar difficult for an audience to interpret. This opens up the possibility that we can be deceptive in a way that does not fundamentally concern truth:

Dishonest CEO

The Chief Executive Officer (CEO) of a company wants to minimise the scale of plastic pollution in order to defend their decision not to invest funds and reduce plastic use in the packaging of their products. They do not, however, want to be caught in a lie. They know that their audience is likely to be less familiar with the geography of the United States than Europe, so they choose to say in their messaging about the environment that the Great Pacific Garbage Patch is twice the size of Texas, rather than say it is three times the size of France.

This is the paradigmatic case of dishonesty, in that the CEO gains an advantage by adversely affecting the epistemic states of others.

However, the sense in which the CEO is deceptive is not fundamentally a matter of distorting or leading their customers away from truth. First, let's distinguish the truth of what is said from the truth of what is implied by what is said. This turns out to be quite complicated.¹⁰⁰ For the purposes of this chapter, note the difference between what is directly expressed by the utterance, which at least includes the proposition that 'the GPGP is twice the size of Texas', and what is implied, which might include things such as 'plastic pollution isn't so bad'. We can also distinguish between what is implied and what is in fact believed by the customers as a result of hearing the CEO's statement. Clearly what is said upholds the truth and the CEO deliberately ensures this. However, it might be argued that what is implied or believed by the customers does not uphold the truth. In other words, the objection is that, while what is said is true, the CEO intends to imply falsehoods, e.g. plastic pollution isn't so bad, or to lead their customers away from true beliefs.

¹⁰⁰ There is difficulty clarifying the notion of 'what is said' that is operative in the distinction between lying and misleading (Saul, 2012).

However, the CEO is dishonest because they intentionally choose an exemplar that is difficult for their audience to interpret and therefore directly obscure their customers' reasons, i.e. reasons for customers to think badly of this organisation and purchase elsewhere. This harm does not hinge on the actual or possible impact on true beliefs, even though the CEO clearly disregards this risk as well. The harm directly concerns undermining reason-responsiveness. Given that I have already argued that we can improve reason-responsiveness with mental representations of the world that deviate from the truth, the source of the intended harm is not to do with accuracy, but rather the salience of an important set of reasons. It seems that dishonesty sometimes involves interfering with understanding via bad exemplars, rather than false representation.

It is also possible to avoid deceiving even when we do intend to lead others away from true representations:

Honest Lecturer

A lecturer considers whether to begin an environmental philosophy lecture by drawing a simple line on the whiteboard to illustrate to their students just how steep the increase in average global temperatures is or whether to opt instead for a list of detailed statistics that communicate the same information more accurately. The lecturer thinks their students might find the first option a bit too hard-hitting; last year, drawing the line on the whiteboard led to uncomfortable questions and responses in the lecture and the lecturer doesn't want to acquire a reputation for scaring students or being negative. The lecturer has good reason to believe the statistical information will have less of an impact on their students and this will make the lecturing experience more enjoyable for themselves. In the end, the lecturer decides that this would be too manipulative and self-oriented and, in the interest of the students, draws the line on the whiteboard during the lecture.

In this case, the lecturer is honest because, despite being tempted to opt for a bad exemplar and thus refuse epistemic access to important facts about climate, they choose a good exemplar. Interestingly, the bad exemplar is more accurate than the good exemplar, but it is difficult for the students to interpret and therefore delivers a poorer quality of understanding in comparison to the good exemplar of the drawn line.

Even if you reject my arguments that understanding cannot be reduced to knowledge and truth, there are still key insights in this chapter that do not depend on these ideas. *Dishonest CEO* presents a way of deceiving others that prioritises truth and accuracy and *Honest Lecturer* demonstrates a way of being honest by distorting the truth. This by itself deserves attention, given that the aims of this chapter are to shed light on the role of virtue in guidance and responsibility. Honesty requires us to choose exemplars that are easily interpreted by those we communicate with, and avoid bad exemplars that are not easily interpreted, which affects the sort of guidance that “be honest” is giving us. In other words, honest inquiries are not guided by faithfulness to the truth, but to other considerations as well to do with choosing good exemplars and judging what we should be trusted to share and emphasise to others. This also raises questions for holding others responsible for failing to be honest because it runs against intuitions that lying is somehow worse than misleading by telling the truth, e.g. because there is something specifically worse about uttering something untrue and therefore directly representing the world as it is not. We may also overlook cases of honesty and dishonesty if we overlook the role of exemplification.

7.5 My Account

At the core of honesty is this concern for trust and taking care not to deceive ourselves or others in contexts where we are rightly trusted not to do so. Moreover, deception can include distortions to one’s understanding, even when this does not concern distorting the truth. We can therefore explain why the CEO is dishonest: they deliberately choose a bad exemplar. They choose an exemplar that interferes with appropriate epistemic access to facts in a scenario in which customers are rightly trust them not to. We can also explain why the lecturer is honest, because they choose not to deceive in this way.

This account of honesty, that centres trust and recognises the role of understanding, preserves some features of other accounts. For example, the solution to the *Unification Challenge* is broadly the same as the solution discussed in Section 7.3, since deception plays a role in each narrow disposition. How we approach the *Unification Challenge* will, however, depend on whether *Dishonest CEO* and *Honest Lecturer* count as cases of misleading or not. If we define a misleading verbal statement as a statement that is not a lie but ‘carries a conversational implicature intended to lead the audience to a false belief’,

then there is another disposition to avoid choosing bad exemplars that is relevant to honesty and that is distinct from ‘forthrightness’ (Miller, 2021: 11). If this is right, accounts of honesty that centre truth solve the original *Unification Challenge*, but I have presented a new challenge which they cannot solve, a challenge that includes a distinct narrow disposition to do with not choosing bad exemplars. Alternatively, we can interpret choosing bad exemplars as a form of misleading and choosing good exemplars as a form of forthrightness. In which case, accounts of honesty that centre truth do not account for all forms of misleading and therefore do not fully meet the *Unification Challenge*.

We might think that I have only made subtle revisions to honesty as it is discussed in Section 7.3 and, in some ways, this is right. That’s why accounts of honesty that centre deception and truth do so well with respect to meeting the desiderata. However, we would be mistaken to think that this change only matters to marginal cases, like the one’s I have introduced. This account has important contributions to make with respect to *Explanation of Value*, *The Scope of Honesty* and *Corresponding Account of Vices*, which I will now unpack respectively.

7.5.1 Explanation of Value and the Scope of Honesty

On my view, we cannot make a clear distinction between moral and epistemic honesty, but this is not a disadvantage of my account. *Explanation of Value* demands that we can make sense of this ambiguity about whether honesty is an epistemic or moral virtue, not that we can successfully make a clear distinction between two virtues, one epistemic and one moral.

Accounts that centre truth have good way of distinguishing epistemic and moral honesty; the morally honest care about the *moral* value of truth, e.g. in relation to autonomy and trust, and the epistemically honest care about the *epistemic* value of truth, e.g. truth for its own sake. However, given that I have argued that honesty concerns avoiding deception when we are rightly trusted not to deceive, and that honesty manifests a concern for this trust, there always seems to be both an epistemic and a moral element to being honest. It might be that an honest person puts greater emphasis on either epistemic or moral values, but it does not seem possible to untangle them entirely or say one is present while the other is not.

There is potentially something particularly valuable about honesty in virtue of the fact it has both an epistemic and moral element. Whenever an agent thinks they are doing wrong, they deceive themselves or others by acting as if the norms they violate do not apply (Miller, 2021: 43-44). My account of honesty reinforces this point. As long as we can always rightly trust ourselves to do what we take to do be the right thing, and when we fail to do what we take to be the right, we betray our own trust in some sense. There is an advantage here with respect to the desideratum of explaining *The Scope of Honesty*.

Honesty has a far reach, compelling the honest person to uphold other virtues. This is a striking feature of honesty that suggests that honesty has a central role in cultivating and practicing *all* virtues because it requires that we do not deny what we take to be normative reality. This might at first seem like an unwelcome consequence, in that, beyond solving the *Unification Challenge*, we have apparently explained how *any* virtuous disposition could be associated with honesty. However, I see this as an advantage of my view. We can now see why honesty is thought to have such a wide scope; saying that someone is an ‘honest person’ often means you can generally depend on them to care about doing the right thing.

7.5.2 Account of Corresponding Vices

Firstly, my account of honesty is compatible with an account of dishonesty. To be dishonest is either to be lacking in a positive orientation towards avoiding deception when we are rightly trusted not to deceive or to be positively oriented towards deception, i.e. being ‘for’ deceiving others when we are rightly trusted not to. We can see why honesty is contrasted strongly with deception, but we also have the resources to theorise about an excess of honesty.

An excess of honesty is when one’s orientation towards avoiding deception (when we are rightly trusted to) leads one astray, e.g. the Tell All who cares too much about deceiving others and includes irrelevant details. Miller (2021: 140) is not convinced that the Tell All fails to be honest, because the Tell All does not intentionally mislead anybody and their character flaw is instead a failure to discern what to share with others. However, this is because Miller sees honesty as a disposition not to engage in deception, rather than a virtue

that resists betrayal via deception and is therefore sensitive to norms about when we can rightly trust others or ourselves not to deceive. On my account, the Tell All can fail to be honest by failing to practice discernment about what to share with others, but *only if their telling all is a result of caring too much about whether they are engaging in objectionable deception*. If telling all expresses a lack of concern for others or a lack in tact, this is not excess honesty, but if telling all expresses a concern for engaging in deception, it may express excess honesty.

It might be objected that it doesn't feel very vicious to care too much about being deceptive when trusted not to be and the Tell All does not express any vice if their telling all is a result of their honest motivations. Firstly, note that some vices, e.g. arrogance or cruelty, are much uglier than others, such as being too kind or too benevolent. This may be because some virtues are closer to the excess or deficiency vices than others. Maybe being too honest isn't so bad in comparison to being dishonest, because honesty is closer to its excess vice. However, it's also important to recognise that caring too much about not deceiving can take a number of different forms, not just telling all. Being excessively honest may also involve sharing other peoples' private information or unnecessarily sharing truths that will hurt others. These manifestations of an excess of honesty do appear to be quite vicious, even if they stem from concerns about betraying others' trust.

7.6 Applications

In this final section, I discuss how honesty applies in environmental contexts. I discuss testimony, paternalism, and self-deception.

7.6.1 Testimony

Because we often rely on experts for information about environmental issues, testimony is crucial to our understanding. My account says honest testimony is not simply telling the truth. Honest testimony must also involve exemplars that are easily interpreted, which suggests a more complicated picture than knowledge transmission. This is important because it affects the guiding advice that might be given to the testifying agent, as well as how they may be responsible.

In some ways, this account of honest testimony is demanding, because what is required is not only that the testifying agent tells the truth with the intention to produce knowledge, but also that they take care with exemplars and the impact they are likely to have on a particular audience. I have developed a subjectivist account of honesty, which means that what matters is that an agent takes themselves to be furthering understanding not that they have success in doing so. An expert who uses bad exemplars by mistake would not necessarily fall short of honesty, although they may fall short of other related virtues.

While honest testimony goes beyond producing knowledge, and often requires producing understanding, a distinction should be drawn between honesty and other virtues that require us to take more of an active interest in the epistemic position of others. One example is *transparency*, which T. Ryan Byerly describes as ‘a tendency to faithfully share one’s perspective on topics of others’ inquiries with these others out of a motivation to promote their epistemic goods’ (2021: 105). While transparency requires that we share information about our perspective due to a concern for another person’s epistemic position, honesty only requires that we inclined not to deceive them. *Frankness* is also contrasted with honesty, in the sense that it is thought to involve sharing information ‘in a supererogatory way, but without oversharing’ (Miller, 2021: 20). Again, frankness requires sharing more than what is required to avoid deception.

Honest testifiers must often go beyond what is required to share knowledge, by not deliberately choosing bad exemplars that bring about misunderstanding. As we have seen, this doesn’t mean that being honest requires an active interest in improving understanding beyond what is required to avoid deception.

7.6.2 Paternalism

One possible objection to honesty as an environmental-epistemic virtue is that it is not always clear that being honest is the right thing to do when one cares about the environment:

The shepherd sees a wolf is about to eat some sheep from the communal flock. She knows that the villagers will want to know that there is a wolf present, but that a cry

of ‘wolf’ will not lead them to form that belief. She also knows, however, that the villagers will respond to a call of ‘a band of wolves’. By the villagers’ own lights, they would prefer to have a belief ‘that the sheep are in danger’ than not, even if that belief is based on a belief – ‘the sheep are in danger from a band of wolves’ (John, 2018: 83).

The shepherd engages in deception in order to bring about a true belief that matters most: that the sheep are in danger. Similarly, Stephen John (2018: 84) suggests that the scientist should opt for the unqualified estimate (e.g. climate change will lead to ice sheet collapse) if it produces a better epistemic situation for their audience than an honest estimate. This forms an argument against thinking of honesty as an environmental virtue.

Another response to a case like this maintains that the scientist is honest, by saying that the nature of the relationship between the scientist and their audience involves communicative norms that permit the revealing of truths that matter most. This view would say we can only rightly trust experts to reveal the truths that matter most, rather than all truths, and the scientist qualifies as honest as long as they comply with these norms. Having said this, there might be a better response.

The true beliefs that matter seem to be the true beliefs that contribute to understanding (especially on my view, where improvements in our grasp are about responsiveness to normative reasons), and so perhaps we can say that the scientist is not deceptive because their aim is to improve understanding. This view does not commit to anything about the norms of the relationship between expert scientists and their audiences, but instead says prioritising true beliefs that make a significant contribution to understanding, at the cost of true beliefs that do not, is not deceptive. To strengthen this argument, consider that it could constitute a form of deception to act otherwise, e.g. to attempt to produce many true beliefs, instead of prioritising the ones that matter, with the intention to bring about misunderstanding.

7.6.3 Self-Deception

Finally, I argued in Section 7.1 that we should take interest in self-deception, but I have not said as much about the self-regarding element of honesty. In fact, we might worry that my account cannot accommodate honesty’s self-regarding element because of its situation

within relationships with other people. However, insofar as we can make sense of what it means to have a relationship with oneself and what it means to say we should be able to trust oneself, we can make sense of what it means to be honest with oneself.

Generally, we should be able to trust ourselves not to engage in self-deceptive belief forming processes listed in Section 7.1 (positive and negative misinterpretation, selective attentiveness and selective evidence-gathering). The basis for this trust might be, as argued by James Andow (2023: 47), that we should not deceive ourselves even when this benefits our overall epistemic position, because this is to treat ourselves as means, rather than an end. But, perhaps there are cases where we cannot rightly trust ourselves not to engage in self-deception.

There might be cases where convincing ourselves of some false belief is morally obligatory, e.g. perhaps there is a case where, if we do not deceive ourselves by dismissing evidence that we might not recover from an illness, then we reduce our chances of recovery.¹⁰¹ This case would share its structure with morally obligatory lies, that I have already discussed in Section 7.4.1, and the same conclusions about honesty apply. We can deny that we are dishonest in such a case, because we cannot rightly trust ourselves not to deceive, and therefore this does not constitute the betrayal involved in dishonesty. Another example could be ‘epistemic innocence’, where a belief is epistemically irrational, but where lacking this belief carries significant epistemic costs (Bortolotti, 2020). Even though epistemically innocent beliefs are formed by ‘self-deceptive’ processes, we might want to say that they do not imply dishonest character, because you cannot rightly trust yourself to avoid epistemically irrational beliefs when they are innocent in this way.

Interestingly, the notion of picking bad exemplars seems more in line with the way that we engage in self-deception. Many epistemologists think we do not have direct control over our beliefs, meaning that lying to ourselves would be quite difficult; we cannot

¹⁰¹ A similar case is discussed by Selim Berker (2013). A case like this is thought to prompt the rejection of epistemic consequentialism. The idea is that beliefs are clearly not justified just in virtue of their future impacts, such as whether they lead to further true beliefs. My interpretation of this case is different because I suggest that this form of self-deception may not be objectionable. If one does not engage in objectionable self-deception, one is not dishonest to believe in one’s recovery, but this need not imply anything about epistemic justification. At least, not unless you think that justification for belief concerns when one can rightly trust oneself not to deceive. Interestingly, this looks like a view that recognises future impact of a belief in non-consequentialist terms.

directly force ourselves to believe what we take to be the truth in the same way that we can simply say something is not true.¹⁰² If this is right, we must conceal the truth from ourselves indirectly and this could take the form of deliberately seeking poor examples, e.g. those that do not make the salient features of a situation readily available and easily interpretable and thus making important facts easy to ignore, minimise or misunderstand. Perhaps this is what is meant by the idea that we can believe the truth in a way that does not really ‘sink in’. I have in mind a case where we believe some truth *p*, such that we would tell others *p* when questioned about whether-*p* and we would think *p* when we reflect on whether-*p*, but our belief in *p* is not integrated into the way we live our lives. To be honest with oneself will involve engaging with exemplars that are easily interpretable, especially when environmental issues concern vast amounts of complication information that might be particularly easy to deny because it is difficult to understand.

7.7 Conclusion

In this chapter, I have argued that honesty is about being positively oriented to not deceive others or oneself when we are rightly trusted not to do so. Honesty is an important environmental-epistemic virtue, as we can see from the discussion in Section 7.6 about its application in environmental context. We can also see that the honest environmental thinker must take care not to deceive oneself and others by distorting understanding with bad exemplars, which may fail to make environmental issues intelligible.

¹⁰² Some argue that we are not in control of our beliefs, like we are in control of our actions (Alston, 1988; Chrisman, 2008). Matthias Steup (2008) argues that there is no reason to think that we have less control of our beliefs than intentional actions.

Chapter 8: Why Use Case Study Research in Developing an Epistemology for the Environment?

8.0 Introduction

This thesis has so far made use of traditional philosophical methods (which might be referred to as “armchair philosophy”). However, the questions that drive my research are applied: they are epistemological questions that arise from environmental challenges. We might worry that the “armchair” approach could end up being out of touch with the experiences of real people. As I suggested in the thesis introduction, I have chosen to focus on the epistemic problem of processing vast amounts of complex information, but my reasons for this focus do not entail that this is the only relevant epistemic problem or even that it is the most important. This is an empirical question that requires insight into peoples’ experiences. In light of this concern, I have carried out case study research in attempt to understand epistemic factors experienced by those that are involved in environmental intervention. This chapter both defends the use of empirical research in developing an epistemology for the environment and defends the choices made in the case study research I have carried out to illustrate this.

Section 8.1 argues for the use of empirical research in epistemology and introduces my research question. I also address some potential concerns about taking an interdisciplinary approach. Section 8.2 explains the philosophical assumptions that influence my approach to the research question and clarifies their consistency with my wider project. Section 8.3 defends my use of the case study research design. Section 8.4 explains my choice of the case study in question: the Nearshore Trawling Byelaw in Sussex and the ongoing Sussex Kelp Recovery Project. Section 8.5 justifies my use of semi-structured interviews and discusses the interview questions. Section 8.6 defends my approach to data analysis. Section 8.7 details ethical considerations taken into account as this research study was carried out.

8.1 Defending the Use of Empirical Research in Applied Epistemology and Clarifying the Research Question

This thesis began with the idea that we should attend to environmental issues in epistemology, since understanding our natural environment and present environmental challenges raises distinct or particularly profound epistemic difficulties. Much of my thesis so far has addressed this issue with traditional philosophical research methods, but this section argues that we should also take interest in the experiences of those who are involved in environmental interventions. I defend the use of empirical research in applied epistemology and introduce my research question.

8.1.1 Applied Epistemology and Empirical Research

Applied epistemology has been described as ‘bringing the tools of contemporary epistemology to bear on particular issues of social concern’ (Lackey, 2021: 3). Applied epistemology can also be said to follow a criticism of traditional epistemology, that says a theory of knowledge has too little practical relevance to real epistemic agents and their interactions.¹⁰³ Following this definition, epistemological questions that arise from environmental issues can be thought of as a matter of applied epistemology. I present three reasons for why applied epistemology should make use of empirical research about social reality: (i) to characterise a practical setting in which to apply normative principles, (ii) to identify appropriate objects of philosophical inquiry, and (iii) to inform the nature of normative principles.¹⁰⁴ I name each of these claims, in order to make it clear when I go on to reference them in data analysis. This is not intended to be an exhaustive list.

First, applied epistemology may concern applying normative principles and concepts to practical scenarios and empirical research can help fully characterise relevant practical scenarios.

¹⁰³ As noted by James Chase and David Coady in the introduction to the Routledge Handbook of Applied Epistemology, applied epistemology did not receive comparable attention to applied ethics in response to this sort of criticism: ‘the absence of an applied turn in epistemology to match the applied turn in ethics is especially puzzling when one considers how many issues of contemporary concern to the general public are epistemic in nature’ (2019: 4). This is the same puzzlement that we might have towards the question of why epistemology has (so far) had comparatively little to say about environmental issues, even though many of the difficulties we face in this context are inherently epistemic problems.

¹⁰⁴ Epistemologists already make reference to empirical research in applied epistemology, so this point itself should not be controversial (Anderson, 2011).

Providing Practical Scenarios: empirical research may help characterise real practical scenarios, in which normative principles and concepts could be applied.

Consider a method of applying normative principles that says that we only need true normative principles and an account of all the relevant non-normative facts; we can then reach a conclusion about the practical problem (Tännsjö, 2013: 4). If this is right, empirical research methods can provide access to relevant non-normative facts, which might include features of our social reality, e.g. relationships, organisations, communications, and general perceptions and shared understandings.

Second, empirical research might reveal what is of importance to the epistemic lives of real agents and thereby justify the direction of epistemological inquiry.

Directing Inquiry: empirical research may be used to direct epistemological inquiry.

If some epistemic problems are more commonly experienced than others or there is a shared understanding of their importance, this could justify a greater emphasis on these problems in our inquiry. We can gain an understanding of what problems are more commonly experienced or considered particularly important using empirical research methods.

Third, empirical research could reveal something about normative reality itself.

Informing the Normative: empirical research may reveal something about the nature of normative principles and concepts.

This idea is motivated by sub-disciplines such as social epistemology and non-ideal epistemology. Non-ideal epistemology criticises the abstraction of epistemic agents and their environments, where doing so adversely affects epistemic theorising (McKenna, 2023). Sometimes arriving at an understanding of real epistemic agents and their environments will impact the outcome of our theorising about normative principles. In other words, detailed accounts of social reality and the experiences of epistemic agents

may inform epistemic norms themselves, rather than just the setting in which they are applied.

In the following section, I defend an interdisciplinary approach to empirical research in epistemology.

8.1.2 Interdisciplinary Approach

As I have argued in Section 8.1.1, empirical research can be useful to applied epistemology, but the question of why we need an interdisciplinary approach still stands. One might object that, although epistemologists need empirical research to construct and answer applied epistemological questions, there is no need for a philosophy project that makes use of empirical research methods within it because appropriate reference to the findings of social research, for example, would be sufficient.¹⁰⁵

The main response to this challenge is that there is something to be gained from empirical research carried out from the perspective of philosophical concerns because it will specifically serve the end of constructing and answering applied epistemological questions. The researcher's position inevitably influences the research project at every stage, i.e. forming questions, designing research, collecting data, and analysing data. The involvement of a philosophy researcher will therefore guide the study to reflect a specific set of concerns and interests. However, given that researcher's influence is inevitable, it is also important to be aware of and mitigate any disadvantages (Hammersley and Gomm, 1997). The advantage, then, is that research can be guided towards philosophically interesting features of the world, but the risk is that, if the researcher is not appropriately aware of their position, the expectations or values present in the philosophical literature will impose too much on the research process – in other words, there is a potential to be

¹⁰⁵ We might worry that to deny that philosophy can make use of empirical methods carries restrictive assumptions about what philosophy is or should be. This objection seems to make the following assumptions, noted by Nora Hangel and Christopher ChoGlueck (2023) in their defence of qualitative research methods in Philosophy of Science. The first assumption is that we can draw a clear line between disciplines like philosophy and sociology, or 'between the "theorists" and the "fact finders"' (2023: 34). The second assumption undermines the work of philosophers who pursue research outside the boundary of traditional philosophical research methods, by denying that philosophical methodology can change over time (2023: 34).

influenced by disciplinary bias towards assumptions operative in philosophy that participants might disagree with.¹⁰⁶

The research project I have undertaken, and that I will go on to motivate, could fall within the sub-discipline of experimental philosophy, which uses empirical tools to make progress on philosophical questions, but there are some notable differences. At present, most experimental studies in philosophy use quantitative research methods, whereas as I will make use of qualitative research methods.¹⁰⁷ In addition to this, my research does not directly ask participants about concepts at the centre of epistemological debates, such as the nature of knowledge and understanding.¹⁰⁸ However, this research project might be considered experimental philosophy, in virtue of its aim to make progress on philosophical questions using empirical methods.

8.1.3 Research Question

The starting point of this research is the concern that my approach to developing an epistemology for the environment is distant from peoples' experiences and that the epistemic problem that I have chosen to focus on is not the only, or most important, epistemic problem we face in understanding the environment. In response to this concern, I suggest the following broad research question:

How do people experience epistemic factors as enabling environmental intervention or presenting barriers to environmental intervention?

This question asks how those directly involved in environmental interventions experience epistemic factors that either positively or negatively interact with these interventions. This is a deliberately broad question that is directed at a possibly wide range of experiences. Although I have focused on epistemic problems in setting up the relevance of this research,

¹⁰⁶ Disciplinary bias is produced by how the researcher's discipline ordinarily answers questions and might result in a tendency to overlook assumptions that participants might want to disagree with (Knott et al., 2022: 6).

¹⁰⁷ Despite not being commonplace, qualitative methods have great potential to contribute to experimental philosophy (Andow, 2016).

¹⁰⁸ One example of a research project that directly concerns participants' judgements about knowing how and understanding why is the experimental study by J. Adam Carter, Duncan Pritchard, and Joshua Shepherd (2019) that tests willingness to attribute knowing how and understanding why and whether these attributions suggest compatibility with epistemic luck.

it is important not to restrict the research question to negative factors, in order to arrive at a more comprehensive picture.

The purpose of this research is to understand how epistemic factors are experienced. The primary purpose of this research might be considered contextual or descriptive; contextual research can be defined as research ‘concerned with identifying what exists in the social world and the way it manifests itself’ (Ritchie and Ormston, 2014: 31). Descriptive research ‘presents a picture of the specific details of a situation, social setting, or relationship’ (Neuman, 2014: 38). My research question most directly aims at a description of how experiences of epistemic factors manifest in a particular case.

Alongside these descriptive aims, this research has an exploratory purpose. The goal of exploratory research is to ‘formulate more precise questions that we can address in future research’ (Neuman, 2014: 38). Research may be exploratory when its subject is new and under-researched, but research can also be exploratory when it contributes to the question of whether a certain approach or type of research project is feasible. Given that this role of empirical research in applied epistemology (especially the qualitative research methods I will be using) is not well established by other similar research projects, there are important questions about how best to carry out this research and whether the results are useful in the way that I have argued they are in Section 8.1.1. Underlying the research question I have specified, there are important questions about whether my approach is effective and whether this research has revealed new research questions or possible lessons for how to conduct future research.

8.2 Philosophy and Qualitative/ Quantitative Research Methods

Social researchers draw upon a range of philosophical ideas in their approach to researching social reality, including ontological considerations, epistemological considerations and assumptions about the relationship between theory and research. Some think we do not need to endorse controversial philosophical ideas in order to make use of, for example, qualitative research methods in philosophy (Andow, 2016: 1138). However, some philosophical considerations influence the research design and guide the process of

data collection and analysis in ways that do suggest some form of endorsement. This section explains how philosophical commitments influence my approach to the research question and addresses possible concerns about whether they are consistent with the rest of this thesis.

8.2.1 Ontological Position

Ontological commitments inform research questions and the way that research is carried out. Qualitative researchers should pay particular attention to the social ontology or, in other words, their views towards the nature of social reality (Pernecky, 2016: 140). A central question for social researchers concerns whether to take an objectivist or constructionist stance on the social entities they study (Clark et al., 2021: 27). Very broadly, one might say that objectivism considers social entities, e.g. organisations and the rules that govern them, to have an external reality that is independent from the perception and influence of the social agents that interact with them, e.g. members of the organisation that are subject to its rules (Clark et al., 2021: 28). In contrast, constructionism claims that social entities are not independent from social agents in this way, as they are always being modified by them (Clark et al., 2021).

Constructionism is normally associated with qualitative research. It is a complex term that can include many ideas I have not expanded upon in the above definition. I do not argue for a version of constructionism in this thesis (as you might expect from a philosophy thesis informed by this ontological position) and I do not myself explicitly commit to a constructionist view (in the way that you might expect from a methods section of a qualitative research study). Instead, I highlight aspects of my research question and research design that are guided by constructionist ontological commitments and, in the interest of consistency, I clarify that these commitments are not at odds with the philosophical commitments in this thesis.

A constructionist stance informs my research questions, which take interest in *shared meanings* about what factors influenced an environmental intervention and shared meanings about the role of epistemic factors in particular. Peter L. Berger and Thomas Luckmann argue in *The Social Construction of Reality* that social realities are intersubjective, in that they are shared with others and interpreted with others, and an

understanding of this social reality concerns itself with ‘what people “know” as “reality” in their everyday, non- or pre-theoretical lives’ (2011: 27). My focus on shared meaning contrasts other possible research questions construed along objectivist lines, e.g. research questions that aim to find out how features of institutions or regulations may have constrained human behaviour. This choice implies shared meanings are worthy of in-depth study and it is in this sense that my research at least carries the constructionist commitment that the meanings we form and share together have a bearing on social reality or form an important part of the nature of social reality.

Another way that this research leans towards constructionism is that my questions will not assume that the environmental intervention discussed (the Nearshore Trawling Byelaw) is the right outcome or indeed that any intervention is right or wrong; this research instead concerns people’s experience of epistemic factors in enabling, slowing down or preventing what they perceive to be the right thing to do. We can start research with the assumption that there is a single, objective reality or we can also begin with another assumption that actions and activities can depend on the perspective (Yin, 2016: 16). Given that my research concerns how epistemic factors are experienced relative to perspective, rather than how they are experienced relative to objective reality about what is a good outcome, my research begins with the assumption of multiple realities.

I acknowledge the utility in this assumption of multiple realities, as an effective way to carry out a qualitative research study but, importantly, this does not mean that multiple realities determine normative reality or that there is no matter of fact about what is a good outcome.¹⁰⁹ With this clarification, my research only assumes that people conceptualise reality differently and that these distinct perspectives are worthy of study. This does not seem to be a particularly controversial ontological assumption.

As discussed in Section 8.1.2, this research does not concern the nature of knowledge and understanding or attributions of these epistemic states and, therefore, by taking a constructionist stance towards some of the objects studied, e.g. organisational structure and

¹⁰⁹ Both of these claims may conflict with some externalist views about what it means to successfully understand (or fail to understand), argued for in Chapter 3, and the commitment to the importance of objective normative reasons, argued for in Chapter 2.

shared meaning, I do not commit to the claim that social agents have a role in determining epistemic states like knowledge and understanding.

8.2.2 Epistemological Position

Social researchers should also consider epistemological stances on what kind of knowledge is acceptable in social science, as this too will inform their research. Writing about the epistemological stance of a qualitative research study, within a thesis that is primarily epistemological inquiry, is difficult to navigate. This section must highlight the epistemological assumptions that inform this research and, again, address any concerns about whether they are consistent with my epistemological commitments.

A central issue for social science is whether the researcher should attempt to understand social reality using the same approach as the natural sciences (Clark et al., 2021: 23). Positivism, in this context, is associated with the view that the methods of natural science should be applied in social science. Although this term incorporates many ideas, positivism has two central theses: (i) we can explain the way people act with objective facts about them and (ii) the explanations that we are looking for are, like in the natural sciences, causal laws (Bevir and Rhodes, 2016: 4). In contrast, interpretivism is associated with the view that social reality cannot or should not be studied in the same way as the natural sciences.¹¹⁰ Like positivism, interpretivism is a broad term intended to capture various approaches. This section clarifies the sense in which my research takes an interpretivist stance.

One interpretivist claim that guides my research is that the study of human beings should be contextualised, rather than in search of ahistorical regularities (Bevir and Blakely, 2018: 9; Yanow, 2015: 8, 2007: 111). In this sense, the interpretivist approach distances from methods in the natural sciences, which aim to identify causal relations and uncover general laws. In contrast, the interpretivist stance suggests an answer to the research question specified in Section. 8.1 demands detailed and contextualised understanding, i.e. an in-depth study of a specific case. This approach assumes that an experience of epistemic factors cannot be abstracted from a particular context and, to understand them, we must

¹¹⁰ Interpretivism draws upon a number of traditions, including phenomenology, hermeneutics, symbolic interactionism and Max Weber's concept *Verstehen* (Clark et al., 2021: 24-27).

seek in-depth descriptions that situate them within a particular context. This interpretivist stance will inform a research strategy that prioritises the opportunity for a rich descriptions.¹¹¹

Another central feature of interpretivism is the emphasis on subjective and shared meanings (Yanow, 2007: 111). Interpretivism sees meaning as essential to human behaviour and therefore seeks an understanding of the way that people make sense of their experiences (Bevir and Blakely, 2018). This view on human endeavour reveals several layers of interpretation, in which the researcher themselves is situated: the initial interpretation of experience by the participant, potential second-hand interpretation provided by participants (i.e. their interpretation of others' interpretation of experience), the interpretation of the researcher in data analysis, and the interpretation of those who read the research report (Yanow, 2015: 20-21). The interpretivist approach recognises that we cannot escape the influence of the researcher, raising the importance of the researcher's reflective awareness about their own position. While the interpretivist's focus on shared meaning need not deny that a positive approach can provide *any* useful data about human behaviour, interpretivism does imply that a positivist approach cannot *fully* understand human behaviour.¹¹²

Interpretivism is not in conflict with any part of my own work in epistemology, at least insofar as it justifies my approach to the research question. To recognise that certain research methods are better at arriving at an understanding of certain objects, such as human experiences, is consistent with my account of understanding and view of epistemic normativity. My approach simply implies that it is important to understand social reality via an attempt to access to peoples' subjective perspectives and that we can understand epistemic problems by acquiring information about peoples' experiences.

¹¹¹ Clifford Geertz (1973) argues that the aim of qualitative research, specifically ethnographic research is to uncover 'thick description', conceived as layers upon layers of symbolic meaning that lies within a social reality.

¹¹² Susan Hodgett and Séverine Deneulin (2009) suggest that there is no clear distinction between quantitative and qualitative research because both require interpretation within a narrative.

8.2.3 Theory and Research

Another important consideration for social researchers is the relationship between theory and research (Clark et al., 2021: 18; Neuman, 2014: 88). The aims of this research are mostly inductive, rather than deductive, meaning that the theory emerges from the data, rather than data as a test for a hypothesis. However, to characterise this research as entirely inductive is misleading. Theory enters the picture, to some extent, before data collection as a starting point; as I have shown in Section 8.1, this research is guided by philosophical literature. That said, the research question aims at an understanding of peoples' experiences and therefore claims in philosophy should not impose too much on the data collected. For example, I do not assume that epistemic factors have a notable impact on environmental interventions and the research process is guided by the participants concerns. In this sense, the aims of this research are inductive.

8.2.4 Qualitative vs Quantitative

This sub-section clarifies the distinction between qualitative and quantitative research methods and begins to develop an argument for my focus on qualitative research methods. Both quantitative and qualitative data is collected and integrated in mixed methods research (Creswell and Plano Clark, 2018: 6). Although primarily qualitative, some of the data I collect is quantitative, so this sub-section will also highlight the benefit of mixed methods research.

Qualitative research can be inclusive of a collection of methods, methodologies and philosophical perspectives and is therefore difficult to define precisely. Generally, qualitative research is thought to prioritise words, over quantification (Denzin and Lincoln, 2005: 6). We could think of qualitative research as:

a type of research strategy that emphasises words, images, and objects when collecting and analysing data. It is broadly inductivist, constructionist, and interpretivist, but it can take wide variety of forms; qualitative researchers do not always subscribe to all three of these features (Clark et al., 2021: 350).

As I have already stated in this section, this research takes a constructionist view of social reality and an interpretivist stance in a number of respects. The research question also

implies inductive research aims. For these reasons, qualitative research methods are most obviously appropriate and the research design reflects the need to focus mostly on qualitative data.

However, using different approaches can help to make sense of social reality. Quantitative research collects and analyses numerical data and it is generally associated with a deductive, positivist and objectivist approach (Clark et al., 2021: 142). Using both quantitative and qualitative research methods can be an effective and natural way to approach a research question and design a study. One of the advantages to mixed methods research is that a researcher may be able to capture strengths of both approaches and minimise their weaknesses (Creswell and Plano Clark, 2018: 13).¹¹³ As I will explain in Section 8.5.1, I use quantitative research methods to enhance the qualitative approach. In what follows, I consider different research designs and further justify my research strategy.

8.3 Research Design

A research design is a specific framework within which the collection and analysis of data takes place (Clark et al., 2021: 39). The research design should reflect the priorities of the research in question. This section considers possible research designs and justifies my use of the case study design.

Experimental designs are used to isolate and manipulate a variable to test its influence on other variables (Clark et al., 2021: 43; Neuman, 2014: 47-48). This involves a high degree of intervention from the researcher, which can take place in real setting (field experiment) or a highly controlled setting (laboratory experiment) (Clark et al., 2021: 43). There may be forms of experimental design that make progress on research questions similar to mine. For example, we could seek to test causal relationships that hold between epistemic barriers and environmental interventions. However, my research question concerns peoples' experience of epistemic factors, rather than testing a theory about a causal

¹¹³ Note that there is disagreement about whether this amounts to (1) mixed methods that treats quantitative methods as central and qualitative as preliminary, (2) mixed methods that treats both as important, but equipped to answer different questions about the same topic, or (3) mixed methods as a form of triangulation, by which we can validate findings by approaching social reality from different perspectives (Ritchie et al., 2013: 40-41).

relationship between specific epistemic factors and environmental interventions. Additionally, my research question concerns how epistemic factors are experienced in real world scenarios and it may be counterproductive to attempt to gain an insight into this by making observations about a significantly manipulated scenario. Finally, experimental designs tend to work well for explanatory purposes, whereas, as outlined in Section 8.1.3, my research aims are primarily descriptive.

Another possibility might have been a cross-sectional research design, which requires collecting (quantifiable) data from a number of cases that represent one point in time and then looking for regularities in this data (Clark et al., 2021: 50; Neuman, 2014: 44). I could have achieved this with a multiple-choice self-completion questionnaire. There may be some benefits to self-completed questionnaire, such as convenience and minimising the impact of the interviewer's presence (Clark et al., 2021: 212), but this design is likely to be too constrained, given the descriptive aims of this research. My research question seeks an in-depth exploration of the experiences of epistemic factors in environmental intervention, which means this research benefits from giving participants the opportunity to expand on their experiences with detailed accounts. Even the opportunity for participants to write answers, i.e. open-ended questions that do not provide fixed answers, would not solve this issue, as participants are unlikely to have dedicated as much time and effort writing their response or to have understood the nuances in the questions.

I could have carried out a structured interview, where the questions would be identical and mostly close-ended, i.e. offer a fixed set of answers (Clark et al., 2021: 191). But, close-ended questions do not offer flexibility. As argued in Section 8.1.2, data collection must be guided towards what a participant takes to be of interest, in order to minimise the risk of disciplinary bias. This requires more flexibility and for the researcher to be in direct contact with the interviewee. The opportunity for follow up questions is also important to this study, to encourage participants to expand on and to clarify certain points. The need for the participant to expand on their answers and the need for the researcher to be able to ask follow-up questions makes a cross-sectional research design inappropriate in this context.

Unlike cross-sectional designs, the longitudinal research design collects data at multiple points over a certain time period (Clark et al., 2021: 54; Neuman, 2014: 44). Whilst it would be interesting to examine changes in the impact of epistemic factors over time,

implementing a longitudinal research project can be costly and time-consuming. This research design falls outside the practical limitations for this doctoral study. Another practically difficult design would be a comparative design, which can use either quantitative or qualitative methods to research phenomena in two or more settings, e.g. in different nations or regions. Although such results might bring clarity to my research question, the time investment and financial costs of this approach are not practical in the context of this thesis with a local subject.

The case study research design is most appropriate to answer my research question. Case study research ‘intensively investigates one or a small set of cases, focusing on many details within each case and the context. In short, it examines both details of each case’s internal features as well as the surrounding situation’ (Neuman, 2014: 42). Case study research can involve analysing one ‘case’ in depth, where case often centres around a location – as is true of this case study (Clark et al., 2021: 59). This approach aims at a rich understanding of the complex and specific details of the case in question (Stake, 1995).

The case study design is appropriate because, as stated above, the concerns raised in applied epistemology look at epistemic problems that arise with real epistemic agents in a particular practical scenario. The case I have selected centres around the Nearshore Trawling Byelaw, the ongoing kelp recovery project, and epistemic factors that were experienced by participants in relation to this fisheries management. I defend my selection of this case in the next section.

8.4 The Nearshore Trawling Byelaw

The aim of this case study is to acquire an understanding of how epistemic factors were experienced in relation to the Nearshore Trawling Byelaw and the Sussex Kelp Recovery Project (SKRP). This section provides context for this case study and justifies its selection.

8.4.1 What is the Nearshore Trawling Byelaw?

The Nearshore Trawling Byelaw was introduced along the Sussex coast to manage trawling, a fishing practice where large nets are used to catch fish near the seabed. The

practice of trawling can damage seabed habitats. This fishing activity has contributed to the loss of 96% of the kelp forest between Selsey and Shoreham-by-Sea (Sussex Kelp Recovery Project, 2024: 15). The Nearshore Trawling Byelaw aims to provide an opportunity for these habitats to recover, by banning the practice of trawling in large areas along the Sussex coast. The Sussex Inshore Fisheries and Conservation Authority (IFCA) made the case for this Byelaw.¹¹⁴

Byelaws are defined as ‘local laws made by a local council under an enabling power contained in a public general act or a local act requiring something to be done – or not done – in a specified area’ (Association of IFCA, n.d.b). To develop a byelaw, IFCA works with the Marine Management Organisation (MMO) and Department of Environment, Food and Rural Affairs (Defra). The process is as follows.¹¹⁵ In Stage 1, IFCA may propose a new byelaw that is considered necessary. Stage 2 includes both informal evidence-gathering (e.g. drop ins, newsletters, community meetings and social media) and formal evidence-gathering. During formal evidence gathering, an Impact Assessment is developed. A formal consultation period also takes place and IFCA considers objections. In Stage 3, a byelaw (along with the Impact Assessment and other relevant documentation) is submitted to the MMO, who works with IFCA to finalise a version of the byelaw, following legal and policy advice. The MMO submits the final version to Defra’s Secretary of State, recommending to approve or reject the byelaw.

By focussing on this case study, I aim to gain an understanding of how epistemic factors are experienced by those involved in the trawling ban in Sussex and the ongoing Sussex Kelp Recovery Project.

8.4.2 What is the Help Our Kelp Campaign and the Sussex Kelp Recovery Project?

A group of organisations came together to support the *Help Our Kelp* campaign, inspired by Big Wave Production’s 2019 film. After the banning of trawling in Spring 2021, the

¹¹⁴ There are 10 IFCAs in the UK. IFCAs were created by the Marine and Coastal Access Act (2009) in the place of Sea Fisheries Committees (SFCs). They took on the responsibilities of former SFCs, as well as the protection of marine ecosystems marine conservation zones.

¹¹⁵ A more detailed description of this process can be found on the Association of IFCA’s website (Association of IFCA, n.d.b)

partners of the *Help Our Kelp* campaign joined with other organisations, including Sussex IFCA, to form the Sussex Kelp Recovery Project partnership. The stated mission is ‘To champion, study and facilitate the recovery of Sussex kelp and other essential fish habitats, through progressive, coherent and collaborative action’ (Sussex Kelp Recovery Project, n.d.b).

8.4.3 Selection of Case Study

Different case studies can be distinguished according to the rationales for their selection and a case study can be selected for a combination of reasons. One reason to select the Nearshore Trawling Byelaw as a case study is that it is *representative* in relevant ways; representative case studies ‘capture the circumstances and conditions of an everyday or commonplace situation’ (Yin, 2009: 48). This introduction of the Nearshore Trawling Byelaw in some ways tells a familiar story; it involves the introduction of an environmental intervention that is subject to debate, due to its conflict with human interests. However, in other ways the Nearshore Trawling Byelaw departs from the typical case. This was a novel and significant piece of legislation that attracted national attention. In this sense, it has elements of a *unique* case study (Clark et al., 2021: 60). Because of the way that the byelaw departs from a typical case, it is of interest to acquire an in depth understanding of whether epistemic factors had any role in the process of campaigning for or implementing the byelaw, as well the attention that this byelaw received.

8.5 Interviews

This section justifies the use of semi-structured interviews, rather than other qualitative research methods, and describes preparation for these interviews and the data collection process.

8.5.1 Research Methods

Ethnography or participant observation involves the researcher immersing themselves in a setting, systematically observing behaviour, keeping a record and producing a write up based on observations (Clark et al., 2021: 392; McNaughton Nicholls et al., 2014: 243; Neuman, 2014: 51). This form of research can be overt, where the researcher makes

themselves known to those observed, or covert, where the researcher conceals their identity as a researcher (Clark et al., 2021: 394; McNaughton Nicholls et al., 2014: 247). Participant observation allows the researcher to collect data that might go beyond limitations of verbal accounts, such as on topics that are difficult to describe accurately (McNaughton Nicholls et al., 2014: 245). This research method was not appropriate in this case study. This case study involves events that happened in the past and cannot be observed in the present. Moreover, this method lacks the structure needed to focus the research process on my research question.

Another possibility might have been to conduct a focus group, which involves group discussion between a small number of participants (Clark et al., 2021: 453; Neuman, 2014: 471). While this would allow the researcher to guide the discussion to some degree, the researcher should not involve themselves or divert the conversation too much. This method therefore may not lead to discussion of the specific research concerns. There may also be practical reasons that this method is not appropriate, such as possible difficulties in arranging focus groups with multiple participants that may have busy schedules.

Another possible qualitative research method is to carry out interviews. Given that this project has a clear focus on epistemic factors, structure is needed to guide the data collection process accordingly. This means that an *unstructured* interview, which mirrors conversation and includes only a few prompts from the researcher, might lack the needed structure (Clark et al., 2021: 426). On the other hand, I argued above that the interview process requires flexibility and the opportunity for follow up questions, which makes a *structured* interview inappropriate. The *semi-structured* interview strikes a balance between these considerations.

The semi-structured interview is ‘flexible, allowing new questions to be brought forward during the interview as a consequence of what the interviewees have said’ (Ruslin et al., 2022: 24). The semi-structured interview allows for a more in-depth explanation than structured interviews and can be guided towards the concerns of participants. However, this form of qualitative interview still ‘holds its direction’, and can be carefully prepared beforehand, meaning that the data collection process will still speak to the specific concerns of my research question (Ruslin et al., 2022: 29). Because the semi-structured interview allows the researcher to guide discussion towards the research question, but not

to an extent that undermines flexibility, it is the most fitting research method. By allowing follow-up questions to vary according to participants' answers and by spending appropriate time on each question according to participant answers, it's possible to collect data that was not anticipated by the researcher and that might have been missed in a highly structured interview. This is especially important, given the exploratory aims of this research.

8.5.2 Planning and Organising Information Gathering

When planning the interviews, the anticipated participant groups included employees and volunteers in conservation groups, academics in marine biology and conservation, employees at media/production companies, environmental activists (i.e. individuals who assisted the *Help Our Kelp* campaign and who are not associated with any other participant groups). This 'purposive sampling' method aims to contact the participants well-placed to provide insight on the trawling ban. This is often thought to be the most appropriate sampling method for a small set of interviews that does not aspire to be representative (Knott et al., 2022: 1). I identified initial participants using contact details for organisations and publicly available contact details. Then, I used snowball sampling to identify further participants, by asking those interviewed for contacts that might be willing to take part in this research (Neuman, 2014: 274-275).

During the pandemic, social researchers had to collect data online or over the phone to abide by social distancing rules (Lobe et al., 2020: 2; Oliffe et al., 2021: 1-2). While social distancing rules have not been in place during the data collection process, the impact of the recent COVID-19 pandemic has affected the way this research was carried out. During the pandemic, video conferencing platforms became widespread in many workplaces and working from home is now more common than it used to be. The widespread use of video conferencing has the potential to make data collection much easier; it is easier to organise online meetings that do not involve travel time or expense and can be fitted into busy schedules, which may increase participation (Gray et al., 2020: 1297; Lobe et al., 2020: 2). However, the use of this technology carries some disadvantages, e.g. potentially excluding those who do not know how to use video conferencing software or do not have access to the internet, speakers, or are lacking a quiet environment etc. (Lobe et al., 2020: 2). Another disadvantage of online interviews is the potential disruption to the interview

process, due to connection issues or distracting environments that are unpredictable and outside of the researcher's control (Olliffe et al., 2021: 4-5). Online interviewing is not simply a replacement for in person interviews. Carrying out interviews online has advantages and disadvantages that should be carefully considered by the researcher. This research project uses both video calls (Microsoft Teams) and audio calls to carry out research interviews in a way that proved convenient for participants.

I aimed to interview 10-30 participants (alongside a review of relevant documents), but due to difficulty arranging interviews and the context of the pandemic, I interviewed 7 participants. All participants were either involved in the *Help Our Kelp* campaign, the Nearshore Trawling Byelaw, or the ongoing Sussex Kelp Recovery Project. The result of interviewing fewer participants than anticipated is that the interview data plays a lesser role in this case study than it might have otherwise. The implications of this are discussed in the section on reflections.

8.5.3 Interview Questions

The interview schedule is semi-structured. I constructed a series of template questions that were asked to each participant but, as I will now describe, the course of the interview is intended to reflect what the participant takes to be important. Follow up questions varied significantly.

Semi-structured interviews centre around a set of topics and questions to guide the interaction between the researcher and the participants. One usual approach is to open with 'relatively easy, open-ended questions. Opening questions should be related to the research topic but broad and easy to answer, so that they help to ease the participant into conversation' (Ruslin et al., 2022: 2). Opening questions also present an opportunity to ask questions about context (Yeo et al., 2014: 188). Following this structure, the first several questions of my interview schedule (in Appendix 1) aimed to understand the background on participants' involvement in the trawling ban (or the Sussex Kelp Recovery Project) and their views on the trawling ban. All changes to my initial set of questions were made after interview 4. Initially the first question read:

What is your understanding of how the Nearshore Byelaw/ trawling ban came about and what was the nature of your involvement?

I changed this question to:

How were you involved/ affected by the trawling ban (or the ongoing Sussex kelp recovery project)

This question was changed after I realised during the interview process that how the ban came about is not relevant to everyone's involvement and can take up a lot of time.

The second question asked participants for their views on the trawling ban. This question is important to determining what participants' perceived to enable, prevent or delay the right sort of intervention, in order to avoid asking questions that minimise opportunity for participants' disagreement. Initially, this question asked what each participant thought was the 'right decision' or 'best response', but after several interviews it was clear that this phrasing was unhelpful. Researchers should ask questions that are clear and do not leave participants uncertain about what information the question is aimed at (Yeo et al., 2014: 192). This question was changed to ask whether banning trawling was the right thing to do:

Do you think that banning trawling was the right thing to do? Are there any advantages/ disadvantages to the trawling ban?

This was a more direct way to ask about the participant's views on the trawling ban.

Depending on what each participant perceived to be the best outcome, I then asked a general question about what may have either enabled, prevented or delayed the best outcome and whether anything got in the way. This interview was semi-structured and I expanded upon the responses each participant gave, so the phrasing of this question changed to reflect participant's responses but roughly followed the form:

What do you think enabled/ prevented/ delayed the [best outcome]?

This general question does not specifically mention knowledge or understanding and provides an opportunity for the participant to guide the interview towards what they found

the most important. It also provides an opportunity to mention important factors that are not related to knowledge and understanding.

Question 4 asks whether knowledge and understanding (or lack of knowledge and understanding) had a role to play in either enabling or getting in the way of the best outcome according to the respondent. My assumption is that the terms knowledge and understanding, or a failure to arrive at knowledge and understanding, will generally track participants' perception of relevant epistemic factors. Asking the question in this form allowed respondents to outline their experience in much more detail, and what they perceived to be possible solutions. Some participants found the question too general and it was useful to ask in two parts, asking whether knowledge and understanding of the kelp forests had a role to play and then asking whether knowledge and understanding of the byelaw and surrounding issues had a role to play.

The last several questions draw upon specific features of the trawling ban and its impact. I asked about visibility of kelp, complexity of the issue, the way that this issue was covered in the media, and conflict with human interests. The exact wording of each question can be found in the interview schedule (Appendix 1). Here is one example to demonstrate the form of these questions:

What role, if any, did the lack of visibility of sea kelp play in delaying or preventing the best outcome?

Please rate the contribution 0 – 5, where 0 is not at all and 5 is significant contribution. follow up: any comments on this? How did this play a role? How was this overcome?

These questions are included in order to bring out any further epistemic factors that I did not want to miss, by not directly asking about them.

However, it was also important to make sure that participants did not feel obligated to comment on these issues if they do not consider them significant (Yeo et al., 2014: 191-192). To address this issue in a consistent way for each participant, I introduced these questions by asking participants to rate the significance of these issues on a Likert scale 0-

5 where 0 is not significant and 5 is very significant. I only asked follow-up questions where participants gave a rating of 4 and 5, (although some participants did respond with more information regardless of their rating). In some cases, the use of a pre-determined scale, e.g. in a survey, may constrain the answers of participants and prioritise breadth of data over in depth descriptions (Knott et al., 2022: 1). In this case, however, the scale is used to direct the conversation towards topics that are of most significance to the participants.

8.6 Approach to Analysis

I used thematic analysis to analyse this data. Virginia Braun and Victoria Clarke (2022) identify 6 phases of thematic analysis, which I broadly follow. The six phases are: (1) familiarisation with the data set, (2) coding, (3) generating initial themes, (4) developing and reviewing themes, (5) refining, defining and naming themes, and (6) writing up (Braun and Clarke, 2022: 35-36).

I first familiarised myself with the data. To do this, I immersed myself in the data as I carried out interviews myself, I transcribed interviews, and I listened back to interviews while reading transcription. I also wrote notes on my initial thoughts, added these notes as annotations on NVivo (a qualitative data analysis software), and I read through interview transcriptions multiple times.

Then, I began the coding process by highlighting parts of the text that are of potential relevance to what may have enabled, delayed or prevented what a participant takes to be a good outcome and all sections that reference epistemic concepts. I consider three signs that a passage references epistemic factors. First, some of my questions specifically refer to knowledge and understanding so, unless participants change the subject, their responses are about what they take to be epistemic factors. Second, a passage may reference epistemic factors if it uses epistemic vocabulary, i.e. words that signify epistemic concepts such knowledge, understanding, belief, evidence etc. Some words indicate something epistemically 'positive', i.e. the acquiring of an epistemic relation (remembering/ raising awareness/ coming to understand), and other words imply something epistemically 'negative', such as losing, failing to obtain knowledge and understanding or otherwise

getting things wrong (forgetting/ mistake/ distraction). Third, the passage describes something indicated by this sort of epistemic vocabulary. For example, “to remember” might be described as “sticking in one’s mind” or a mistake might be described in detail without actual use of the word “misunderstanding” or “false belief”. This coding process was initially carried out by highlighting and annotating a word document and then I used NVivo to make use of coding tools, text search and word cloud functions.

I identified some initial themes. Given the aims of this research, it is important that the themes were organised to reveal epistemic factors. While developing themes, researchers should reflect on the central organising concept of each theme, i.e. the potentially implicit idea that unifies the meaning of a theme, and this may lead to collapse of several themes or division of one theme into multiple (Braun and Clarke, 2022: 35). For example, the data suggests that the *Help Our Kelp* film (released during IFCA’s consultation process) is important in this case study, but several central organising concepts were revealed relating to the film, such as awareness, trust, simplicity, and understanding value. Each of these had distinct epistemic relevance. When I returned to the wider data set and grey literature, these epistemic themes were sometimes mentioned in other contexts beyond the film. In this case, the importance of the film was divided into multiple themes.

Paying attention to epistemic and non-epistemic themes that emerge from the data will answer to the central research question detailed in Section 8.1.3, which asks how epistemic factors were experienced. However, the arguments in this chapter prompt further analysis. Firstly, I began this chapter with the concern that the first 7 chapters of this thesis could be out of touch with the experiences of real epistemic agents. The research question is intended to find out about these experiences, but the analysis of the data should make comparisons concerning the discussion in the first 7 chapters (as well as theory that informed this discussion) and the epistemic themes revealed by the data. Secondly, I argued in Section 8.1.1 that empirical research is useful to applied epistemology in at least the following ways: (i) providing practical scenarios, (ii) directing inquiry and (iii) informing the normative. We should also keep in mind in analysis whether the data acquired is useful with respect to any of these aims. Thirdly, the exploratory aims of this research concern whether a qualitative research project like this is feasible in applied epistemology. The answer to whether this research is feasible may be informed by whether the findings offer any insights with respect to the above two questions.

8.7 Ethics

In carrying out this research study, I referred to the University of East Anglia's 'Guidelines on Good Practice in Research', 'Research Ethics Policy', and 'Research Data Management Policy', and the 'British Sociological Association Statement of Ethical Practice' to help with the practical experience of interviewing (British Sociological Association, 2017; Members of the Research Data Management Project Team and Scott, 2021; The University's Research Integrity Manager, Research and Innovation Services and University Research Ethics Committee, 2021, 2022). The purpose of the BSA's statement is to bring awareness to Sociologists about ethical issues that may arise during their research and I have chosen to refer to this statement, alongside UEA's Ethics guidance, as I am using some sociological research methods (British Sociological Association, 2017: 2). This research was approved by the university's Faculty of Arts and Humanities Research Ethics Subcommittee.

8.7.1 Professional Integrity

I considered professional integrity at each stage of the process. One issue that arose during the data collection process was that this research would be missing important data without participants from local fishing communities. It would have been preferable to carry out these research interviews in person. Despite attempts to make connections, this proved difficult to achieve within the given time constraints. As addressed in Chapter 9, in hindsight a different approach to planning, and teleological considerations raised by the pandemic, may have improved the overall data collection process.

The BSA ethical statement notes that researchers should be aware of what kind of project they are able to carry out and the boundaries of their competence (British Sociological Association, 2017: 4). Given the practical limitations and time constraints on this project, the best resolution is to recognise the limitation of the data collected in this case study. The data collected reflects the views and experiences of a group of participants associated with conservation groups, local authorities and filmmakers involved in the Sussex Kelp Recovery Project, as well as information available in publicly available documents, and does not expand to all those involved in or affected by the trawling ban, which would include a wider range of participant groups. This data set does not therefore provide a

comprehensive picture of views and experiences related to the trawling ban, but it does give some indication of a range of epistemic themes experienced by those involved in an environmental intervention, which reflects the initial aim of the research project.

8.7.2 Informed Consent

This research was carried out with informed consent of participants. No covert research was carried out. Consent was established by sharing a participant information sheet that detailed what was involved in taking part in the research project. Participants were informed of the kind of data that would be collected, including audio recording of the interviews, the manner in which data would be anonymised, means of storage, access, and the date information would be deleted. Participants were informed of how data would be analysed. Participants were made aware that they would receive a transcript of their interview and reserve the right to withdraw consent at any point before data analysis. Consent for participants involvement was acquired through a written consent form and verbal consent before recording was started. Consent forms also asked whether participants would like to receive feedback about the overall results of this study. Participant information sheet and consent forms are available in Appendix 2 and Appendix 3.

8.7.3 Anonymity

Anonymity of participants was protected as far as possible. During transcription of recorded interviews, participants were anonymised using pseudonyms (numbers from 1-7). For context, I have included the type of organisation that participants are associated with. Given that the case study in question is a notable event that involved a limited number of organisations, there is a risk that participants may be identified based on the information they choose to share about their involvement. Both the BSA Ethical Statement and the previously listed UEA documents stress the importance of anonymity wherever possible and the need to make sure participants are aware of how much anonymity will be possible. It was explained on the information sheet that no directly identifiable information (such as names, or names of the associated organisation) would be included, but due to the nature of this research, there is risk of identification. Participants were also informed that they are free to withdraw consent after reviewing transcripts. Transcripts and audio recordings of interviews are not published and shared only with project supervisors.

8.7.4 Safety of Participants

It is generally agreed that researchers should consider potential consequences of their research that could be harmful and to minimise the potentially harmful effects. Before data collection began, I considered the potential impact of normative concepts at the centre of some of my research, which includes the idea of moral responsibility and of praise and blame. My research does not aim to apply these concepts to any particular group or individual that was involved in or affected by the trawling ban; the aim is to identify and analyse epistemic factors that were experienced. However, given that these concepts are mentioned in my thesis as a whole and are part of the general discussion, this was communicated to participants in the information sheet to guard against any discomfort during the data collection process and to ensure informed consent for how the data would be analysed. This research did not involve groups that may be vulnerable due to age, disability, or physical and mental health. This research also did not include children.

8.7.5 Data Storage

It was also important to consider data storage during this research process. This research was carried out in line with UEA's 'Research Data Management Policy' (Members of the Research Data Management Project Team and Scott, 2021). Data collected includes audio recordings of 20-minute research interviews with participants and a transcription of each interview. Participants were informed that data would be stored on a UEA OneDrive and will be deleted within 3 years. Potential uses of data were shared with participants on the participant information sheet. Participants were informed of possible publication of findings in the future.

8.8 Conclusion

In this chapter, I have defended the use of empirical methods in the development of an epistemology for the environment and I have defended my approach to the case study research into the Nearshore Trawling Byelaw that I carried out to illustrate this. In Chapter 9, I present my findings about epistemic and non-epistemic themes that emerged from this

case study research. I also draw conclusions and share my reflections on the research process.

Chapter 9: Themes and Reflections

9.0 Introduction

Chapter 8 argued for the use of qualitative research in this thesis and defended my choices in the case study research I carried out on the Nearshore Trawling Byelaw. As explained in Section 8.6 on my approach to analysis, epistemic themes emerge from codes categorised as epistemic according to the rationale I have argued for. This chapter presents findings on epistemic themes and other important themes, by referencing the number of participants who discuss a given theme and including important quotes from transcripts and relevant documents. Section 9.1 discusses the following epistemic themes that emerged during this research: evidence, awareness, simplicity, hope, understanding value, trust and memory. In discussion of each of these themes, I reflect on the use of this data for applied epistemology. For a more complete picture of the case study, Section 9.2 presents my findings with respect to other important themes, objections to the byelaw and collaboration, and their relationship to epistemic themes. Section 9.3 draws some conclusions based on these findings, in light of questions outlined in the discussion of my approach to data analysis (in Chapter 8, Section 8.6). Section 9.4 shares some reflections on this research and potential future projects.

9.1 Epistemic Themes

This section discusses my findings with respect to the following epistemic themes: evidence, awareness, simplicity, hope, understanding value, trust and memory. In Chapter 8 Section 8.6, I outlined three signs that a passage refers to epistemic factors: (i) it is a response to a question that specifically asks about knowledge and understanding and there is no reason to think the participant has changed the subject, (ii) use of epistemic vocabulary, and (iii) description of phenomena indicated by epistemic vocabulary. Themes are categorised as epistemic when they emerge from codes that are epistemic for these reasons.

9.1.1 Evidence

An epistemic theme that emerges from the data is evidence. Evidence could be considered central to many prominent debates in epistemology and plays a key role in theorising, e.g. the evidentialist thesis that epistemic justification is determined by one's evidence (Conee and Feldman, 2004). Although what constitutes evidence may be debated, evidence for p can be thought of as basically synonymous with a good reason to believe p. This may include a range of entities, including belief or credence in propositions, as well as feelings and experiences that speak in favour of p (Conee and Feldman, 2004: 2). The theme of evidence emerged frequently in at least two contexts: evidence in the consultation process and the role of anecdotal evidence.

Evidence in the Consultation Process

Evidence is raised in the context of the consultation process. IFCA is responsible for the proposal of new byelaws like the Nearshore Trawling Byelaw and their involvement was, as you might expect, discussed by all participants (n=7). Many participants specifically mention the case for the byelaw made by Sussex IFCA and the work of individuals associated with Sussex IFCA as an enabling factor to the introduction of the Byelaw (n=5). As explained in Chapter 8, Section 8.4 IFCA is required to undertake an informal and formal stage of evidence-gathering prior to submitting a byelaw to the MMO.

Sussex IFCA produced an Impact Assessment of the Nearshore Trawling Byelaw, following informal public consultation (June 2018) and formal public consultation (September and October 2019), and its purpose is to evaluate costs and benefits of this management option in comparison to others (Sussex IFCA, 2020: 12). IFCA aims to update previous trawling restrictions in line with 'best available evidence' (Sussex IFCA, 2020: 1). The Impact Assessment references the Government's recommendation to take a precautionary approach to fisheries management, which means that a lack in scientific evidence does not count against conservation and management measures (Sussex IFCA, 2020: 23). The evidential standard that the Impact Assessment strives for is the 'best available evidence'.¹¹⁶

¹¹⁶ The SKRP mention that researchers face inherent challenges in seeking to acquire evidence about marine environments (Sussex Kelp Recovery Project, 2024: 25).

The evidence presented in the Impact Assessment takes a holistic approach. Following the Marine and Coastal Access Act 2009, the IFCA's were given more of a role in conservation than the Sea Fisheries Committees (SFCs) that they replaced.¹¹⁷ Sussex IFCA is required to review byelaws and the practice of trawling was reviewed as part of this process. Sussex IFCA took an ecosystem approach (following Government guidance) which 'balances ecological wellbeing with human and societal wellbeing, and aims to stop the degradation of underlying natural capital assets to ensure people can continue to enjoy the full range of services and benefits into future' (Sussex IFCA, 2020: 23). Sussex IFCA thereby uses the best available evidence to establish:

- (1) the impact of bottom trawling on habitats
- (2) type and extent of natural capital assets (aspects of nature that provide value to people) (Sussex IFCA, 2020: 23), their sensitivity (i.e. how vulnerable habitats were to pressures, measured by their resistance and capacity to recover) (Sussex IFCA, 2020: 31), diversity, ecosystem goods and services, and value.¹¹⁸
- (3) importance of Nearshore area
- (4) current natural capital assets condition and risks
- (5) impacts of management proposal.¹¹⁹

This was informed by habitat mapping and environmental mapping to determine sensitivity of habitats to different fishing activities and the value of those habitats:

It was a combination of the kind of environmental value and the sensitivity of the habitats... that along with where the historical kelp beds used to be, informed the basis for the Nearshore Trawling Byelaw (P5, Fisheries and Conservation Authority).

Note that, although possibility of kelp recovery is part of the evidence base for the byelaw and the primary focus of the campaign *Help Our Kelp*, the evidence for the byelaw pertains to the wider impact of trawling.

Participants made both positive and negative comments about the consultation process. Some reflected on the practical costs of the consultation process:

¹¹⁷ The Marine and Coastal Access Act (2009) lays out the duties of IFCA's (see UK Government, 2025). When IFCA's replaced SFCs, their duties extended to conservation (Association of IFCA, n.d.a).

¹¹⁸ Ecosystem services are defined as 'functions and products from nature that can be turned into human benefits with varying degrees of human input' (Sussex IFCA, 2020: 24).

¹¹⁹ This list can be found in the Impact Assessment (Sussex IFCA, 2020).

The level of evidence required, that the impact assessment produced by Sussex IFCA is extraordinarily good, but must have taken quite a considerable amount of resource and effort to compile (P1, Conservation Organisation).

I feel like the consultation processes take a long time, but we can't not have them otherwise people don't have an opportunity to contribute from different perspective. And I think particularly in the marine environment. It seems incredibly complex (P4, Conservation Organisation).

In the latter quote, the participant notes that the consultation process takes time, but also that its necessity for fairness, which is also expanded upon in the following quotes.

The complexity of it does mean it's very hard to do things quickly and do them well, with sufficient engagement of different people and, yeah, make it fair in terms of access and co-design and all that (P4, Conservation Organisation).

It would be wrong to say a consultative process got in the way. Because that's the whole point. So I don't think anything got in the way of it because the whole point of a consultation is to get peoples' opinions...the purpose is to make a measured decision. So, I would say in terms of the Byelaw, it has to have that process. So you can't look at it in terms of it getting in the way. It is the point of it (P7, Conservation Organisation).

So while the evidence-gathering process was recognised to take time, this was also considered to be a necessary part of an effective and fair byelaw development process.

Anecdotal Evidence

Participants emphasised the importance of anecdotal evidence. Different forms of evidence were used to support the byelaw, including reports from fishers and divers about the kelp loss (Sussex IFCA, 2020: 28). Some participants (n=3) mentioned divers recorded the kelp and witnessed its decline. This is also acknowledged on the Marine Conservation Society's (MCS) Sussex Byelaw timeline, where the key milestones begin with 'evidence of expansive kelp forests from historical diver surveys and anecdotal stories' (Marine Conservation Society, 2023: 13). The MCS's report on lessons and recommendations from the implementation of the Nearshore Trawling Byelaw, suggests utilising diverse evidence,

including citizen science data and anecdotal stories, as this approach ‘aligns with the holistic management approach and creates an evidence base which is inclusive innovative, and robust’ (Marine Conservation Society, n.d.: 31).

The term “anecdotal evidence” has many uses and in this context anecdotal evidence is often contrasted with scientific expertise.¹²⁰ For example, where anecdotal reports were cited in the Impact Assessment, the Impact Assessment states ‘although far from scientific, this evidence is indicative of the potential gains of habitat restoration’ (Sussex IFCA, 2020: 76). The MCS report references a similar contrast, between a narrow conception of evidence from ‘trained researchers gathering ‘facts’ about the world’ and alternative methods, e.g. citizen science data and anecdotal stories (Marine Conservation Society, n.d.: 31). The 2023 SKRP report also contrasts anecdotal and scientific evidence, where anecdotal evidence includes observations and data from local sea users (e.g. divers, fishers, beach walkers, and sailors) (Sussex Kelp Recovery Project, 2024: 20).

As well as featuring in MCS lessons and recommendations from the byelaw, anecdotal evidence is highly valued by the SKRP as ‘a vital and complementary source of information for the project’ (Sussex Kelp Recovery Project, 2024: 20). Anecdotal evidence is important given the large space that this project expands into (300km²). Scientific evidence can only be gathered in specific sites and times, given the size of the area and the expense of this type of research. To complement this scientific research, SKRP carries out qualitative interviews and surveys and uses photos and videos from Sussex Underwater, a group of volunteer divers, and other sea users. One participant says:

We know the flaws of anecdotal evidence. But it is, every time it is the thing that tells us that change is happening...It’s not the surveys being done every year, year in year out by the academics. They will show us in due course, the change. But the change is noted first by divers and fishers...if you read the scientific literature, anecdotal evidence can be quite pooh-poohed. But this project would be absolutely lost without it. Our narrative is...dominated by an anecdotal narrative, which is then underpinned by the scientific narrative (P7, Conservation Organisation).

¹²⁰ There is a lively debate in applied epistemology about what constitutes an expert and what defines expertise. There is fundamental disagreement on the methodology that one should make use of to arrive at an understanding of expertise. Some seek a conceptual analysis (Goldman, 2018; Croce, 2019). Christian Quast (2018) uses practical explication to argue that expertise is a ‘service-relation’. Oliver R. Scholz (2018) identifies ‘symptoms’ of expertise.

This participant highlights the importance of anecdotal evidence, in providing a complete picture of what is happening over large area more quickly. While the practical costs of research in marine environments could act as a barrier to knowledge and understanding of them, the availability of anecdotal evidence is experienced by this participant as beneficial to overcoming such a barrier.

A familiar tension appears to underly both the consultation process and use of anecdotal evidence: a tension between the epistemic benefits of scientific evidence or lengthy evidence-gathering procedures and the practical costs associated with them, especially given time pressures and funding restrictions.¹²¹ There are some important questions that can be raised here in the context of epistemology. One point to make is that discussion of evidence in epistemology, such as the evidentialist thesis mentioned in the introduction of this section, focusses almost entirely on the relationship between evidence and belief, and does not illuminate much about epistemic norms that might govern the process of gathering evidence. One possible explanation for this is that evidence gathering is thought to be relative to practical interests and therefore the norms that govern them would be practical, rather than epistemic, but this idea is resisted in recent work (Flores and Woodard, 2023). This would suggest that even where epistemology engages with the concept of evidence, it is silent on the epistemic factors identified by participants above. Although we cannot draw many concrete conclusions from this case alone, if a similar result was produced across other case studies, this could be a useful finding for *Directing Inquiry* (see Chapter 8, Section 8.1). If the norms that govern evidence-gathering are important, this might justify directing our attention towards this issue in epistemological inquiries.

Another point is that this tension between higher evidential standards and practical costs engages with an important epistemological question already noted in environmental discourse. The government, we know, advises a precautionary approach to conservation measures in light of scientific uncertainty and this is the approach the Impact Assessment takes. But some philosophers note the ambiguity of this sort of principle, including the

¹²¹ This is not to say that there are not other reasons to value anecdotal evidence, such as inclusivity and collaboration, or that scientific evidence is the only robust sort of evidence.

epistemological question of what degree of evidence is required to take a threat of harm seriously (Hopster, 2023: 2435).¹²² This requires input from epistemologists. In this case study, the complexity of this epistemological question manifests itself in a practical scenario, where there is a need to establish an understanding of impacts and ensure that the proposed measure is fair and well-evidenced, against the challenges of producing evidence in good time to prevent or reverse environmental harm. The complexity of both the epistemological and ethical questions at the centre of this sort of concern, and the multitude of considerations which might be relevant, reveals itself in this case study. This might then be an example of data emerging from case studies that could be helpful in *Providing Practical Scenarios* (see Section 8.1) to which principles could be applied.

9.1.2 Awareness

The Nearshore Trawling Byelaw received a large amount of public attention for a piece of fisheries management, particularly in relation to the impact of trawling on historic kelp beds. Following the *Help Our Kelp* campaign for the byelaw (which involved releasing media and lobbying the government), IFCA received 2500 responses to this consultation (Marine Conservation Society, 2023: 16). Participants consider the success in raising awareness for the byelaw to be a significant factor in public support.

Out of Sight, Out of Mind

Many participants note the tendency for what happens in marine environments to go unnoticed (n=4). One participant says:

I think everybody was slow to realise how much of the sea forest had gone... so it was out of sight, out of mind, you know, had the same thing happened on land there would have been a national outcry... I think as people had started to realise what we had lost, I think that did help to sway public opinion and it certainly was why thousands of people signed up to support the trawling ban (P6, Production Company).

¹²² There is significant agreement that precaution is appropriate only when the risk of harm is a realistic possibility (Carter and Peterson, 2015; Gardiner, 2006; Hopster, 2023).

The area of historic kelp is *sub-tidal*, which means it is always underwater. The impact of trawling on the kelp was therefore not visible to most people, and ‘out of sight, out of mind’. Raising awareness was therefore particularly important to bringing this issue to the wider public attention.

The Film and Visual Imagery

All participants mention the role of visual imagery of kelp (n=7). The film *Help Our Kelp* (2019), launched on 30th September, was created to ‘raise public awareness of the importance of underwater kelp forests which are declining nationally, and have all but disappeared off our Sussex coast in the last 30 years’ (Big Wave Productions, n.d). The film was broadcast on both national and local TV: the One Show, BBC local news and BBC Inside Out nationwide. It was also shared on BBC Online and BBC World News. The film communicates about what kelp is, its role in the ecosystem and the pressures of trawling, as well as details about how viewers can get involved to support a trawling ban. The involvement of Sir David Attenborough, who narrated the film, was also thought to increase its impact (n=5).

Participants emphasise that the film was effective, not just because it raised awareness of the facts, but because of the way that it raised awareness of the facts:

The waters in Sussex are very turbid so the images don’t really do justice to the landscape beneath the waves. So the film I think was particularly well done and inspiring... brought it all to life (P1, Conservation Organisation).

Any visual content, any sort of media that helps spread the message... will greatly make people realise...why they need to care and why it’s so important (P6, Production Company).

The *Help Our Kelp* film made the kelp and its beauty *visible* to the public. The use of visual content, which showed images of kelp and species that live within it, was an important factor in public engagement. This sub-theme appears in the earlier chapters of this thesis, with the argument that it is not enough to merely become aware of an important environmental issue that is ‘out of sight, out of mind’, we must become aware with vivid

representations that stick in our minds, and therefore have a lasting impact on us (Jenni, 2003).

One participant mentions the process of discovery itself as significant enabling factor:

If you're going to talk British wildlife...you take the boring old grey sea that nobody thinks is very interesting and you show them some footage of what is *actually* under the waves, it blows peoples' minds. It's so beautiful, and so amazing, and people don't know it's there. So...there is part of an unfolding of a genuine learning (P7, Conservation Organisation).

Other participants too note realisation or recognition of what has been lost as a significant factor in public engagement (n=4). Consider the opening line of the *Help Our Kelp* film: 'Off our coast, there are magical underwater forests. A world of giant seaweed, known as kelp...' (Help Our Kelp, 2019). The film presents itself as a reveal of something magical and this process of becoming aware of the kelp is itself part of the appeal of *Help Our Kelp* campaign. In becoming aware of the kelp forest, viewers discover that something previously thought to be boring is in fact beautiful and interesting.

Raising awareness of the kelp by making it visible to a wider audience is recognised as crucial to gaining public support for this management. This case study may be *Providing a Practical Scenario* (Section 8.1), that helps to concretise comments about which representations lead to an understanding, but also about being outward-looking. Participants use the phrase of "out of sight, out of mind" to describe the tendency to be distant from what is happening under the waves, because it is not obvious from one's surroundings, unless it is directly relevant to their lives – such as it is to the divers and fishers, in this case study. In this case, it matters that the film was created, as well as the imagery that continues to be provided by divers in Sussex Underwater.¹²³ Because of this effort, what is not obvious from peoples' immediate surroundings, was brought to their attention and into sharp focus with underwater video footage. The film was experienced by participants as helping the public overcome the distance or disconnect with what happens

¹²³ Sussex Underwater share underwater footage on social media to record its recovery and 'show local people that what's out of sight, isn't out of mind' (Sussex Kelp Recovery Project, 2024: 57). Sussex Underwater have accounts on Facebook, Instagram and Youtube.

beneath the waves, by raising awareness of the kelp forest through communication of accessible information and inspiring visual imagery.

One thing that stands out is that extending awareness to what is happening under the water involves collaboration. The Sussex Kelp Recovery Project states that:

The vast kelp beds were lost because they weren't visible enough. Now, due to the incredible efforts of local people, scientists, conservationists, documentary makers and others, awareness of the marine habitats and species of Sussex is growing, and in tandem, a local sense of pride in this fascinating ecosystem (Sussex Kelp Recovery Project, 2024: 21).

This quote frames the collaboration involved as the story of a community becoming aware, rather than an individual's quest for virtue. I draw further conclusions about this in Section 9.3.

9.1.3 Simplicity

Although the Nearshore Trawling Byelaw had a wider reach, the historic kelp beds have been the focus of the *Help Our Kelp* campaign and the Sussex Kelp Recovery Project. Many participants (n=6) mentioned that there were benefits to a focus on kelp in communications, such as the *Help Our Kelp* film, some of which were epistemic.

By focusing on kelp in communications, the message surrounding the byelaw was simplified. One reason that this was considered a benefit is that simplifying the issue meant that it was easy to grasp:

The byelaw process was all the statutory documents, it was like 90 page document full of detail so I think we can be assured that not many people have read all that... but the fact that all of that information can be consolidated into a 7 minute film narrated by, you know, eco superstar, um, means that everybody had a grasp of it (P7, Conservation Organisation).

The information relevant to the byelaw could be communicated simply in the *Help Our Kelp* campaign. In part, this is because of a focus on kelp, as well as the fact the byelaw

targets one pressure (trawling). It was therefore possible to condense the message into something that could be communicated in a short film that focussed on kelp.

As focus on the single species of kelp, rather than the more complicated picture of the byelaw's wider benefits, was also beneficial because of something interesting about kelp:

I think from a communications perspective, um, kelp is quite a, what do you call it, charismatic megaf flora [laughs], almost like a panda in that it's quite an iconic species, but people find it easy to form an attachment to. So it's (...) it can be kind of like a picture species for a movement, in the same way that like a panda can be for like, kind of, other conservation...and because people have quite often seen kelp in documentaries or gone snorkelling themselves in other parts of the UK, so they can kind of relate to it. So I think that kind of public understanding of kelp habitat was what enabled that big public buy in that for us was on a much bigger level than just kelp, but because of the kelp and because of the kind of local connection, I think, that was really useful in the kind of process this time (P5, Fisheries and Conservation Authority).

Kelp is a charismatic species and it is something that people can kind of latch on to and you know there's a lot of narrative around terrestrial forests so having a conversation about ocean (..) you know forests in the sea is not a kind of tricky kind of line for people to draw in terms of understanding why it's important (P2, Conservation Organisation).

The notion of a "charismatic" species is first seen in conservation literature as a trait of *flagship* species, which are popular species that can drive awareness and encourage action in conservation campaigns due to their capacity to evoke interest and emotional responses. In the article 'Non-Human Charisma', different factors are identified as part of charisma such detectability/distinctiveness (i.e. existing awareness), socio-economic biases (i.e. existing reputation or views), aesthetics, and potential to generate satisfaction (Lorimer, 2007). It is notable that all these factors are directly mentioned at least once by a participant in this study: the recognition of kelp, views about kelp and its reputation, the beauty of how it looks under the water, and the satisfaction of curiosity that people have when learning about the kelp. These factors were recognised as part of the reason that the Nearshore Trawling Byelaw gained a significant amount of attention.

However, despite the interest and intelligibility of a simple message about kelp, the issue itself is in many respects quite complex. For example, participants note that there are more pressures on kelp beds beyond trawling, including sedimentation (Sussex Kelp Recovery Project, 2022). Many participants also note the complexity of marine environments in general and the interactions between various stakeholders. The Nearshore Trawling Byelaw is also expected to have wider benefits beyond kelp recovery. One participant mentioned that there were some epistemic disadvantages to the focus on kelp:

I get that in order to connect with wider, kind of, society things need to be made super straightforward and simple and that's where kelp being a kind of charismatic species and something that is lost and recovery and climate crisis by diversity crisis, you know, it is (..) um something that sticks in peoples' minds, but is not possible to happen everywhere. So, some people now think that they might be able to get kelp (..) when they won't (P2, Conservation Organisation).

The participant describes a risk of misunderstanding that is produced by the perception that the Nearshore Byelaw was solely aimed at kelp recovery. The trawling ban extends along the Sussex coastline, including areas that have never had kelp, because the justification for this legislation takes an ecosystem approach, as defined in the discussion of the epistemic theme of evidence in Section 9.1.1. The participant emphasises that, while a focus on kelp may improve communication, it is not entirely unproblematic; the risk of misunderstanding the aims of the byelaw meant that some perceived the byelaw to be unjustified in areas that cannot support kelp beds.

When simplifying a complex issue, something is gained by increasing intelligibility and interest for wider public engagement, but complexity is lost in this process. This suggests a tension between presenting a message that is easy to grasp and attracts attention and communicating the complexity of an issue accurately. In their report, the MCS suggests that information should be tailored to meet different needs and interests: the fishing community needs 'accessible but detailed information, whereas the local community need a story with a simple premise to capture their imaginations' (Marine Conservation Society, n.d.: 4). Following the insights in this report, the appropriate focus may depend on the purpose of communications.

Some parallels can be drawn here with my arguments in earlier chapters, in that what matters is not necessarily how much is known, but intelligibility (Chapter 1-4). If this finding was consistently found in other case studies, this would be the sort of data that could help with *Informing the Normative* (see Section 8.1). I have argued that thinking about the environment involves processing vast amounts of complex information in a way that renders it intelligible. I also argued that, because understanding provides intelligibility rather than propositional knowledge, it is understanding that is of most importance to an epistemology for the environment. The presence of these issues in this case study supports the premise that intelligibility is a central concern in environmental contexts.

9.1.4 Hope

The story of the Nearshore Trawling Byelaw and the promise of kelp recovery is hopeful, which is recognised as having a positive impact on its audience. Hope might not strike the reader as obviously epistemic, but there is an epistemic dimension to hope because it is characterised by a specific sort of belief. Hope is accompanied by a less than confident belief that a desired outcome is possible (Moellendorf, 2022: 9; Williston, 2012: 168). To see this, consider the contrast with hopelessness, which is accompanied with a belief that an outcome is *impossible* and the contrast with wishful thinking, where an agent manifests an irrational confidence that some outcome obtains or is possible. Hope maintains a belief that some positive outcome is possible, but at the same time, unlike wishful thinking, responds to evidence against the desired outcome (Williston, 2012: 172). Hope, then, requires that an agent believes in the possibility of the desired outcome, but not so confidently that they do not fear that the desired outcome won't happen or so irrationally that their belief persists in the face of contrary evidence.

Hope can be found throughout the narrative of this piece of fisheries management. The Nearshore Trawling Byelaw removes a significant pressure on the kelp and therefore may enable its successful return. Sir David Attenborough's response to the ban captures this message:

Sussex's remarkable kelp forests will now have a chance to regenerate and provide a home for hundreds of species, creating an oasis of life off the coast, enhancing

fisheries and sequestering carbon in our fight against climate change (Sussex Kelp Recovery Project, 2022: 2).

Relieving the area of the pressure of trawling presents the possibility of recovery. Attenborough paints this picture of a flourishing future ecosystem and its benefits. In a blog post, the Blue Marine Foundation describes the recovery of kelp as a ‘story of hope’ (Blue Marine Foundation, 2023). Recently, SKRP have also released a short film *Glimmers of Hope* that pictures signs of recovering marine life (Sussex Kelp Recovery Project, n.d.a).

Moreover, some participants consider the belief in the possibility of kelp recovery to be important to public reception of the Byelaw:

I think that made people just recognise that they had lost something very important, off their coastline and, actually, they wanted to get behind...giving it a chance to recover (P1, Conservation Organisation).

I think all that people felt when they heard about the trawling ban coming into place was that they liked that the forest had a chance to regrow and that the environment had a chance to recover (P6, Production Company).

Some participants clearly felt hope, specifically hopeful narrative, distinguishes the Nearshore Trawling Byelaw and the Sussex Kelp Recovery Project from other environmental projects and this partly accounts for the significant amount of attention it received from the wider public:

in a world of miserable stories about the environment, it is genuinely, um, exciting and hopeful (P7, Production Company).

There’s a certain way of storytelling isn’t there, where it used to be great, and then it got really bad, and then the opportunity for it to come *back* ... clearly the story of Eric Smith was really, really important there because he used to dive as a kid in the kelp forest ...And then the hope that he has as an old man seeing, hopefully, the kelp forest kind of come back in his lifetime (P3, Local Authority).¹²⁴

¹²⁴ Eric Smith has been diving in Sussex since he was 11. He wrote letters and articles about changes to marine life in the area as a result of trawling and campaigned to ban trawling in Sussex (Smith, n.d).

It seems that the hopeful story, where the kelp is lost and then given the opportunity to recover and return increased public interest.¹²⁵ To garner this degree of public support in the Nearshore Trawling Byelaw, it was important to invoke hope and belief in the possibility of the kelp's return.

In this case study, some experienced hope having a positive impact in engaging people in this issue and mobilising them to take action. This reflects the view that hope can be an important part of motivating people to take action on challenging environmental issues, e.g. climate change (Moellendorf, 2022; Williston, 2012). Hope is therefore shown to be epistemically relevant in two ways: (1) by being characterised by a specific belief, i.e. less than confident belief of some desired outcome and (2) by encouraging interest and raising awareness. We might worry, however, that believing that the desired outcome is possible would have the opposite effect, e.g. leading to complacency. But recall that the epistemic dimension of hope would suggest otherwise; when we hope for a desired outcome, we have a *less than confident belief*, such that we can still fear that the desired outcome may not obtain. It would make sense that this sort of epistemic condition would ground motivation to take action, given that we believe there is possibility of a desired outcome that we cannot take for granted.

9.1.5 Understanding Value

It was not only important to understand and raise awareness about what was happening under the water, but also to understand relevant *values*. References to understanding of value include the process of arriving at an understanding of value, sharing an understanding of value, and misunderstanding of values and their interactions.

Arriving at an Understanding of Value

The Sussex Kelp Recovery Project states as its second aim: 'to understand the ecological, social and economic value of kelp and other essential fish habitats in Sussex' (Sussex Kelp Recovery Project, 2024: 33-42). Moreover, the Impact Assessment produced by IFCA aims

¹²⁵ A related finding is that some participants (n=2) note the benefits of positive framing, where we gain something rather than take something away.

to establish the environmental value of different habitats in its district. To do so, IFCA refers to the concept of ‘Natural Capital’, which is defined by the Natural Capital Committee as ‘the elements of nature that directly or indirectly produce value to people, including ecosystems, species, freshwater, land, minerals, the air and oceans, as well as natural processes and functions’ (Sussex IFCA, 2020: 23).¹²⁶ As well as ‘Ecosystem Services’, defined as ‘the functions and products from nature that can be turned into human benefits with varying degrees of human input’ (Sussex IFCA, 2020: 24). Fish habitats are part of the natural capital of marine environments because they deliver many ecosystem services. The IFCA bases its argument for the Nearshore Trawling Byelaw on these calculations of value.

Sharing Understanding of Value

Another important element raised by participants is not just the effort to determine value, but the effort to communicate this value to the wider public.

In a sense the value was established through the David Attenborough narrated film which was, if you like, a way of getting past the numbers into really understanding the story of the value of the kelp forest (P3, Local Authority).

[the *Help Our Kelp* film] definitely raised the profile...I think one of the key things it shows is that if people understand why something is important or how it relates directly to their lives, um, and their quality of life, then they will respond to that (P4, Conservation Organisation).

The film is thought by these participants to have contributed to communicating the value of the kelp forest to the wider public. So, alongside the establishing or determining of value that took place as part of the consultation process (which we might think of as arriving at an understanding of value), understanding value also emerged in the context of revealing value to a wider audience.

¹²⁶ The Natural Capital Committee was an independent advisory committee, advising the government on natural capital between 2012 and December 2020 (UK Government, n.d).

Misunderstanding Value

Many participants recognise that there is commonly a misunderstanding of value that takes the form of a false dichotomy, putting environmental and economic factors against each other:

What is so challenging about the current assessment process is that it also pits environmental value against economic value, when actually there is so much more around that...there's all sorts of values that you can place on a space that the current sort of impact assessment doesn't kind of allow for... that really got in the way, because it didn't enable discussion around non-financial value...For me, it's kind of quite an inherently flawed kind of process as it stands (P2, Conservation Organisation)

I think for a lot of people, they're so disconnected from nature that they don't actually understand the foundation that their lives are reliant on that. You know, our economy is reliant on that... without that, you know, there isn't anything else. I don't think people really understand that...I think one of our greatest challenges in conservation, I think it's beginning to be understood, but maybe too late, is how very dependent on nature we are (P4, Conservation Organisation).

A false dichotomy misrepresents available options in weighing different values, e.g. believing two values to be in conflict more than they are. In this sense, participants describe a mistake, false belief or misunderstanding, where economic and environmental values are seen to be in conflict, where they are not. This was considered an epistemic barrier.

9.1.6 Trust

As mentioned within the theme of awareness, Attenborough's involvement brought both attention to this campaign and helped the *Help Our Kelp* film reach a wider audience. Another element of Attenborough's involvement was the credibility that he brought to the campaign:

Of course, David Attenborough, you can't do any better than having his voice to validate how important a project is. (P7, Production Company).

nobody's going to argue with David Attenborough, right? (P3, Conservation Organisation).

The involvement of Sir David Attenborough narrating the film not only brought more attention to the issue but validated the message.

Participants above speak about Attenborough's involvement as if it requires no further explanation why he has the capacity to validate the message. But when asked what it is about David Attenborough that makes people listen, one participant says:

certainly in this country, he feels very relevant to people, he's been very consistent, he's been a consistent presence communicating you know the same principle for decades, you know, for somebody growing up in the UK, sort of, the BBC natural world, all of that stuff. It's what a lot of people in the UK have grown up with, if you think about nature conservation in the UK, you would think David Attenborough. So it's the consistency I think it's the longevity of his role and, um, yeah trust, I guess, people trust him (P4, Conservation Organisation).

Attenborough is a figure in environmental discourse that people are familiar with and he is considered a credible voice. His involvement is thought to have not only helped the film reach a wider audience, but also enhanced trust in the importance and urgency of this issue.

The trust of the wider public is crucial to communication of environmental issues because our evidence about environmental issues is highly dependent on our trust in others (Almassi, 2012). We are often distanced from evidence of environmental changes, e.g. in this case the kelp forest decline was not initially visible to most and required expertise to understand. Some will be better positioned than others to testify this evidence. When we cannot rely on our own understanding or experience of a phenomenon, public understanding of environmental issues faces significant social-epistemic problems that arise from how to identify and trust in experts from the non-expert perspective.¹²⁷ The role of trust in facilitating the sharing of knowledge and understanding about environmental issues appears to be present in this case study, where trust is thought to play a role in the

¹²⁷ Trust is also seen as important in not only the reception of information in the case study, but the deliverance of that information. The MCS write the following: 'mutual trust and respect is essential throughout the process and stakeholders should be engaged from the outset to discuss options and facilitate the input of their insight and expertise' (Marine Conservation Society, n.d.: 4).

effectiveness of the messaging around the kelp. Arguably, the presence of this epistemic theme in this case study could be useful in *Directing Inquiries*, because it validates the fact this issue receives considerable attention in applied epistemology (Goldman, 2001; Guerrero, 2016).

9.1.7 Memory

This case study centrally involves memory in several respects. Memory plays a role in how public support was achieved, but memory also plays a role in potential resistance to the byelaw. There is some overlap here with the sub-theme of anecdotal evidence, which often includes memories from the fishing community and divers, but the discussion of anecdotal evidence referred to the use of different methods of data collection and their benefits, whereas discussion of memory here refers to the functioning of memory in our epistemic lives and the positive and negative associations of memories.

Participants note that the divers drew attention to the kelp decline. Many participants mention Eric Smith as a key figure in the *Help Our Kelp* campaign and specifically refer to his memories of the kelp and their impact on public understanding of kelp loss. Eric Smith has since worked with Big Wave TV to produce another film, *Our Sea Forest*, to tell the story of kelp decline and recovery (this took place after the campaign). Here is a quote from Eric Smith in this film:

I was 11 years old when this was taken, in 1959. This is what it looked like. Beautiful clear water. Reminds me of my youth. Everywhere you went was kelp, kelp, kelp. Cleaning the water, filtering the water. And it's such a home to everything. As a matter of fact, you had a job getting in the water sometimes, there was so much kelp. The area between the kelp and the sand was a superb nursery full of tiny fish. It was a forest. An absolute forest ecosystem. GoneIt brings back so many memories, it's untrue. Right, sorry. Pull myself together a bit, that was a bit dramatic, but. It's just...that's my past, gone (*Our Sea Forest*, 2023).

Eric Smith recalls his memory of the kelp with much emotion. Many felt that sharing these memories of the kelp forest was important in this case study.

In stark contrast to Eric Smith's fond recollections of the kelp forest, participants note that some members of the public may have negative memories of the kelp on beaches (n=3):

People's reference points were linked to probably historic memories of kelp on the beaches where it wasn't necessarily considered to be a positive. So piles of smelly kelp on the beaches, you know, prior to its loss. It was a problem. It was a problem for local authorities. It was a problem for the tourists and visitors. So I don't know that people really missed it, who weren't divers, or who weren't fishermen...I think the film, for example, was critical in giving them an insight and a window into what kelp looks like and what it supports when it's growing naturally in the water, before it gets washed up (P1, Conservation Organisation)

From a public's point of view, say in Worthing, and I'm sure it's true elsewhere in the coast, is that the public mind was split, some people wanted it back for sustainability reasons, others didn't want it back because it used to *stink*...I remember it as a kid and it used to be piled high, it stank and we couldn't go to the beach and there were flies everywhere...and that was the Worthing Council report which was like how the hell are we going to manage all of this kelp on the beach...they had a problem and the local people *remember* all of that... and they would have lobbied the MPs to that effect as well as people lobbying for it to come back for sustainability (P3, local authority).¹²⁸

The interesting, sort of, counter to protecting the kelp... the only sort of negative... I mean apart from, obviously the trawlers who are annoyed because they can't trawl. Was this sort of idea of the kelp washing up on the beach, and it will smell bad and it's, you know, ugly (P5, Fisheries and Conversation Authority).

This formed, alongside the trawlers, another potential objection to the trawling ban revolving around the possible impact of kelp on residents of coastal towns and tourism. One participant notes that the film was important to show the public what kelp looked like underwater, offering a different reference point beyond these memories.

¹²⁸ There are many articles in the *Worthing Herald* that discuss seaweed as a problem. One example of this is a front page news story entitled 'Seaweed On-Sea' on 25th August 1989 (Hollebone, 1989). Seaweed removal, in the interest of tourism, on beaches such as Worthing was discussed in Parliament on 18th February 1987.

It's unsurprising that memory is a key epistemic theme in this case study, given that the kelp was abundant in the past and then it was lost. We often know about the past by remembering the past. The accuracy, the vividness, activation and the positive or negative connotations of our memories and their capacity to motivate us are plausibly important to any environmental issue where the past is contrasted with the present and possible futures. It strikes me that there are interesting philosophical questions to ask here.¹²⁹ Perhaps by highlighting the importance of memory, this finding is useful for *Directing Inquiries* in this area. What is interesting about memory in this case study is that the role of memory is closely tied to emotion. Although what we remember and how we remember is very plausibly connected to what we feel, perhaps this calls into question that we can have a meaningful understanding of the epistemic relevance of memory abstracted from the emotions of people who are remembering. I return to this idea in Section 9.3.

9.2 Other Important Themes

There were other important themes that are not epistemic according to the rationale I have provided in my approach to analysis. In this section, I introduce these and highlight some connections with the epistemic themes in Section 9.1.

9.2.1 Objections

Most participants mention objections to the byelaw from the trawling community (n=6). IFCA are required to consider all objections raised during public consultation. Some participants say that, due to these objections to the trawling ban, there were members of the fishing community that were hesitant to speak out in favour of it (n=2). Some participants expressed the view that there could have been more support for the transition of those affected by the ban and who suffered significant financial losses as a result (n=2):

There is no government conversation around, kind of, wider support to, kind of, cover losses. So, the management measures themselves... I guess cover the areas they need to but they weren't communicated perfectly and the roll out has been extremely challenging for some (P2, Conservation Organisation).

¹²⁹ Perhaps applying debates about memory in epistemology (Burge, 1993; Michaelian, 2011; Myers, 2024).

In interviews carried out with stakeholders, it was recognised that there had not been support consequent to the impact of the byelaw for trawlers who had to adapt to other areas or to use different fishing methods (Marine Conservation Society, n.d.: 25).

Objections to the byelaw is an important theme in this case study but is not itself epistemic according to my rationale. There are some things to note, however, in relation to the aims of this research. The first is that some participants suggest that significant public support, in part generated by the epistemic themes identified above, meant that objections did not slow the process down as much as they might have:

When you're introducing management, you're obviously stopping people from fishing in areas where they previously have fished. So if we hadn't had so much public buy in, and a voice, then the other voices from other stakeholders might have shouted louder (P5, Fisheries and Conservation Authority).

Some participants note too that mutual understanding of each other's perspective in response to objections was important:

Oversimplifying and going back to the, kind of, age old us versus them is actually quite harmful to the process. And not, um, highlighting where this is, kind of, common ground and, you know, mutual understanding means that um that also causes a lot of stress for the fishing community and also they feel kind of isolated and arguably unnecessarily (P2, Conservation Organisation)

The whole point of a consultation is to get peoples' opinions. You're not necessarily, going to, everybody agree with each other. But that's not really the point, the purpose is to make a measured... decision (P7, Conservation Organisation).

Understanding objections and other peoples' opinions is considered an important part of the consultation process.

9.2.2 Collaboration

The Nearshore Trawling Byelaw was led by IFCA, but it was supported by a group of other organisations that formed the Help Our Kelp Partnership:

The process is the same for us whatever, but normally we don't have... a large group of organisations pushing to get a byelaw through at the same time (P5, Fisheries and Conservation Authority).

This was recognised as being important to increasing public support for the byelaw and the ongoing SKRP:

Another thing that I think is really special to this project is the level of participation and the amount of people involved and...it's fantastic how many people were involved and, you know, no-one leads it, no-one owns it, but it's a shoulder-by-shoulder experience. So you've got to keep the culture, of we are all contributing into something alongside each other. If it was on land, the land would belong to someone and the landowner would lead the process, but it's at sea. So... it's a completely different type of leadership...SKRP ties to lead through collaboration (P7, Conservation Organisation).

SKRP describes itself as a partnership that aims to understand and facilitate kelp recovery and other fish habitats through collaboration (Sussex Kelp Recovery Project, 2024: 17 and 21).

Although collaboration is not an epistemic theme according to my criteria, it is important to emphasise that collaboration contributes to knowledge and understanding in this case study. For example, there are several projects associated with the SKRP that aim at engaging people to participate in improving our understanding of kelp recovery. This includes the Sussex Kelp Recording Scheme, which invites those who use the beach or sea in the area to download an app and record when kelp is seen. This will help to map kelp recovery. Sussex Seasearch (part of Marine Conservation Society's national programme) involves volunteer divers and snorkellers to create a database of sightings of marine life. Arriving at an understanding of kelp recovery is a collaborative effort.

9.3 Further Conclusions

Recall that in Chapter 8, Section 8.1 I argued that empirical research could inform applied epistemology in the following ways: (i) provide a practical setting in which to apply principles (*Providing Practical Scenarios*), (ii) direct epistemological inquiry

(*Directing Inquiries*) or (iii) influence how we construct normative principles (*Informing the Normative*). In discussion of the epistemic themes, I have indicated some ways in which the data could be relevant to these uses. However, it is clear from this discussion that to draw any significant conclusions about these points, much more data would be required and comparisons should be made across different case studies. This case study research only provides insights about the experience of epistemic factors within this case.

Despite the fact that this case study cannot provide a comprehensive picture, we could still raise questions concerning the relationship between this study and the first 7 chapters of this thesis. As addressed in Chapter 8, Section 8.6, this would speak to the concern that motivated this case study research, i.e. the question of whether I have chosen to focus on epistemic problems that are not the only, or most important, problems experienced by those involved in environmental interventions. It is important to acknowledge that there is a possibly large category of relevant information that is not found in either my thesis so far or the data that I have collected. Still, we can make some relevant comparisons by asking the following questions:

- (1) what is in the data and in the theory?
- (2) what is in the data, but not in the theory?
- (3) what is not in the data, but is in the theory?

Question (1) asks whether the concerns of my thesis are present in this case study, most notably whether the epistemic themes that emerge are relevant to processing vast amounts of complex information and the emphasis of understanding over knowledge. If what matters to my thesis is present in the case study, this lends some support to the idea that the theory in my thesis in chapters 1-7 reflects at least some experiences of real agents.

Question (2) asks what is present in this case study that is not in the theory. These might include epistemic factors that are as important or more important, but that are not in focus in my thesis. Question (3) asks what is prominent in the theory that does not appear in the data, which would highlight epistemic factors that are not experienced in this case study but are of importance in the thesis.

It turns out these comparisons are not as simple as naming epistemic themes which appear in the data and not in the theory, or vice versa, as the relevant comparisons are not

contained within the epistemic themes. It is instead *elements* of these epistemic themes, and the manner in which they were experienced, that are of concern in this comparison.

9.3.1 What is in the Data and the Theory?

Some of the epistemic themes that emerge from the data are also prominent in my thesis or the theory that informs my thesis. For example, as discussed in Section 9.1, the role of evidence is prominent in epistemology, environmental philosophy discusses understanding value (and calls for involvement of epistemologists), and the role of trust arises in prominent debates in applied epistemology, as well as questions about expertise. This suggests that some issues in the philosophical discourse that inform my thesis also arise in this case study.

Importantly, the issue that centrally motivates my arguments for a shift towards understanding in an epistemology for the environment is also present in this case study. As addressed in Section 9.1.3, I began this thesis with the claim that a focus on propositional knowledge is unlikely to suffice for an epistemology for the environment. I argued that understanding, which centres grasp and representations of the world that selectively emphasise what is important, is the epistemic condition of most relevance in this context. We could draw some similar conclusions from the prominence of the epistemic theme simplicity and the role of the film in providing visual representations of the kelp. It was important both that the message was intelligible for the wider public and that the film represented the kelp in a way that made its value apparent to a wider audience. A focus on understanding has the resources to elucidate this, but it is not clear that a focus on propositional knowledge does.

9.3.2 What is in the Data but not the Theory?

One conclusion we might draw from how the epistemic themes were experienced in this case study, particularly hope, memory and the focus on charismatic species, is that their epistemic relevance appears to be entangled in emotions. I have mentioned the role of emotion in some areas of this thesis, e.g. in discussion of outward-lookingness (although this focused on the negative impact of emotion) and in discussion of the emotions related to curiosity and interest. However, we might conclude from the findings of this case study

that emotions play a central role in our epistemic lives. In particular, this case study suggests that emotion not only influences what we pay attention to, what we learn about and what we remember, but also is deeply entangled in how we pay attention, how we learn and how we remember things.

Another observation about what is in the data and not the theory is that I have chosen to focus on individual epistemic agents, not because there is no interest in group epistemic agents and their epistemic states, but because these issues are outside the scope of the thesis. In the data, collaboration is an important theme. Although I have characterised collaboration as a non-epistemic theme, collaboration is consistently discussed in the context of acquiring knowledge and understanding. This suggests that a focus on the individual may not in the end suffice for a comprehensive epistemology for the environment.¹³⁰

9.3.3 What is not in the Data but is in the Theory?

One lesson we could draw from the relationship between my theorising and what is revealed in the data is that epistemologists do not typically refer to epistemic factors in the same way that participants in this case study did. The issues discussed in Chapters 1-4, e.g. the distinction between knowledge and understanding, externalism and internalism, objectivism and subjectivism, factive and non-factive etc., do not emerge from the data. However, many epistemic factors did arise in the data, as seen in the quotes included in the discussion of epistemic themes, with many rich descriptions of epistemological issues. It's unclear what conclusions should be drawn from a potential gap between the terminology in epistemology and the terminology that we generally use to describe our epistemic lives, but it could lend support to the idea that epistemologists ought to take an interest in a wider range epistemological questions of interest to real epistemic agents.

It is challenging to draw many concrete conclusions from this study. Some elements of the theory also appeared in the data, which validates the idea that these dynamics are

¹³⁰ We might take interest in group epistemic agents, although this would require that we have a good idea about what a group epistemic agent is. Some think that a group believes as an epistemic agent in its own right, when members jointly commit to a belief (Gilbert, 1987). Alexander Bird (2014) distinguishes the distributed model, according to which individuals in the group are given different roles in gathering or coordinating information to produce knowledge.

experienced in this case. Similarly, some factors experienced in the case study were not central to the theory and *vice versa*. Optimistically, this could suggest that empirical research is an important part of an epistemology for the environment that seeks relevance to real epistemic agents who are involved in environmental interventions. Regarding the exploratory aims of this project, there is some evidence that qualitative research has the potential to reveal valuable insights into applied epistemological issues. However, adjustments might need to be made to how such research is carried out, which I will now outline sharing my reflections on this research project.

9.4 Reflections

As pointed out in Section 9.3, this case study suggests that qualitative methods have the potential to contribute to debates in applied epistemology. However, this was a case study of one piece of legislation. Very few inferences can be made relating to other cases. One area for future study may be to look at other case studies, perhaps those that are similarly representative of possible epistemic barriers and unique in the sense that these epistemic barriers were overcome. It would then be possible to compare findings and see if there are any parallels in other cases. I am not aware of other case studies using qualitative methods for the specific purpose of furthering this debate, but there are of course examples of other qualitative case studies about our understanding of environmental issues (Norgaard, 2011). It would be easier to make direct comparisons with a case study that uses a similar methodology.

Given that this project is interdisciplinary, in the sense that it combines qualitative research methods with epistemology, more issues arose than might be expected, particularly in the way that the questions were framed. My questions specifically aimed at those involved in development or campaigning for the Nearshore Trawling Byelaw and valuable data may have been missed by this focus, such as the experiences of those directly affected by the byelaw and the public. The initial questions also had to be altered to improve communication with participants. These issues may have been prevented by carrying out a few preliminary unstructured interviews to inform the first questions asked, i.e. a pilot study, which would be especially useful in a larger study. This would also have helped to prepare questions more effectively and conduct wider research into views about

the ban and understandings of surrounding issues. Preliminary interviews may be particularly important given the interdisciplinary nature of this case study. The context of the Covid pandemic made this unlikely in these particular circumstances.

What went well, which I would repeat in similar projects, is the use of semi-structured interviews. In particular, to begin with an open-ended question and then ask specific questions about epistemic factors afterwards. This seemed to lead to conversation and better encouraged participants' views. Although the epistemic and non-epistemic themes were present in the publicly available documents that I reviewed as part of this case study, the interview data led to in depth descriptions. In addition, some themes, such as hope and memory, and sub-themes such as the role of visual imagery, became clearer from the interview data. I conclude from this that one of the findings with respect to the exploratory aims of this case study is that qualitative interviews have potential to reveal more detail on epistemic factors that might be missed by solely reviewing publicly available documents. Collecting even more data via interviews would obviously have been beneficial and would be recommended for future projects, especially those beyond the Covid restrictions and aftermath.¹³¹

9.5 Conclusion

An important finding from this case study is that, despite the potential for various obstacles to get in the way of the passing of the Byelaw, e.g. objections, difficulty acquiring evidence, and lack of visibility of the kelp decline, many epistemically relevant factors aligned and contributed to significant attention and support for the Nearshore Trawling Byelaw, as well as the success of the ongoing Sussex Kelp Recovery Project which continues to maintain public support. Case study research into similar important cases may prove to be an important part of developing an epistemology for environment.

¹³¹ As discussed in Chapter 8, Section 8.5.2, the COVID-19 pandemic meant that qualitative researchers had to adapt their data collection methods (Lobe et al., 2020; Oliffe et al., 2021; Keen, et al., 2022)

Conclusion

In this thesis, I have explored some approaches to developing an epistemology for the environment. I began with an epistemic problem that arises in this context. In order to do the right thing by the environment, we need to think about the subject in the right way, and this requires that we process vast amounts of complex information. I have argued that an epistemology for the environment should therefore be understanding-oriented; understanding, in virtue of its grasping element, renders the world intelligible.

In Chapter 1, I argued that truth is not the sole fundamental epistemic aim or concern and neither is it a particularly important epistemic aim or concern. This veritist view could not explain the importance of grasping, which I argued centrally concerns reason-responsiveness. In Chapter 2, I defended some of the commitments that follow from the claim that grasping is an epistemic success that facilitates responsiveness to facts that objectively count in favour of certain responses, i.e. objective normative reasons. In Chapter 3, I argued that the degree to which one grasps corresponds to the degree to which one is in an epistemic position to respond to objective normative reasons. In Chapter 4, I argued that one need not truly represent the world in order to be in an epistemic position to respond to facts and one might even improve one's position to respond to facts by deviating from the truth, thus validating both anti-veritist arguments and a non-factive view of understanding. At the core of understanding is its grasping element. Grasping renders the world intelligible by making features of the world that present reasons salient. An epistemology for the environment should therefore be understanding-oriented because this is the sort of epistemic condition that allows us to process vast amounts of complex information concerning the environment.

I have then argued, in Chapter 5, that a focus on responsibilist epistemic virtues will also be important to an epistemology for the environment, by providing both guiding advice to the epistemic agent who seeks to do right by the environment and a way to evaluate epistemic agents. In Chapter 6, I argued that outward-lookingness, among other virtues, is an environmental-epistemic virtue because it contributes to what I have called big picture thinking. In Chapter 7, I argued for an account of honesty as the avoidance of deception

when we are rightly trusted not to deceive. This account of honesty is also understanding-oriented, given that deception can also take the form of distorting understanding and this does not always centre around the distortion of truth.

Since the epistemic problem I have focused on may not be the only epistemic problem that matters in an environmental context, I suggested in Chapter 8 that we should seek to better understand the experiences of those who are involved in environmental interventions by carrying out empirical research. To demonstrate this, I undertook case study research on the Nearshore Trawling Byelaw in Sussex. Chapter 8 also outlined this investigation and defended the methodological choices in this case study research. In Chapter 9, I presented the findings, concluding that case study research like this could potentially be important to an emergent epistemology for the environment.

This thesis also raises further important questions that present opportunities for future research. I will now discuss some possible future research projects.

I have argued that the degree to which one grasps corresponds to the degree to which one is in an epistemic position to respond to facts that objectively count in favour of certain responses. I have also argued that this means grasp is not immune to epistemic luck. However, as I discussed in Section 3.1.3, my arguments are not equally powerful for all forms of luck, rather only with respect to Gettier luck that interferes between the belief and fact. It is less clear that environmental luck (e.g. in barn façade cases) interferes with one's responsiveness to a fact. However, others have argued that understanding is immune to environmental luck, but not to Gettier luck (Pritchard, 2008). Possible future research might explore in greater detail whether my approach to grasp supports this thesis and what this means for the distinction between understanding and knowing.

I have argued too that responsibilist virtues are important to an epistemology for the environment, which would support further research to individuate and characterise epistemic virtues through the environmental lens. Some specific questions arise from my account of being honest. First, there are questions about what we can and cannot be rightly trusted to do. Second, from this account honesty has quite a wide scope, especially if one can *always* rightly trust oneself to do what one takes to be the right thing. If all virtues involve this effort to do what one takes to be right, then all virtues have some important

relation to honesty. This raises the question of whether honesty is more fundamental than other virtues, which might be explored in further research.

Finally, it is fair to say that few conclusions could be taken from the case study research due to the limited amount of data and the focus on one case study. Given that the limited data collected has the potential to be useful to an epistemology for the environment (and applied epistemology in general), future research might use a similar (if even adapted) methodology to collect more data and compare findings for the benefit of future theory.

I will end this thesis by summarising the important conclusions. An adequate epistemology for the environment may take different approaches to define and to answer the epistemological questions that arise from our relationship to the natural environment. I have explored some potentially valuable approaches, such as considering our epistemic aims and concerns, identifying and characterising environmental-epistemic virtues, and carrying out interdisciplinary research. Given my account of grasp as an epistemic relation that helps us make sense of a complex world, an epistemology for the environment should focus its attention on understanding and how it is achieved.

Appendices

Appendix 1

Background Questions

(1) Original Question: What is your understanding of how the Nearshore Byelaw/ trawling ban came about and what was the nature of your involvement?

Revised Question: How were you involved/ affected by the trawling ban (or the ongoing Sussex kelp recovery project)

Main Questions

(2) Original question: Do you think the nearshore trawling byelaw is the best decision (in response to the depletion of kelp forests in Sussex)?

Revised question: Do you think that banning trawling was the right thing to do? Are there any advantages/ disadvantages to the trawling ban?

(3) Based on past answers, what do you think enabled/ prevented/ delayed the [best outcome]?

(4) Original question: Do you think knowledge/ understanding of the kelp forests and environmental issues had a role to play in enabling/ preventing/ delaying the best outcome?

Revised questions: (4a) Do you think that knowledge/ understanding (or lack of knowledge/ understanding) of the kelp forests had a role to play in enabling/ preventing/ delaying the trawling ban or any improvements to this response?

(4b) Do you think that knowledge/ understanding (lack of knowledge/ understanding) of the trawling ban and its consequences had a role to play in enabling/preventing/ delaying trawling ban or any of the improvements to this response?

Other Questions

(5) What role, if any, did the lack of visibility of sea kelp play in delaying or preventing the best outcome?

Please rate the contribution 0 – 5, where 0 is not at all and 5 is significant contribution.

(6) What role, if any, did the complexity of the issue play in delaying or preventing the best outcome?

Please rate the contribution 0 – 5, where 0 is not at all and 5 is significant contribution.

(7) What role, if any, did the way that this issue is covered in the media play in preventing or delaying the best outcome?

Please rate the contribution 0 – 5, where 0 is not at all and 5 is significant contribution.

(8) What role, if any, do you think conflict with human interests played in preventing or delaying the best outcome?

Please rate the contribution 0 – 5, where 0 is not at all and 5 is significant contribution.

Other Participants

(9) Do you have any contact details for others who may be willing to take part in this research?

Appendix 2

The Sussex Kelp Restoration Project and the Moral Responsibility to Understand Our Environment

PARTICIPANT INFORMATION SHEET

(1) What is this study about?

You are invited to take part in a research study about the kelp restoration project in Sussex. If you choose to take part, I will ask you about your views on the trawling ban in Sussex. I hope to find out what you think helps us do the right thing for the environment and what you think gets in the way.

My wider research is about what prevents us from meeting our moral responsibilities (and what helps us meet them), particularly towards the environment. I research our understanding of the environment and how it affects our moral responsibilities and attitudes such as praise and blame. I am interested in the general nature of these concepts and will not be making moral judgments about any particular groups, individuals, or actions.

The data collected from these interviews will provide some examples of what might enable, prevent or delay action to understand and protect our environment. You have been invited to participate in this study because you were involved in the campaign for a trawling ban, you might be affected by the implementation of the trawling ban or you are involved in the Sussex Kelp Restoration Project. This Participant Information Sheet tells you about the research study. Knowing what is involved will help you decide if you want to take part in the study. Please read this sheet carefully and ask questions about anything that you don't understand or want to know more about.

Participation in this research study is voluntary. By giving consent to take part in this study you are telling us that you:

- Understand what you have read.
- Agree to take part in the research study as outlined below.
- Agree to the use of your personal information as described.
- You have received a copy of this Participant Information Sheet to keep.

(2) Who is running the study?

The study is being carried out by the following researcher(s): Miss Alice Wheatley. This will take place under the supervision of Dr James Andow (j.andow@uea.ac.uk,).

(3) What will the study involve for me?

Your involvement in this study would include an interview at an agreed time between January and July 2023 and can take place in person or via video call.

The interview questions will ask you about what you believe to be the best outcome of the Help Our Kelp campaign/ Sussex Kelp Restoration Project and whether you feel this

outcome was prevented or delayed. I am particularly interested in what you feel has enabled, prevented or delayed the best outcome and your views on what helps or prevents us from understanding our environment.

An audio/video recording will be taken.

You will have the opportunity to review information generated about you prior to publication.

(4) How much of my time will the study take?
The interview will be no longer than 20 minutes long.

(5) Do I have to be in the study? Can I withdraw from the study once I have started?
Being in this study is completely voluntary and you do not have to take part. Your decision whether to participate will not affect your current or future relationship with the researchers or anyone else at the University of East Anglia now or in the future.

If you decide to take part in the study, you can withdraw your consent up to the point that your data is fully anonymised. You can do this by contacting alice.wheatley@uea.ac.uk

(6) What are the consequences if I withdraw from the study?
You are free to stop the interview at any time. Unless you say that you want us to keep them, any recordings will be erased and the information you have provided will not be included in the study results. You may also refuse to answer any questions that you do not wish to answer during the interview. If you decide at a later time to withdraw from the study your information will be removed from our records and will not be included in any results, up to the point we have analysed and published the results.

(7) Are there any risks or costs associated with being in the study?
Aside from giving up your time, we do not expect that there will be any risks or costs associated with taking part in this study.

(8) Are there any benefits associated with being in the study?
The potential benefits of contributing to this study include the opportunity to discuss and share your experience of what enables or interferes with our ability to meet our moral responsibilities towards the environment.

(9) What will happen to information provided by me and data collected during the study?
The data collected from these interviews will provide some examples of factors that enable, prevent or delay action to protect our environment, particularly where this is relevant to knowledge and understanding. The interview will be audio recorded and stored. An anonymised transcript will be produced and stored; no individuals or organisations will be named on the transcript. Participants may request access to this transcript and can withdraw their contribution at any time prior to data analysis. Any data used in the final thesis or future publications will remain anonymised.

Your personal data and information will only be used as outlined in this Participant Information Sheet, unless you consent otherwise. Data management will follow the Data Protection Act 2018 (DPA 2018) and UK General Data Protection Regulation (UK GDPR), and the University of East Anglia's Research Data Management Policy.

The information you provide will be stored securely and your identity will be kept strictly confidential, except as required by law. Study findings may be published. Although every effort will be made to protect your identity, there is a risk that you might be identifiable due to the nature of the study and/or results.

Study data may also be deposited with a repository to allow it to be made available for scholarly and educational purposes. The data will be kept for at least 10 years beyond the last date the data were accessed. The deposited data will not include your name or any directly identifiable information about you, but there is a risk that you might be identifiable due to the nature of the study and/or results.

(10) What if I would like further information about the study?

When you have read this information, Miss Alice Wheatley (alice.wheatley@uea.ac.uk) will be available to discuss it with you further and answer any questions you may have.

(11) Will I be told the results of the study?

You have a right to receive feedback about the overall results of this study.

You can tell us that you wish to receive feedback by ticking the relevant box on the consent form.

This feedback will be in the form of a one page summary.

This feedback will be given to participants at the end of study (2024)

(12) What if I have a complaint or any concerns about the study?

If there is a problem please let me know. You can contact me via the University of East Anglia at the following address:

Miss Alice Wheatley
School of Politics, Philosophy, Languages and Communication Studies
University of East Anglia
NORWICH NR4 7TJ
alice.wheatley@uea.ac.uk

If you are concerned about the way this study is being conducted or you wish to make a complaint to someone independent from the study, please contact the Head of School of Politics, Philosophy, Languages and Communication Studies:

Dr Kim Ridealgh
K.ridealgh@uea.ac.uk
01603 593408

(13) How do I know that this study has been approved to take place?

To protect your safety, rights, wellbeing and dignity, all research in the University of East Anglia is reviewed by a Research Ethics Body. This research was approved by the HUM S-REC (Faculty of Arts and Humanities Research Ethics Subcommittee).

(14) What is the general data protection information I need to be informed about?

According to data protection legislation, we are required to inform you that the legal basis for processing your data as listed in Article 6(1) of the UK GDPR is because this allows us to process personal data when it is necessary to perform our public tasks as a University.

In addition to the specific information provided above about why your personal data is required and how it will be used, there is also some general information which needs to be provided for you:

- The data controller is the University of East Anglia.
- For further information, you can contact the University's Data Protection Officer at dataprotection@uea.ac.uk
- You can also find out more about your data protection rights at the Information Commissioner's Office (ICO).
- If you are unhappy with how your personal data has been used, please contact the University's Data Protection Officer at dataprotection@uea.ac.uk in the first instance.

(15) OK, I want to take part – what do I do next?

You need to fill in one copy of the consent form and return to alice.wheatley@uea.ac.uk
Please keep the letter, information sheet and the second copy of the consent form for your information.

(16) Further information

This information was last updated on 25 November 2022.

If there are changes to the information provided, you will be notified by email.

This information sheet is for you to keep

Would you like to receive feedback about the overall results of this study?
YES NO

If you answered YES, please indicate your preferred form of feedback and address:

Postal: _____

Email: _____

.....
Signature

.....
PRINT name

.....
Date

Works Cited

- Adam Carter, J. and Peterson, M. (2015) 'On the Epistemology of the Precautionary Principle', *Erkenntnis*, 80(1), pp. 1–13. Available at: <https://doi.org/10.1007/s10670-014-9609-x>.
- Adams, R.M. (2006) *A Theory of Virtue: Excellence in Being for the Good*. Oxford: Oxford University Press.
- Almassi, B. (2012) 'Climate Change, Epistemic Trust, and Expert Trustworthiness', *Ethics and the Environment*, 17(2). Available at: <https://dx.doi.org/10.2979/ethicsenviro.17.2.29>.
- Alston, W.P. (1988) 'The Deontological Conception of Epistemic Justification', *Philosophical Perspectives*, 2, pp. 257–299. Available at: <https://doi.org/10.2307/2214077>.
- Alston, W.P. (1993) 'Epistemic Desiderata', *Philosophy and Phenomenological Research*, 53(3), pp. 527–551. Available at: <https://doi.org/10.2307/2108079>.
- Alston, W.P. (2005) *Beyond 'Justification': Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Alvarez, M. (2010) *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford: Oxford University Press.
- Anderson, E. (1995) 'Feminist Epistemology: An Interpretation and a Defense', *Hypatia*, 10(3), pp. 50–84. Available at: <https://doi.org/10.1111/j.1527-2001.1995.tb00737.x>.
- Anderson, E. (2011) 'Democracy, Public Policy, and Lay Assessments of Scientific Testimony', *Episteme*, 8(2), pp. 144–164. Available at: <https://doi.org/doi:10.3366/epi.2011.0013>.
- Andow, J. (2016) 'Qualitative Tools and Experimental Philosophy', *Philosophical Psychology*, 29(8), pp. 1128–1141. Available at: <https://doi.org/10.1080/09515089.2016.1224826>.
- Andow, J. (2023) 'Epistemic Deontology, Epistemic Trade-offs, and Kant's Formula of Humanity', *Synthese*, 201(2). Available at: <https://doi.org/10.1007/s11229-023-04043-3>.
- Anscombe, G.E.M. (1958) 'Modern Moral Philosophy', *Philosophy*, 33(124), pp. 1–19.
- Arpaly, N. (2011) 'Open-Mindedness as a Moral Virtue', *American Philosophical Quarterly*, 48(1), pp. 75–85.
- Arpaly, N. and Schroeder, T. (2014) *In Praise of Desire*. New York, NY: Oxford University Press.
- Association of IFCA (n.d.a) *About Us*. Available at: <https://association-ifca.org.uk/about-us/> (Accessed: 12 June 2024).

Association of IFCA (n.d.b) *Byelaws*. Available at: <https://association-ifca.org.uk/byelaws/> (Accessed: 9 June 2024).

Audi, R. (2013) 'Knowledge, Justification, and the Normativity of Epistemology', *Res Philosophica*, 90(2), pp. 127–145. Available at: <https://doi.org/10.11612/resphil.2013.90.2.2>.

Baehr, J. (2011) *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford: Oxford University Press.

Baier, A.C. (1990) 'Why Honesty Is a Hard Virtue', in O. Flanagan and A.O. Rorty (eds.) *Identity, Character, and Morality: Essays in Moral Psychology*. Cambridge, MA: The MIT Press.

Baron, M. (1988) 'What is Wrong with Self-Deception', in B. P. McLaughlin and A. O. Rorty (eds.) *Perspectives on Self-Deception*. Berkeley, CA: University of California Press.

Basu, R. (2021) 'A Tale of Two Doctrines: Moral Encroachment and Doxastic Wronging', in J. Lackey (ed.) *Applied Epistemology*. Oxford: Oxford University Press, pp. 99–118.

Battaly, H. (2014) 'Intellectual Virtues', in S. van Hooft (ed.) *The Handbook of Virtue Ethics*. 177th–187th edn. Oxon: Routledge.

Battaly, H. (2019) 'Vice Epistemology has a Responsibility Problem', *Philosophical Issues*, 29(1), pp. 24–36. Available at: <https://doi.org/10.1111/phis.12138>.

Battaly, H. (2024) 'Honesty Isn't Always a Virtue', *Analysis*, 84(2), pp. 414–424. Available at: <https://doi.org/10.1093/analys/anad080>.

Bengson, J. (2015) 'A Noetic Theory of Understanding and Intuition as Sense-Maker', *Inquiry*, 58(7–8), pp. 633–668. Available at: <https://doi.org/10.1080/0020174X.2015.1083473>.

Bengson, J. (2017) 'The Unity of Understanding', in S.R. Grimm (ed.) *Making Sense of the World: New Essays on the Philosophy of Understanding*. New York, NY: Oxford University Press.

Berger, P.L. and Luckmann, T. (2011) *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. New York, NY: Open Road Integrated Media.

Berker, S. (2013) 'Epistemic Teleology and the Separateness of Propositions', *The Philosophical Review*, 122(3), pp. 337–393. Available at: <https://doi.org/10.1215/00318108-2087645>.

Bevir, M. and Blakely, J. (2018) *Interpretive Social Science: An Anti-Naturalist Approach*. Oxford: Oxford University Press.

Bevir, M. and Rhodes, R.A. (2016) 'Interpretive Political Science: Mapping the Field', in R.A. Rhodes and M. Bevir (eds) *Routledge Handbook of Interpretive Political Science*. London: Routledge.

- Big Wave Productions (n.d) *Help Our Kelp*. Available at: <https://www.bigwavetv.com/productions/help-our-kelp/> (Accessed: 11 June 2024).
- Bird, A. (2014) ‘When is There a Group that Knows?’, in J. Lackey (ed.) *Essays in Collective Epistemology*. Oxford: Oxford University Press, pp. 42–63.
- Bloomfield, P. (2022) ‘Virtues are Excellences’, *Ratio*, 35(1), pp. 49–60. Available at: <https://doi.org/10.1111/rati.12323>.
- Blue Marine Foundation (2023) ‘A Story of Hope - Recovery of Sussex Underwater Forests’, 21 March. Available at: <https://www.bluemarinefoundation.com/2023/03/21/a-story-of-hope-recovery-of-sussex-underwater-forests> (Accessed: 6 June 2024).
- BonJour, L. (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Bortolotti, L. (2020) *The Epistemic Innocence of Irrational Beliefs*. Oxford: Oxford University Press.
- Bourget, D. (2017) ‘The Role of Consciousness in Grasping and Understanding’, *Philosophy and Phenomenological Research*, 95(2), pp. 285–318. Available at: <https://doi.org/10.1111/phpr.12208>.
- Bourget, D. (2025) ‘A Phenomenal Theory of Grasping and Understanding’, in A.I. Mărășoiu and M. Dumitru (eds) *Understanding and Conscious Experience: Philosophical and Scientific Perspectives*. New York, NY: Routledge, pp. 99–132.
- Brady, M. (2009) ‘Curiosity and the Value of Truth’, in A. Haddock, A. Millar, and D. Pritchard (eds) *Epistemic Value*. Oxford: Oxford University Press, pp. 265–283.
- Braun, V. and Clarke, V. (2022) *Thematic Analysis: A Practical Guide*. London: Sage.
- Bricker, A.M. (2025) ‘Knowledge as a (Non-factive) Mental State’, *Erkenntnis*, 90, pp. 263–284. Available at: <https://doi.org/10.1007/s10670-023-00697-0>.
- British Sociological Association (2017) *BSA Statement of Ethical Practice*. Available at: https://www.britisoc.co.uk/media/24310/bsa_statement_of_ethical_practice.pdf (Accessed: 25 March 2025).
- Brogaard, B. (2005) ‘I Know. Therefore, I Understand’, Unpublished work.
- Burge, T. (1993) ‘Content Preservation’, *Philosophical Review*, 102(4), pp. 457–488. Available at: <https://doi.org/10.2307/2185680>.
- Byerly, T.R. (2021) *Intellectual Dependability: A Virtue Theory of the Epistemic and Educational Ideal*. New York: Routledge.
- Cafaro, P. and Sandler, R.L. (eds) (2005) *Environmental Virtue Ethics*. Lanham, MD: Rowman & Littlefield Publishers.

- Cambridge Dictionary (2025) 'Honesty', *Cambridge Dictionary*. Cambridge University Press. Available at: <https://dictionary.cambridge.org/thesaurus/honesty> (Accessed: 29 May 2025).
- Carr, D. (2014) 'The Human and Educational Significance of Honesty as an Epistemic and Moral Virtue', *Educational Theory*, 64(1), pp. 1–14. Available at: <https://doi.org/10.1111/edth.12047>.
- Carroll, L. (1895) 'What the Tortoise said to Achilles', *Mind*, 4(14), pp. 278–280. Available at: <https://doi.org/10.1093/mind/IV.14.278>.
- Carter, J.A. (2024) 'Epistemic Normativity is Not Independent of our Goals', in B. Roeber, E. Sosa, M. Steup and J. Turri (eds.) *Contemporary Debates in Epistemology*, Third Edition. Hoboken, NJ: Wiley-Blackwell.
- Carter, J.A., Pritchard, D. and Shepherd, J. (2019) 'Knowledge-how, Understanding-why and Epistemic Luck: an Experimental Study', *Review of Philosophy and Psychology*, 10(4), pp. 701–734. Available at: <https://doi.org/10.1007/s13164-018-0429-9>.
- Cassam, Q. (2016) 'Vice Epistemology', *The Monist*, 99(2), pp. 159–180. Available at: <https://doi.org/10.1093/monist/onv034>.
- Chrisman, M. (2008) 'Ought to Believe', *The Journal of Philosophy*, 105(7), pp. 346–370. Available at: <https://doi.org/10.5840/jphil2008105736>.
- Clark, T., Foster, L., Sloan, L. and Bryman, A. (2021) *Bryman's Social Research Methods*. Sixth Edition. Oxford: Oxford University Press.
- Coady, D. (2019) 'Epistemology and Climate Change', in M. Fricker, P. J., Graham, D. Henderson and N. J. L. L. Pedersen (eds.) *The Routledge Handbook of Social Epistemology*. New York, NY: Routledge, pp. 466–473.
- Coady, D. and Chase, J. (2019) 'The Return of Applied Epistemology', in D. Coady and J. Chase (eds.) *The Routledge Handbook of Applied Epistemology*. London: Routledge.
- Code, L. (1984) 'Toward a "Responsibilist" Epistemology', *Philosophy and Phenomenological Research*, 45(1), pp. 29–50. Available at: <https://doi.org/10.2307/2107325>.
- Code, L. (1987) *Epistemic Responsibility*. Reprint, Albany: State University of New York Press, 2020.
- Cohen, L.J. (1995) *An Essay on Belief and Acceptance*. Oxford: Oxford University Press.
- Conee, E. and Feldman, R. (2004) *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Creswell, J.W. and Plano Clark, V.L. (2018) *Designing and Conducting Mixed Methods Research*. Third Edition. Thousand Oaks, CA: Sage.
- Croce, M. (2019) 'On What it Takes to be an Expert', *The Philosophical Quarterly*, 69(274), pp. 1–21. Available at: <https://doi.org/10.1093/pq/pqy044>.

- Dancy, J. (2000) *Practical Reality*. Oxford: Oxford University Press.
- Denzin, N.K. and Lincoln, Y.S. (2005) *The SAGE Handbook of Qualitative Research*. Third Edition. Thousand Oaks, CA: Sage.
- DePaul, M. (2001) 'Value Monism in Epistemology', in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. New York: Oxford University Press, pp. 170–183.
- Driver, J. (2003) 'The Conflation of Moral and Epistemic Virtue', *Metaphilosophy*, 34(3), pp. 367–383. Available at: <https://doi.org/10.1111/1467-9973.00279>.
- Druckman, J.N. and McGrath, M.C. (2019) 'The Evidence for Motivated Reasoning in Climate Change Preference Formation', *Nature Climate Change*, 9(2), pp. 111–119. Available at: <https://doi.org/10.1038/s41558-018-0360-1>.
- Dzwonkowska, D. (2025) *Virtue Ethics and the Environment*. Oxon: Routledge.
- Egeland, J. (2019) 'The Demon that Makes Us Go Mental: Mentalism Defended', *Philosophical Studies*, 176(1), pp. 3141–3158. Available at: <https://doi.org/10.1007/s11098-018-1167-7>.
- Egeland, J. (2020) 'The Problem with Trusting Unfamiliar Faculties: Accessibilism Defended', *Logos & Episteme*, 11(4), pp. 447–471. Available at: <https://doi.org/10.5840/logos-episteme202011434>.
- Elgin, C.Z. (2017) *True Enough*. Cambridge, MA: MIT Press.
- Fantl, J. and McGrath, M. (2010) *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Feldman, R. (2005) 'Epistemological Duties', in P.K. Moser (ed.) *The Oxford Handbook of Epistemology*. Oxford: Oxford University Press, pp. 362–384.
- Feldman, R. and Conee, E. (2001) 'Internalism Defended', *American Philosophical Quarterly*, 38(1), pp. 1–18.
- Flores, C. and Woodard, E. (2023) 'Epistemic Norms on Evidence-Gathering', *Philosophical Studies*, 180(9), pp. 2547–2571. Available at: <https://doi.org/10.1007/s11098-023-01978-8>.
- Foley, R. (1987) *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Fricker, M. (2007) *Epistemic Injustice*. Oxford: Oxford University Press.
- Friedman, J. (2018) 'Junk Beliefs and Interest-Driven Epistemology', *Philosophy and Phenomenological Research*, 97(3), pp. 568–583. Available at: <https://doi.org/10.1111/phpr.12381>.
- Friedman, J. (2020) 'The Epistemic and the Zetetic', *The Philosophical Review*, 129(4), pp. 501–536. Available at: <https://doi.org/10.1215/00318108-8540918>.

- Gardiner, G. (2012) 'Understanding, Integration, and Epistemic Value', *Acta Analytica*, 27(2), pp. 163–181. Available at: <https://doi.org/10.1007/s12136-012-0152-6>.
- Gardiner, S.M. (2006) 'A Core Precautionary Principle', *Journal of Political Philosophy*, 14(1), pp. 33–60. Available at: <https://doi.org/10.1111/j.1467-9760.2006.00237.x>.
- Geertz, C. (1973) *The Interpretation of Cultures: Selected Essays*. New York, NY: Basic Books.
- Gettier, E.L. (1963) 'Is Justified True Belief Knowledge?', *Analysis*, 23(6), pp. 121–123. Available at: <https://doi.org/10.1093/analys/23.6.121>.
- Gibbons, J. (2010) 'Things That Make Things Reasonable', *Philosophy and Phenomenological Research*, 81(2), pp. 335–361. Available at: <https://doi.org/10.1111/j.1933-1592.2010.00373.x>.
- Gibbons, J. (2013) *The Norm of Belief*. Oxford: Oxford University Press.
- Gilbert, M. (1987) 'Modelling Collective Belief', *Synthese*, 73(1), pp. 185–204. Available at: <https://doi.org/10.1007/BF00485446>.
- Goldman, A.I. (1976) 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy*, 73(20), pp. 771–791. Available at: <https://doi.org/10.2307/2025679>.
- Goldman, A.I. (1999) *Knowledge in a Social World*. Oxford: Oxford University Press.
- Goldman, A.I. (2001) 'Experts: Which Ones Should You Trust?', *Philosophy and Phenomenological Research*, 63(1), pp. 85–110. Available at: <https://doi.org/10.1111/j.1933-1592.2001.tb00093.x>.
- Goldman, A.I. (2009) 'Internalism, Externalism, and the Architecture of Justification', *The Journal of Philosophy*, 106(6), pp. 309–338. Available at: <https://doi.org/10.5840/jphil2009106611>.
- Goldman, A.I. (2018) 'Expertise', *Topoi*, 37(1), pp. 3–10. Available at: <https://doi.org/10.1007/s11245-016-9410-3>.
- Gordon, E.C. (2012) 'Is There Propositional Understanding?', *Logos & Episteme*, 3(2), pp. 181–192. Available at: <https://doi.org/10.5840/logos-episteme20123234>.
- Graham, P.A. (2010) 'In Defense of Objectivism about Moral Obligation', *Ethics*, 121(1), pp. 88–115. Available at: <https://doi.org/10.1086/656328>.
- Gray, L.M., Wong-Wylie, G., Rempel, G. R. and Cook, K. (2020) 'Expanding Qualitative Research Interviewing Strategies: Zoom Video Communications', *The Qualitative Report*, 25(5), pp. 1292–1301. Available at: <https://doi.org/10.46743/2160-3715/2020.4212>.
- Greco, J. (1994) 'Virtue Epistemology and the Relevant Sense of "Relevant Possibility"', *Southern Journal of Philosophy*, 32(1). Available at: <https://doi.org/10.1111/j.2041-6962.1994.tb00703.x>.

- Grimm, S.R. (2006) 'Is Understanding a Species of Knowledge?', *The British Journal for the Philosophy of Science*, 57(3), pp. 515–535.
- Grimm, S.R. (2008) 'Epistemic Goals and Epistemic Values', *Philosophy and Phenomenological Research*, 77(3), pp. 725–744. Available at: <https://doi.org/10.1111/j.1933-1592.2008.00217.x>.
- Grimm, S.R. (2010) 'The Goal of Explanation', *Studies in the History and Philosophy of Science*, 41(4), pp. 337–344. Available at: <https://doi.org/10.1016/j.shpsa.2010.10.006>.
- Grimm, S.R. (2011) 'What is Interesting?', *Logos & Episteme*, 2(4), pp. 515–542. Available at: <https://doi.org/10.5840/logos-episteme2011243>.
- Grimm, S.R. (2012) 'The Value of Understanding', *Philosophy Compass*, 7(2), pp. 103–117. Available at: <https://doi.org/10.1111/j.1747-9991.2011.00460.x>.
- Grimm, S.R. (2014) 'Understanding as Knowledge of Causes', in A. Fairweather (ed.) *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of Science*. Cham: Springer.
- Grimm, S.R. (2017) 'Understanding and Transparency', in S.R. Grimm, C. Baumberger, and S. Ammon (eds.) *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. New York, NY: Routledge, pp. 333–361.
- Guenin, L.M. (2005) 'Intellectual Honesty', *Synthese*, 145(2), pp. 177–232. Available at: <https://doi.org/10.1007/s11229-005-3746-3>.
- Guerrero, A.A. (2016) 'Living with Ignorance in a World of Experts', in R. Peels (ed.) *Perspectives on Ignorance from Moral and Social Philosophy*. New York, NY: Routledge, pp. 156–185.
- Guerrero, A.A. (2017) 'Intellectual Difficulty and Moral Responsibility', in P. Robichaud and J.W. Wieland (eds.) *Responsibility: The Epistemic Condition*. Oxford: Oxford University Press.
- Hammersley, M. and Gomm, R. (1997) 'Bias in Social Research', *Sociological Research Online*, 2(1), pp. 7–19. Available at: <https://doi.org/10.5153/sro.55>.
- Hangel, N. and ChoGlueck, C. (2023) 'On the Pursuitworthiness of Qualitative Methods in Empirical Philosophy of Science', *Studies in History and Philosophy of Science*, 98, pp. 29–39. Available at: <https://doi.org/10.1016/j.shpsa.2022.12.009>.
- Harman, G. (1999) *Reasoning, Meaning and Mind*. New York, NY: Oxford University Press.
- Help Our Kelp* (2019). Big Wave Productions.
- Hieronymi, P. (2005) 'The Wrong Kind of Reason', *The Journal of Philosophy*, 102(9), pp. 437–457.
- Hieronymi, P. (2014) 'Reflection and Responsibility', *Philosophy & Public Affairs*, 42(1), pp. 3–41. Available at: <https://doi.org/10.1111/papa.12024>.

- Hills, A. (2016) 'Understanding Why', *Nous*, 50(4), pp. 661–688. Available at: <https://doi.org/10.1111/nous.12092>.
- Hodgett, S. and Deneulin, S. (2009) 'On the Use of Narratives for Assessing Development Policy', *Public Administration*, 87(1), pp. 65–79. Available at: <https://doi.org/10.1111/j.1467-9299.2008.01742.x>.
- Hollebone, J. (1989) 'Seaweed On-Sea', *Worthing Herald*, 25 August, p. 1.
- Hopster, J. (2023) 'Climate Uncertainty, Real Possibilities and the Precautionary Principle', *Erkenntnis*, 88(6), pp. 2431–2447. Available at: <https://doi.org/10.1007/s10670-021-00461-2>.
- Hursthouse, R. (1999) *On Virtue Ethics*. Oxford: Oxford University Press.
- Hursthouse, R. (2006) 'Applying Virtue Ethics to Our Treatment of the Other Animals', in J. Welchman (ed.) *The Practice of Virtue: Classic and Contemporary Readings in Virtue Ethics*. Indianapolis, IN: Hackett Publishing Company, pp. 136–55.
- Hyman, J. (2010) 'The Road to Larissa', *Ratio*, 23(4), pp. 393–414. Available at: <https://doi.org/10.1111/j.1467-9329.2010.00475.x>.
- James, S.P. (2015) *Environmental Philosophy: An Introduction*. Oxford: Polity Press.
- James, W. (1987) 'The Will to Believe', in *The Will to Believe: And Other Essays in Popular Philosophy*. Reprint, New York: Cambridge University Press, 2014.
- Jamieson, D. (1998) 'Animal Liberation is an Environmental Ethic', *Environmental Values*, 7(1), pp. 41–57. Available at: <https://doi.org/10.3197/096327198129341465>.
- Jamieson, D. (2007) 'The Moral and Political Challenges of Climate Change', in S.C. Moser and L. Dilling (eds.) *Creating a Climate for Change: Communicating Climate Change and Facilitating Social Change*. Cambridge, UK: Cambridge University Press, pp. 475–482.
- Jamieson, D. (2014) *Reason in a Dark Time: Why the Struggle Against Climate Change Failed and What It Means for our Future*. New York, NY: Oxford University Press.
- Jenni, K. (2003) 'Vices of Inattention', *Journal of Applied Philosophy*, 20(3), pp. 279–295. Available at: <https://doi.org/10.1046/j.0264-3758.2003.00253.x>.
- John, S. (2018) 'Epistemic Trust and the Ethics of Science Communication: Against Transparency, Openness, Sincerity and Honesty', *Social Epistemology*, 32(2), pp. 75–87. Available at: <https://doi.org/10.1080/02691728.2017.1410864>.
- Johnson King, Z. (2022) 'Radical Internalism', *Philosophical Issues*, 32(1), pp. 46–64. Available at: <https://doi.org/10.1111/phis.12235>.
- Jones, K. (1999) 'Second-Hand Moral Knowledge', *The Journal of Philosophy*, 96(2), pp. 55–78. Available at: <https://doi.org/10.2307/2564672>.

- Jones, K. and Schroeter, F. (2012) 'Moral Expertise', *Analyse & Kritik*, 34(2), pp. 217–230. Available at: <https://doi.org/10.1515/auk-2012-0204>.
- Kawall, J. (2009) 'The Epistemic Demands of Environmental Virtue', *Journal of Agricultural and Environmental Ethics*, 23(1–2), pp. 109–128. Available at: <https://doi.org/10.1007/s10806-009-9183-4>.
- Kawall, J. (2018) 'Virtue Epistemology and the Environment', in *The Routledge Handbook of Virtue Epistemology*. New York, NY: Routledge, pp. 392–406.
- Kawall, J. (2021) 'Information and Virtue in the Anthropocene', *Ethics, Policy & Environment*, 24(1), pp. 1–15. Available at: <https://doi.org/10.1080/21550085.2021.1904533>.
- Keen, S., Lomeli-Rodriguez, M. and Joffe, H. (2022) 'From Challenge to Opportunity: Virtual Qualitative Research During COVID-19 and Beyond', *International Journal of Qualitative Methods*, 21, pp. 1–11. Available at: <https://doi.org/10.1177/16094069221105075>.
- Kelly, T. (2003) 'Epistemic Rationality as Instrumental Rationality: A Critique', *Philosophy and Phenomenological Research*, 66(3), pp. 612–640. Available at: <https://doi.org/10.1111/j.1933-1592.2003.tb00281.x>.
- Kelp, C. (2015) 'Understanding Phenomena', *Synthese*, 192(12), pp. 3799–3816. Available at: <https://doi.org/10.1007/s11229-014-0616-x>.
- Khalifa, K. (2017) *Understanding, Explanation, and Scientific Knowledge*. Cambridge, UK: Cambridge University Press.
- Khalifa, K. (2020) 'Understanding, Truth, and Epistemic Goals', *Philosophy of Science*, 87(5), pp. 944–956. Available at: <https://doi.org/10.1086/710545>.
- Kidd, I.J. (2020) 'Epistemic Corruption and Social Oppression', in I.J. Kidd, H. Battaly, and Q. Cassam (eds.) *Vice Epistemology: Theory and Practice*. New York, NY: Routledge.
- Kiesewetter, B. (2011) "'Ought" and the Perspective of the Agent', *Journal of Ethics and Social Philosophy*, 5(3), pp. 1–25.
- Kiesewetter, B. (2017) *The Normativity of Rationality*. Oxford: Oxford University Press.
- Kiesewetter, B. (2018) 'How Reasons are Sensitive to Available Evidence', in C. McHugh, J. Way, and D. Whiting (eds.) *Normativity: Epistemic and Practical*. Oxford: Oxford University Press, pp. 90–114.
- King, N.L. (2021) *The Excellent Mind: Intellectual Virtues for Everyday Life*. New York, NY: Oxford University Press.
- Knott, E., Rao, A. H., Summers, K. and Teeger, C. (2022) 'Interviews in the Social Sciences', *Nature Reviews Methods Primers*, 2(1), article no: 73. Available at: <https://doi.org/10.1038/s43586-022-00150-6>.

- Kolodny, N. (2005) 'Why Be Rational?', *Mind*, 114(455), pp. 509–563. Available at: <https://doi.org/10.1093/mind/fzi509>.
- Kornblith, H. (1993) 'Epistemic Normativity', *Synthese*, 94(3), pp. 357–376. Available at: <https://doi.org/10.1007/BF01064485>.
- Korsgaard, C.M. (1996) *The Sources of Normativity*. Cambridge, UK: Cambridge University Press.
- Kvanvig, J.L. (2003) *The Value of Knowledge and the Pursuit of Understanding*. New York, NY: Cambridge University Press.
- Lackey, J. (2021) *Applied Epistemology*. Oxford: Oxford University Press.
- Lawler, I. (2021) 'Scientific Understanding and Felicitous Legitimate Falsehoods', *Synthese*, 198(7), pp. 6859–6887. Available at: <https://doi.org/10.1007/s11229-019-02495-0>.
- Le Bihan, S. (2021) 'Partial Truth Versus Felicitous Falsehoods', *Synthese*, 198(6), pp. 5415–5436. Available at: <https://doi.org/10.1007/s11229-019-02413-4>.
- Lehrer, K. and Cohen, S. (1983) 'Justification, Truth, and Coherence', *Synthese*, 55(2), pp. 191–207. Available at: <https://doi.org/10.1007/BF00485068>.
- Levy, N. and Brownstein, M. (2021) 'Philosophy's Other Climate Problem?', *Journal of Social Philosophy*, 52(4). Available at: <https://doi.org/10.1111/josp.12396>.
- Littlejohn, C. (2012) *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.
- Littlejohn, C. (2017) 'How and Why Knowledge is First', in J.A. Carter, E.C. Gordon, and B. Jarvis (eds.) *Knowledge First: Approaches in Epistemology and Mind*. New York, NY: Oxford University Press, pp. 19–45.
- Lobe, B., Morgan, D. and Hoffman, K.A. (2020) 'Qualitative Data Collection in an Era of Social Distancing', *International Journal of Qualitative Methods*, 19, pp 1-8. Available at: <https://doi.org/10.1177/1609406920937875>.
- Logins, A. (2018) 'Subjective Unpossessed Reasons', *Thought: A Journal of Philosophy*, 7(4), pp. 262–270. Available at: <https://doi.org/10.1002/tht3.395>.
- Lord, E. (2018) *The Importance of Being Rational*. Oxford: Oxford University Press.
- Lorimer, J. (2007) 'Nonhuman Charisma', *Environment and Planning D: Society and Space*, 25(5), pp. 911–932. Available at: <https://doi.org/10.1068/d71j>.
- Lynch, M.P. (2005) *True to Life: Why Truth Matters*. Cambridge, MA: MIT Press.
- Marine Conservation Society (2023) 'Key Milestones: The Sussex IFCA Nearshore Trawling Byelaw', in *Sussex Kelp Recovery Project Progress and Impact Report*.

- Marine Conservation Society (n.d) *Safeguarding Nearshore Marine Habitats at Scale: Lessons and Recommendations from the Implementation of the Sussex Inshore Fisheries and Conservation Authority Nearshore Trawling Byelaw*. Available at: https://s3.eu-west-1.amazonaws.com/media.mcsuk.org/documents/Sussex_IFCA_Nearshore_Trawling_Byelaw_Report.pdf. (Accessed: 16 June 2025).
- Mason, E. (2015) ‘Moral Ignorance and Blameworthiness’, *Philosophical Studies*, 172(11), pp. 3037–3057. Available at: <https://doi.org/10.1007/s11098-015-0456-7>.
- McKenna, R. (2023) *Non-Ideal Epistemology*. Oxford: Oxford University Press.
- McNaughton Nicholls, C., Mills, L. and Kotecha, M. (2014) ‘Observation’, in J. Ritchie, J. Lewis, C. McNaughton Nicholls and R. Ormston (eds.) *Qualitative Research Practice: A Guide for Social Science Students and Researchers*. London: Sage.
- McShane, K. (2007) ‘Anthropocentrism vs. Nonanthropocentrism: Why Should We Care?’, *Environmental Values*, 16(2), pp. 169–185. Available at: <https://doi.org/10.3197/096327107780474555>.
- Mele, A.R. (2001) *Self-Deception Unmasked*. Princeton, NJ: Princeton University Press.
- Members of the Research Data Management Project Team (University of East Anglia) and Scott, R. (2025) ‘Research Data Management Policy’. Available at: <https://assets.uea.ac.uk/f/185899/x/8932506b17/research-data-management-policy.pdf> (Accessed: 16 June 2025).
- Michaelian, K. (2011) ‘The Epistemology of Forgetting’, *Erkenntnis*, 74(3), pp. 399–424. Available at: <https://doi.org/10.1007/s10670-010-9232-4>.
- Miller, C.B. (2021) *Honesty: The Philosophy and Psychology of a Neglected Virtue*. New York, NY: Oxford University Press.
- Miscevic, N. (2007) ‘Virtue-Based Epistemology and the Centrality of Truth (Towards a Strong Virtue-Epistemology)’, *Acta Analytica*, 22(3), pp. 239–266. Available at: <https://doi.org/10.1007/s12136-007-0011-z>.
- Miscevic, N. (2015) ‘Curiosity—The Basic Epistemic Virtue’, in *Moral and Intellectual Virtues in Western and Chinese Philosophy*. New York, NY: Routledge, pp. 155–173.
- Mizrahi, M. (2012) ‘Idealizations and Scientific Understanding’, *Philosophical Studies*, 160(2), pp. 237–252. Available at: <https://doi.org/10.1007/s11098-011-9716-3>.
- Moellendorf, D. (2022) *Mobilizing Hope: Climate Change and Global Poverty*. Oxford: Oxford University Press.
- Montmarquet, J. (1992) ‘Epistemic Virtue and Doxastic Responsibility’, *American Philosophical Quarterly*, 29(4), pp. 331–341.
- Myers, J. (2024) ‘The Epistemic Role of Vividness’, *Analysis*. Available at: <https://doi.org/10.1093/analys/anae002>.

- Nagel, J. (2017) 'Factive and Nonfactive Mental State Attribution', *Mind & Language*, 32(5), pp. 525–544. Available at: <https://doi.org/10.1111/mila.12157>.
- Nelkin, D. (2012) 'Responsibility and Self-Deception: A Framework', *Humana Mentis Journal of Philosophical Studies*, 5(20), pp. 117–139.
- Neuman, W.L. (2014) *Social Research Methods: Qualitative and Quantitative Approaches*. Seventh Edition. Harlow: Pearson.
- Norgaard, K.M. (2011) *Living in Denial: Climate Change, Emotions, and Everyday Life*. Cambridge, MA: MIT Press.
- Norton, B.G. (1991) *Toward Unity Among Environmentalists*. New York, NY: Oxford University Press.
- Oliffe, J. L., Kelly, M. T., Gonzalez Montaner, G., and Yu Ko, W. F., (2021) 'Zoom Interviews: Benefits and Concessions', *International Journal of Qualitative Methods*, 20. Available at: <https://doi.org/10.1177/16094069211053522>.
- O'Neill, J. (1992) 'The Varieties of Intrinsic Value', *The Monist*, 75(2), pp. 119–137. Available at: <https://doi.org/10.5840/monist19927527>.
- Oreskes, N. and Conway, E.M. (2010) *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. New York, NY: Bloomsbury Press.
- Our Sea Forest* (2023). Big Wave Productions.
- Parfit, D. (1984) *Reasons and Persons*. Oxford: Oxford University Press.
- Parfit, D. (2011) *On What Matters: Volume One*. Oxford: Oxford University Press.
- Pernecky, T. (2016) *Epistemology and Metaphysics for Qualitative Research*. London: Sage.
- Pritchard, D. (2008) 'Knowing the Answer, Understanding and Epistemic Value', *Grazer Philosophische Studien*, 77(1). Available at: <https://doi.org/10.1163/18756735-90000852>.
- Pritchard, D. (2009) 'Knowledge, Understanding and Epistemic Value', *Royal Institute of Philosophy Supplement*. 2009/05/27 edn, 64, pp. 19–43. Available at: <https://doi.org/10.1017/S1358246109000046>.
- Pritchard, D. (2010) 'Knowledge and Understanding', in *The Nature and Value of Knowledge: Three Investigations*. Pritchard, D., Millar, A. and Haddock, A. (eds.) Oxford: Oxford University Press.
- Pritchard, D. (2014) 'Truth as the Fundamental Epistemic Good', in J. Matheson and R. Vitz (eds.) *The Ethics of Belief*. Oxford: Oxford University Press.
- Pritchard, D. (2023) 'Truth as a Fundamental Value', in *Practices of Truth in Philosophy: Historical and Comparative Perspectives*. P. Gori and L. Serini (eds.) New York, NY: Routledge, pp. 8–23.

- Quast, C. (2018) 'Expertise: A Practical Explication', *Topoi*, 37(1), pp. 11–27. Available at: <https://doi.org/10.1007/s11245-016-9411-2>.
- Railton, P. (1994) 'Truth, Reason, and the Regulation of Belief', *Philosophical Issues*, 5, pp. 71–93. Available at: <https://doi.org/10.2307/1522874>.
- Regan, T. (1985) 'The Case for Animal Rights', in P. Singer (ed.) *In Defense of Animals*. New York, NY: Blackwell.
- de Regt, H.W. (2015) 'Scientific Understanding: Truth or Dare?', *Synthese*, 192(12), pp. 3781–3797. Available at: <https://doi.org/10.1007/s11229-014-0538-7>.
- Rice, C.C. (2016) 'Factive Scientific Understanding Without Accurate Representation', *Biology & Philosophy*, 31(1), pp. 81–102. Available at: <https://doi.org/10.1007/s10539-015-9510-2>.
- Riggs, W.D. (2003) 'Understanding "Virtue" and the Virtue of Understanding', in M. DePaul and L. Zagzebski (eds.) *Intellectual Virtue: Perspectives in Ethics and Epistemology*. Oxford: Clarendon.
- Riggs, W.D. (2016) 'Open-Mindedness, Insight, and Understanding', in J. Baehr (ed.) *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology*. New York, NY: Taylor & Francis.
- Ritchie, J. and Ormston, R. (2014) 'The Applications of Qualitative Methods to Social Research', in J. Ritchie, R. Ormston, C. McNaughton Nicholls and J. Lewis (eds.) *Qualitative Research Practice: a Guide for Social Science Students and Researchers*. London: Sage.
- Roberts, R.C. and West, R. (2020) 'The Virtue of Honesty: A Conceptual Exploration', in C.B. Miller and R. West (eds.) *Integrity, Honesty, and Truth Seeking*. Oxford University Press, pp. 97–126. Available at: <https://doi.org/10.1093/oso/9780190666026.003.0004>.
- Roberts, R.C. and Wood, W.J. (2007) *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Oxford University Press.
- Robichaud, P. (2017) 'Is Ignorance of Climate Change Culpable?', *Science and Engineering Ethics*, 23, pp. 1409–1430. Available at: <https://doi.org/10.1007/s11948-016-9835-5>.
- Rosen, G. (2004) 'Skepticism about Moral Responsibility', *Ethics*, 18(1), pp. 295–313. Available at: <https://doi.org/10.1111/j.1520-8583.2004.00030.x>.
- Ross, L. (2020) 'The Virtue of Curiosity', *Episteme*, 17(1), pp. 105–120. Available at: <https://doi.org/10.1017/epi.2018.31>.
- Ruslin, R., Mashuri, S., Rasak, M. S. A., Alhabsyi, F. and Syam, H. (2022) 'Semi-Structured Interview: A Methodological Reflection on the Development of a Qualitative Research Instrument in Educational Studies', *IOSR Journal of Research and Method in Education*, 12(1), pp. 22–29. Available at: <https://doi.org/10.9790/7388-1201052229>.

- Russell, B. (1948) *Human Knowledge: Its Scope and Limits*. New York, NY: Simon and Schuster.
- Sandler, R.L. (2007) *Character and Environment: a Virtue-Oriented Approach to Environmental Ethics*. New York, NY: Columbia University Press.
- Saul, J.M. (2012) *Lying, Misleading, and What Is Said : An Exploration in Philosophy of Language and in Ethics*. Oxford: Oxford University Press.
- Scanlon, T.M. (2014) *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Scholz, O.R. (2018) ‘Symptoms of Expertise: Knowledge, Understanding and Other Cognitive Goods’, *Topoi*, 37(1), pp. 29–37. Available at: <https://doi.org/10.1007/s11245-016-9429-5>.
- Schroeder, M. (2008) ‘Having Reasons’, *Philosophical Studies*, 139(1), pp. 57–71. Available at: <https://doi.org/10.1007/s11098-007-9102-3>.
- Schroeder, M.A. (2007) *Slaves of the Passions*. Oxford: Oxford University Press.
- Shah, N. and Velleman, J.D. (2005) ‘Doxastic Deliberation’, *The Philosophical Review*, 114(4), pp. 497–534. Available at: <https://doi.org/10.1215/00318108-114-4-497>.
- Silva, P. (2021) ‘Possessing Reasons: Why the Awareness-First Approach is Better than the Knowledge-First Approach’, *Synthese*, 199(1-2), pp. 2925–2947. Available at: <https://doi.org/10.1007/s11229-020-02916-5>.
- Singer, P. (1991) *Animal Liberation*. Second Edition. London: Thorsons.
- Sliwa, P. (2015) ‘IV—Understanding and Knowing’, *Proceedings of the Aristotelian Society*, 115(1), pp. 57–74. Available at: <https://doi.org/10.1111/j.1467-9264.2015.00384.x>.
- Smith, E. (n.d) *Our Story*. Available at: <https://sussexunderwater.uk/our-story/> (Accessed: 15 July 2025).
- Smithies, D. (2012) ‘Moore’s Paradox and the Accessibility of Justification’, *Philosophy and Phenomenological Research*, 85(2), pp. 273–300. Available at: <https://doi.org/10.1111/j.1933-1592.2011.00506.x>.
- Sosa, E. (1991) *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge, UK: Cambridge University Press.
- Sosa, E. (2003) ‘The Place of Truth in Epistemology’, in M. DePaul and L. Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press, pp. 155–179.
- Srinivasan, A. (2020) ‘Radical Externalism’, *The Philosophical Review*, 129(3), pp. 395–431. Available at: <https://doi.org/10.1215/00318108-8311261>.
- Stafford, S.P. (2010) ‘Intellectual Virtues in Environmental Virtue Ethics’, *Environmental Ethics*, 32(4). Available at: <https://doi.org/10.5840/enviroethics201032439>.

- Stake, R.E. (1995) *The Art of Case Study Research*. Thousand Oaks, CA: Sage Publications.
- Sterman, J.D. (2011) ‘Communicating Climate Change Risks in a Skeptical World’, *Climatic Change*, 108(4), pp. 811–826. Available at: <https://doi.org/10.1007/s10584-011-0189-3>.
- Steup, M. (2003) ‘A Defense of Internalism’, in L.P. Pojman (ed.) *A Theory of Knowledge: Classical and Contemporary Readings*. Belmont, CA: Wadsworth.
- Steup, M. (2008) ‘Doxastic Freedom’, *Synthese*, 161(3), pp. 375–392. Available at: <https://doi.org/10.1007/s11229-006-9090-4>.
- Strawson, P.F. (1962) ‘Freedom and Resentment’, *Proceedings of the British Academy*, 48, pp. 187–211.
- Strevens, M. (2017) ‘How Idealizations Provide Understanding’, in S. R. Grimm, C. Baumberger and S. Ammon *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. New York, NY: Routledge, pp. 76–84.
- Strevens, M. (2024) ‘Grasp and Scientific Understanding: A Recognition Account’, *Philosophical Studies*, 181(4), pp. 741–762. Available at: <https://doi.org/10.1007/s11098-024-02121-x>.
- Stroud, B. (1979) ‘Inference, Belief, and Understanding’, *Mind*, 88(350), pp. 179–196. Available at: <https://doi.org/10.1093/mind/LXXXVIII.1.179>.
- Sussex IFCA (2020) ‘Sussex IFCA Nearshore Trawling Byelaw 2019 Impact Assessment’. Available at: <https://secure.toolkitfiles.co.uk/clients/34087/sitedata/files/Byelaws/Nearshore-Trawling-IA.pdf>. (Accessed: 16 June 2025).
- Sussex Kelp Recovery Project (2024) ‘Progress and Impact Report 2023’.
- Sussex Kelp Recovery Project (n.d.a) ‘Glimmers of Hope for Recovery of Sussex Sea Forests’. Available at: <https://sussexkelp.org.uk/news/glimmers-of-hope-for-recovery-of-sussex-sea-forests> (Accessed: 12 June 2024).
- Sussex Kelp Recovery Project (n.d.b) *Our Aims*. Available at: <https://sussexkelp.org.uk/about-the-skrp/our-aims> (Accessed: 23 February 2024).
- Sylvan, K. (2015) ‘What Apparent Reasons Appear to Be’, *Philosophical Studies*, 172(3), pp. 587–606. Available at: <https://doi.org/10.1007/s11098-014-0320-1>.
- Sylvan (Routley), R. (2008) ‘Is There a Need for a New, an Environmental, Ethic?’, in R. Attfield (ed.) *The Ethics of the Environment*. London: Routledge.
- Tännsjö, T. (2013) *Understanding Ethics*. Edinburgh: Edinburgh University Press.
- The University of East Anglia's Research Integrity Manager, Research and Innovation Services and University Research Ethics Committee (UREC) (2025) ‘Research Ethics

Policy'. Available at: <https://assets.uea.ac.uk/f/185899/x/2fecbbc6cb/university-research-ethics-policy.pdf> (Accessed: 16 June 2025).

The University of East Anglia's Research Integrity Manager, Research and Innovation Services (RIN), Research Ethics Committee (UREC) and Research Culture Group (2025) 'Guidelines on Good Practice in Research'. Available at: <https://assets.uea.ac.uk/f/185899/x/6e506f694a/university-guidelines-on-good-practice-in-research.pdf>. (Accessed: 16 June 2025).

Thomson, J.J. (1990) *The Realm of Rights*. Cambridge, MA: Harvard University Press.

Treanor, B. (2014) *Emplotting Virtue: a Narrative Approach to Environmental Virtue Ethics*. Albany, NY: State University of New York Press.

Treanor, N. (2014) 'Trivial Truths and the Aim of Inquiry', *Philosophy and Phenomenological Research*, 89(3), pp. 552–559. Available at: <https://doi.org/10.1111/j.1933-1592.2012.00612.x>.

UK Government (2025) *Marine and Coastal Access Act 2009*, GOV.UK. Available at: <https://www.legislation.gov.uk/ukpga/2009/23/contents> (Accessed: 1 July 2025).

UK Government (n.d) *Natural Capital Committee (NCC)*, GOV.UK. Available at: <https://www.gov.uk/government/groups/natural-capital-committee> (Accessed: 13 April 2025).

Warenski, L. (2021) 'Epistemic Norms: Truth Conducive Enough', *Synthese*, 198(3), pp. 2721–2741. Available at: <https://doi.org/10.1007/s11229-019-02242-5>.

Way, J. (2009) 'Two Accounts of the Normativity of Rationality', *Journal of Ethics and Social Philosophy*, 4(1), pp. 1–9.

Wedgwood, R. (2002a) 'Internalism Explained', *Philosophy and Phenomenological Research*, 65(2), pp. 349–369. Available at: <https://doi.org/10.1111/j.1933-1592.2002.tb00206.x>.

Wedgwood, R. (2002b) 'The Aim of Belief', *Philosophical Perspectives*, 16, pp. 267–297. Available at: <https://doi.org/10.1111/1468-0068.36.s16.10>.

White, P.Q. (2022) 'Honesty and Discretion', *Philosophy & Public Affairs*, 50(1), pp. 6–49. Available at: <https://doi.org/10.1111/papa.12203>.

Whiting, D. (2010) 'Should I Believe the Truth?', *Dialectica*, 64(2), pp. 213–224. Available at: <https://doi.org/10.1111/j.1746-8361.2009.01204.x>.

Wilkenfeld, D.A. (2013) 'Understanding as Representation Manipulability', *Synthese*, 190(6), pp. 997–1016. Available at: <https://doi.org/10.1007/s11229-011-0055-x>.

Williams, B. (1979) 'Internal and External Reasons', in Ross. Harrison (ed.) *Rational Action: Studies in Philosophy and Social Science*. Cambridge, UK: Cambridge University Press.

Williams, B. (2006) *Ethics and the Limits of Philosophy*. London: Routledge.

- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williston, B. (2012) 'Climate Change and Radical Hope', *Ethics & the Environment*, 17(2), pp. 165–186. Available at: <https://doi.org/10.2979/ethicsenviro.17.2.165>.
- Wilson, A.T. (2017) 'Avoiding the Conflation of Moral and Intellectual Virtues', *Ethical Theory and Moral Practice*, 20(5), pp. 1037–1050. Available at: <https://doi.org/10.1007/s10677-017-9843-9>.
- Wilson, A.T. (2018) 'Honesty as a Virtue', *Metaphilosophy*, 49(3), pp. 262–280. Available at: <https://doi.org/10.1111/meta.12303>.
- Wolf, S. (2011) 'Blame, Italian Style', in R.J. Wallace, R. Kumar, and S. Freeman (eds.) *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. Oxford: Oxford University Press.
- Woodward, J. (2005) *Making Things Happen: A Theory of Causal Explanation*. New York, NY: Oxford University Press.
- Yanow, D. (2007) 'Interpretation in Policy Analysis: On Methods and Practice', *Critical Policy Studies*, 1(1), pp. 110–122. Available at: <https://doi.org/10.1080/19460171.2007.9518511>.
- Yanow, D. (2015) 'Thinking Interpretively: Philosophical Presuppositions and the Human Sciences', in D. Yanow and P. Schwartz-Shea (eds.) *Interpretation and Method: Empirical Research Methods and the Interpretive Turn*. Second Edition. Abingdon, Oxon: Routledge, pp. 5–26.
- Yeo, A., Legard, R., Keegan, J., Ward, K. McNaughton, Nicholls, C., Lewis, J. (2014) 'In-Depth Interviews', in J. Ritchie, J. Lewis, C. McNaughton Nicholls, R. Ormston (eds.) *Qualitative Research Practice: A Guide for Social Science Students and Researchers*. Second Edition. London: Sage.
- Yin, R.K. (2009) *Case Study Research: Design and Methods*. Fourth Edition. Thousand Oaks, CA: Sage.
- Yin, R.K. (2016) *Qualitative Research from Start to Finish*. Second Edition. New York, NY: The Guilford Press.
- Zagzebski, L. (2001) 'Recovering Understanding', in M. Steup (ed.) *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. New York, NY: Oxford University Press, pp. 235–252.
- Zagzebski, L. (2003) 'The Search for the Source of Epistemic Good', *Metaphilosophy*, 34(1–2), pp. 12–28. Available at: <https://doi.org/10.1111/1467-9973.00257>.
- Zagzebski, L. (1996) *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge, UK: Cambridge University Press.
- Zimmerman, M.J. (2008) *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge, UK: Cambridge University Press.

